HANKs fr th Mmrs (Even If They Weren't So Great)

The Distributional Cost of Recessions in a Heterogeneous Agent New Keynesian Model with Search Frictions and Human Capital

Joseph Kopecky*

March 2026

Abstract

I build a two-asset Heterogeneous Agent New Keynesian model with search-and-matching unemployment and endogenous human capital dynamics to study *who* bears the cost of recessions. Human capital depreciates during unemployment, and workers with low human capital face both lower re-employment rates and lower re-employment wages, generating a self-reinforcing scarring trap. When human capital depreciation intensifies during severe downturns, this produces persistence and a sharp distributional divide. In a Great Recession-calibrated shock (3pp separation rate increase), hand-to-mouth (HtM) households suffer consumption-equivalent (CE) welfare losses of -4.1%, nearly 60% larger than high-liquid households' -2.6%, with a distributional gradient (HtM minus high-liquid CE) of -1.5pp. This gradient is *convex* in shock size: it is approximately five times the average-recession (1pp) value of -0.29pp, suggesting strong non-linearities in the distributional costs of labor market scarring. An empirically scaled training subsidy (peak fiscal cost $\sim$3.6% of steady-state output at the worst quarter) nearly eliminates the distributional gradient at Great Recession severity.

JEL classification codes: E24, E32, E52, J24, J64

Keywords: HANK, distributional incidence, scarring, human capital, search and matching, training subsidies, inequality

*Trinity College Dublin (jkopecky@tcd.ie)

Say a prayer, but let the good times roll. In case God doesn't show, let the good times roll, let the good times roll...

— "Thnks fr th Mmrs," Fall Out Boy

1. Introduction

Recessions leave scars, but not equally. Workers displaced during downturns suffer earnings losses of 10–25% that persist for a decade or more (Jacobson, LaLonde, and Sullivan, 1993; Davis and von Wachter, 2011; Huckfeldt, 2022). These losses are sharply heterogeneous: low-earners, the less educated, and older workers bear the heaviest burden (Yagan, 2019; Oreopoulos, von Wachter, and Heisz, 2012; Fallick and Krolikowski, 2018). The "memories" that recessions leave on individual labor market histories, accumulated skill loss, degraded match quality, and weakened job-finding prospects, plausibly depend on what resources workers bring into the downturn. Yet existing models of macroeconomic scarring abstract from precisely the wealth heterogeneity that determines who is most vulnerable.

The Heterogeneous Agent New Keynesian (HANK) literature (Kaplan, Moll, and Violante, 2018; Auclert, Rognlie, and Straub, 2025) provides rich models of how household heterogeneity shapes the transmission of aggregate shocks and policy. But unemployment in HANK is typically a temporary state with no lasting consequences; there is no "memory" of the recession in individual careers. The hysteresis literature (Blanchard and Summers, 1986; Galí, 2022; Acharya, Bengui, Dogra, and Wee, 2022) models how temporary shocks become permanent, but does so in representative-agent frameworks that cannot speak to the distributional incidence of scarring. The emerging work combining the two (Alves and Violante, 2024, 2025; Acabbi, Alati, and Mazzone, 2024) has focused on whether scarring amplifies *aggregate* dynamics.

I build a HANK model with Diamond-Mortensen-Pissarides (DMP) search-and-matching unemployment, endogenous human capital dynamics, and match quality heterogeneity, and use it to instead study the *distributional incidence* of recessions. The key mechanism is a "scarring trap": human capital depreciates during unemployment, and workers with low human capital face lower job-finding rates *and* lower re-employment wages. When combined with incomplete markets, the model predicts that wealth-poor workers are most vulnerable; they cannot self-insure, spend longer unemployed, and lose more human capital. An extension incorporating match quality (Appendix C) confirms that match quality resets upon re-employment reinforce this channel, though the aggregate effects are small. I show that even though the aggregate implications are modest, there are important distributional consequences of labor market scarring. Further, I show that monetary policy does little to

offset these distributional effects, while targeted training subsidies can fully eliminate them.

The model is solved using the Sequence-Space Jacobian (SSJ) method of Auclert, Bardoczy, Rognlie, and Straub (2021). I use this framework to address a number of questions. First, I show that human downgrading in recessions has modest aggregate effects, with only small amplification of recessions relative to a conventional HANK model with unemployment. I then ask the degree to which these small aggregate effects mask important distributional movements: how are the welfare costs of recessions distributed across the wealth and human capital distribution? Which workers bear the largest and most persistent scars, and what determines vulnerability: liquid wealth or human capital? These turn out to be potentially large, with non-linear amplification in shock size. Finally, I ask how policy can reshape the distribution of recession costs, asking how do training subsidies, which directly offset accelerated human capital depreciation during recessions, reshape the distribution of recession costs? Can they substitute for the liquid wealth buffer that the most vulnerable workers lack?

The model features a two-asset structure with liquid bonds and illiquid savings subject to portfolio adjustment costs. This generates "wealthy hand-to-mouth" households, agents with positive illiquid wealth but near-zero liquid assets, who face binding liquidity constraints despite having positive net worth. Approximately 27% of households are hand-to-mouth in the calibrated steady state, consistent with the 24–31% range from SCF data (Kaplan, Violante, and Weidner, 2014). These households are especially vulnerable to the scarring mechanism: their inability to self-insure with liquid assets during unemployment amplifies the consequences of human capital depreciation. The HtM share is higher among the unemployed (35.2%) than the employed (26.7%), providing the right ingredient for heterogeneous scarring vulnerability.

Three sets of results emerge. First, the distributional costs of scarring are *convex* in shock size. In an average recession (1pp separation shock), state-dependent scarring ($\kappa = 2$) nearly doubles aggregate CE welfare costs to -1.0% and generates a moderate distributional gradient of -0.3pp (HtM minus high-liquid CE). In a Great Recession-calibrated shock (3pp), aggregate costs reach -3.2% and the gradient widens to -1.49pp, approximately five times the average-recession value ($5.1\times$). A shock-size sweep over $[0.5, 4]$pp confirms this convexity: the distributional divide accelerates nonlinearly as the scarring trap binds more tightly at larger shocks. Second, the distributional incidence at Great Recession severity is stark. HtM households suffer CE losses of -4.1%, nearly 60% larger than high-liquid households' -2.6%, driven by the interaction of binding liquidity constraints, elevated unemployment, and the wage-human capital penalty. Third, a training subsidy calibrated to observed active labor market policy (ALMP) spending (peak fiscal cost $\sim$3.6% of steady-

state output at the worst quarter) nearly eliminates the distributional gradient at 3pp, reducing it from -1.5pp to $+0.1$pp by preserving human capital for the most vulnerable workers. Monetary policy has limited leverage over labor market shocks, motivating direct intervention.

Section 2 discusses the related literature. Section 3 presents the model, including the wage-human capital channel, state-dependent scarring, and training subsidy mechanisms. Section 4 describes the calibration, including the human capital parameter calibration targeting the GE-feasible earnings loss profile and the recession-expansion displacement differential. Section 5 characterizes the steady state. Section 6 presents aggregate and distributional dynamics. Section 7 contains the policy experiments: a monetary policy evaluation followed by training subsidies as the headline analysis. Section 8 concludes.

2. Related Literature

This paper engages three literatures: HANK models with labor market frictions, macroeconomic hysteresis, and the emerging intersection of the two.

2.1. HANK Models with Search and Matching

The HANK paradigm, originating with Kaplan et al. (2018), embeds incomplete-markets households into New Keynesian models with sticky prices. The key insight is that monetary policy transmission operates primarily through *indirect* general-equilibrium channels (labor income, profits) rather than direct intertemporal substitution, because a large fraction of households are "hand-to-mouth" with high marginal propensities to consume (Kaplan et al., 2014). Bilbiie (2020, 2024) provide analytical foundations, and Auclert, Rognlie, and Straub (2024) develop the Intertemporal Keynesian Cross as a general framework for aggregate demand in HANK. The broader incomplete-markets tradition, surveyed by Heathcote, Storesletten, and Violante (2009) and Krueger, Mitman, and Perri (2016), provides the micro foundations for consumption insurance that these models build on; Heathcote, Storesletten, and Violante (2014) develop an analytical framework for partial insurance that is closely related to the consumption-smoothing mechanisms at work here.

A growing literature combines HANK with DMP search-and-matching frictions ("HANK & SAM"). Ravn and Sterk (2021) show analytically that unemployment risk and precautionary saving create powerful amplification: when unemployment rises, even employed workers increase saving, depressing demand (the "unemployment fears" channel). Kekre (2023) studies unemployment insurance (UI) in HANK+SAM, finding large stabilization benefits; see also Lentz (2009) on optimal UI design with human capital depreciation in

a job search framework. Graves (2025) builds a quantitative two-asset model with endogenous unemployment risk. Broer, Druedahl, Harmenberg, and Öberg (2025) find that endogenous separations approximately double the unemployment-risk channel. Bayer, Born, and Luetticke (2024) estimate a medium-scale HANK model and find that shocks, frictions, and inequality interact to shape business cycle dynamics. Gornemann, Kuester, and Nakajima (2021) show contractionary monetary policy disproportionately hurts low-wealth workers through job losses. Consolo and Hansel (2024) combine two-asset HANK with DMP frictions and briefly examine hysteresis after energy shocks. In a complementary vein, Guerrieri and Lorenzoni (2017) study how credit crises and precautionary savings generate liquidity traps, providing a balance-sheet channel for persistent recessions that operates alongside the human capital channel I emphasize.

My contribution relative to this literature is adding persistent consequences of job loss, individual labor market "memories," and studying how the resulting scarring is distributed across the wealth distribution. In existing DMP-based HANK+SAM models, re-employed workers return to their prior productivity; there is no memory of the recession in individual careers, and hence no role for wealth in determining the long-run cost of displacement.[1]

2.2. Macroeconomic Hysteresis

Blanchard and Summers (1986) introduced modern hysteresis through insider-outsider dynamics. Ljungqvist and Sargent (1998) show that turbulence in human capital loss combined with generous unemployment benefits can generate persistent unemployment traps, a mechanism closely related to my scarring trap, though their model lacks the wealth heterogeneity that determines who falls into the trap. Galí (2022) extends hysteresis in a NK model with an endogenous natural rate, showing optimal monetary policy calls for strong employment stabilization. Acharya et al. (2022) build a model where skill loss during unemployment generates unemployment traps at the zero lower bound. Garga and Singh (2021) study hysteresis through endogenous growth in a NK framework. Fornaro and Wolf (2023) analyze supply-side scarring from monetary policy.

The empirical evidence on hysteresis is now extensive. Cerra, Fatas, and Saxena (2023) provide a comprehensive survey documenting persistent output losses from recessions across 190+ countries. At the individual level, Davis and von Wachter (2011) document earnings losses of 19% upon displacement, still 10% after twenty years. Yagan (2019) shows Great Recession exposure permanently reduced employment. The structural mechanisms

[1]Some recent work explores aggregate persistence/hysteresis in HANK settings; e.g., Consolo and Hansel (2024) briefly examine insider-outsider dynamics, but these do not model worker-level post-displacement productivity states conditional on re-employment, which is the relevant object for my estimated distributional incidence.

underlying individual scarring have been studied by Jarosch (2023) (job ladder cascades), Huckfeldt (2022) (human capital and occupational downgrading), Krolikowski (2017) (match-specific capital), Pissarides (1992) (skill depreciation), and Baley, Figueiredo, and Ulbricht (2022) (cyclical mismatch dynamics).

My contribution relative to this literature is embedding these mechanisms in a heterogeneous-agent framework where the severity and persistence of scarring depends on household wealth and self-insurance capacity, and using the resulting model to study the distributional incidence of recession costs.

2.3. HANK Meets Hysteresis

The intersection of HANK and hysteresis is thin but growing. Alves and Violante (2024) build the closest existing model: a HANK with three labor market states where displacement causes lasting earnings losses through calibrated skill depreciation. Alves and Violante (2025) extend this to evaluate the Fed's lower-for-longer strategy. Acabbi et al. (2024) document that recessions distort worker-firm sorting and impair human capital accumulation, accounting for ∼30% of cumulative output losses; their model lacks nominal rigidities but provides important empirical targets. Consolo and Hansel (2024) combine two-asset HANK with DMP frictions and briefly examine hysteresis after energy shocks.

A common thread in this emerging literature is the pursuit of *aggregate* amplification: does scarring make recessions deeper or longer? My model finds that the answer is "modestly at best": adding human capital dynamics and match quality changes aggregate impulse responses by less than 5%. This is consistent with the relatively small aggregate effects in Alves and Violante (2024).

While aggregate implications are a first-order concern, these models also reveal important insights for the *distributional incidence* of scarring. I differ from Alves and Violante (2024, 2025) in three ways. First, scarring in my model arises through the interaction of human capital depreciation, h-dependent job-finding rates, h-dependent re-employment wages, and match quality heterogeneity, rather than through calibrated exogenous skill loss. The joint operation of the finding rate and wage channels is calibrated to match the empirical earnings loss profile of displaced workers (Davis and von Wachter, 2011; Jacobson et al., 1993). Second, the two-asset structure generates wealthy hand-to-mouth households whose scarring vulnerability is driven by *liquid wealth*: the ability to maintain human capital and draw good matches upon re-employment depends on liquid asset holdings, not net worth. Third, I focus on the distributional incidence of recession costs, who bears the burden, and how policy can reshape this distribution, rather than on aggregate amplification. I also introduce training subsidies as a novel policy instrument that directly targets the micro-

founded scarring mechanism. To my knowledge, no existing HANK model with labor market frictions studies state-contingent human capital subsidies or their distributional incidence across the wealth distribution.

3. Model

The economy consists of a continuum of households, a frictional labor market with search-and-matching, monopolistically competitive firms with sticky prices, and a government that conducts monetary and fiscal policy. Time is discrete, with one period corresponding to one quarter.

3.1. Households

A continuum of infinitely-lived households is heterogeneous in liquid assets b, illiquid assets a, human capital h, within-employment productivity e, and labor market status $s \in \{E, U\}$ (employed or unemployed).

Preferences. Households maximize expected lifetime utility:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \, u(c_t), \qquad u(c) = \log c, \tag{1}$$

where β is the discount factor and I use log utility (EIS $= 1$).

Budget constraint. Households hold two assets: liquid bonds b and illiquid savings a. Adjusting the illiquid portfolio incurs a convex cost $\Psi(a', a)$. Employed households at human capital level h and within-employment productivity e earn a wage that depends on both the aggregate wage w and their individual human capital, while unemployed households receive UI benefits based on e, their match-specific productivity:

$$c + b' + a' + \Psi(a', a) = (1 + r^b)\, b + (1 + r^a)\, a + y(s, h, e) \tag{2}$$

$$y(s, h, e) = \begin{cases} w(h) \cdot e \cdot h - T & \text{if } s = E \\ b_{ui} \cdot e - T & \text{if } s = U \end{cases} \tag{3}$$

where the individual effective wage is:

$$w(h) = w \cdot \left(\frac{h}{h_{\text{ref}}} \right)^{\phi_w}, \tag{4}$$

r^b is the liquid return (equal to the policy rate r), $r^a = r + \xi$ is the illiquid return (with illiquidity premium ξ), T is a lump-sum tax, w is the aggregate wage rate from Nash bargaining, b_{ui} is the UI benefit level, and h_{ref} is the mean human capital of employed workers in steady state (defined in Section 3.2 below). The parameter $\phi_w \geq 0$ governs the elasticity of individual wages with respect to human capital. When $\phi_w = 0$, all workers earn the same base wage (the baseline specification); when $\phi_w > 0$, workers with depreciated human capital earn lower wages upon re-employment, generating a *wage scarring* channel alongside the job-finding rate channel. This is motivated by the empirical finding that displacement causes persistent wage losses, not just longer unemployment spells (Davis and von Wachter, 2011; Jacobson et al., 1993). The variable h enters the employed income expression twice: once through the effective wage $w(h)$ and once as a productivity multiplier, so the total income penalty from low h is amplified. Liquid assets are subject to a borrowing constraint $b \geq 0$, and illiquid assets satisfy $a \geq 0$.

Adjustment costs. The portfolio adjustment cost follows Kaplan et al. (2018):

$$\Psi(a', a) = \frac{\chi_1}{\chi_2} \left| a' - (1 + r^a)a \right| \cdot \left(\frac{|a' - (1 + r^a)a|}{(1 + r^a)a + \chi_0} \right)^{\chi_2 - 1}, \tag{5}$$

where χ_0 shifts the cost at low wealth, χ_1 scales the overall cost level, and χ_2 controls curvature. This cost makes illiquid assets difficult to access quickly, generating "wealthy hand-to-mouth" households who hold positive illiquid wealth but near-zero liquid assets.

Optimality. The household problem is solved via a two-asset endogenous grid method (EGM) adapted from Auclert et al. (2021). The optimality conditions for the two assets are:

$$u'(c_t) \geq \beta(1 + r^b_{t+1}) \, \mathbb{E}_t \left[u'(c_{t+1}) \right], \tag{6}$$

$$(1 + \Psi_{1,t}) \, u'(c_t) \geq \beta \, \mathbb{E}_t \left[(1 + r^a_{t+1} - \Psi_{2,t+1}) \, u'(c_{t+1}) \right], \tag{7}$$

where $\Psi_{1,t} = \partial \Psi(a_{t+1}, a_t) / \partial a_{t+1}$ is the marginal cost of saving evaluated at time-t choices (known at t), and $\Psi_{2,t+1} = \partial \Psi(a_{t+2}, a_{t+1}) / \partial a_{t+1}$ is the marginal reduction in next period's adjustment cost from higher beginning-of-period illiquid wealth (not known at t, hence inside the expectation), with equality when the respective borrowing constraints do not bind. The expectation is taken over the joint transition of employment status, human capital, and within-employment productivity. When the liquid constraint binds ($b' = 0$) but the illiquid does not, only illiquid savings adjust (see Section 3.7 for details).

3.2. Human Capital Dynamics

Human capital $h \in [h_{\min}, h_{\max}]$ evolves deterministically conditional on employment status:

$$h' = \begin{cases} h + \phi_0 \left(h_{\max} - h\right)^{\phi_2} & \text{if employed (learning-by-doing)} \\ \max\left(h_{\min},\, h \cdot (1 - \delta_h)\right) & \text{if unemployed (depreciation)} \end{cases} \tag{8}$$

where ϕ_0 governs the rate of on-the-job learning, ϕ_2 controls curvature (diminishing returns to high-h workers), and δ_h is the quarterly depreciation rate during unemployment. Human capital is discretized to n_H grid points (see Section 3.7 for the discretization procedure).

3.3. Labor Market

The labor market features DMP search-and-matching with a Cobb-Douglas matching function:

$$M = \chi_m \cdot U^{\alpha_m} \cdot V^{1-\alpha_m}, \tag{9}$$

where U is the measure of unemployed workers, V is vacancies, α_m is the matching elasticity, and χ_m is matching efficiency. Defining market tightness $\theta = V/U$, the job-finding rate and vacancy-filling rate are:

$$f(\theta) = \chi_m\, \theta^{1-\alpha_m}, \qquad q(\theta) = \chi_m\, \theta^{-\alpha_m}. \tag{10}$$

Scarring through h-dependent job finding. The job-finding rate depends on human capital:

$$f(h) = f \cdot \left(\frac{h}{h_{\text{ref}}}\right)^{\zeta}, \tag{11}$$

where $f \equiv f(\theta_t)$ is the aggregate job-finding rate from (10), h_{ref} is a reference human capital level, and $\zeta \geq 0$ is the scarring elasticity. When $\zeta = 0$, the model nests a standard HANK+SAM without scarring. When $\zeta > 0$, unemployed workers with low human capital face lower finding rates, spending more time unemployed and losing more human capital, generating a self-reinforcing scarring trap.

The reference level h_{ref} is set to the mean human capital of employed workers in steady state ($h_{\text{ref}} \approx 1.8$), so that at $h = h_{\text{ref}}$ the individual finding rate equals the aggregate rate: $f(h_{\text{ref}}) = f(\theta)$. Aggregate matches are accounted for via $M_t = \int f_t(h)\, d\Gamma_{U,t}(h)$, where $\Gamma_{U,t}$ is the measure of unemployed workers over h. The matching function (9) determines $f(\theta_t)$ as a baseline hazard used on the firm side; $f(h)$ introduces a reduced-form heterogeneity wedge on the worker side, with the $(h/h_{\text{ref}})^{\zeta}$ factor redistributing finding rates across the

unemployed pool. This formulation is standard in heterogeneous-worker search models (e.g., Lise and Robin, 2017).

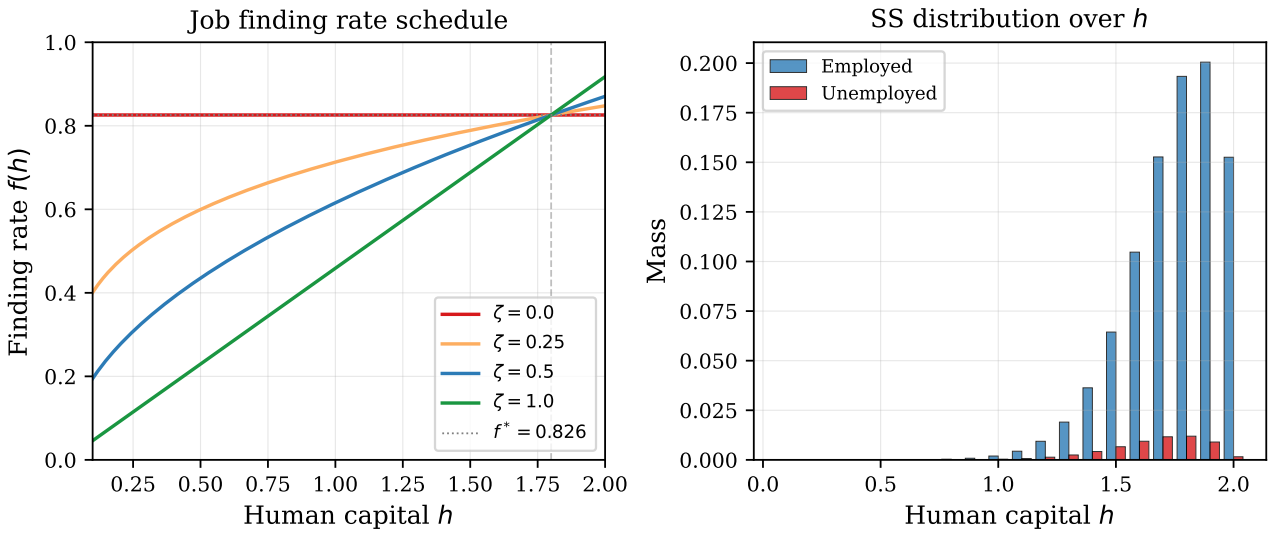

Figure 1: *Left: Job-finding rate schedule $f(h) = f^*(h/h_{ref})^{\zeta}$ for different values of the scarring elasticity ζ. When $\zeta = 0$, the finding rate is constant (no scarring). Higher ζ lowers the finding rate for workers with depreciated human capital. Right: Steady-state distribution over human capital for employed and unemployed workers. Most employed mass is concentrated at high h (near $h_{\max}$), while unemployed workers are spread across lower h values.*

Separations. Each period, employed workers are separated with exogenous probability s. Upon separation, the worker enters unemployment and human capital begins depreciating.

Wages. The aggregate wage is determined by Nash bargaining between the worker and firm:

$$w = \nu \left(mc \cdot Z + \kappa_v \, \theta \right) + (1 - \nu) \, b_{ui}, \tag{12}$$

where ν is the worker's bargaining power, mc is real marginal cost, Z is aggregate productivity, and κ_v is the vacancy posting cost. The worker's outside option is the UI flow value b_{ui}. The aggregate wage w does not depend on individual assets; this simplification ensures tractability and is common in the HANK+SAM literature. However, the *effective* wage $w(h)$ that enters individual income does depend on human capital through Eq. (4), so that workers with depreciated h earn lower wages upon re-employment.

Free entry. Firms post vacancies until the expected value of a filled vacancy equals the posting cost:

$$\frac{\kappa_v}{q(\theta_t)} = (mc_t \cdot Z_t - w_t) + \beta_{\text{firm}}(1 - s)\frac{\kappa_v}{q(\theta_{t+1})}, \tag{13}$$

where β_{firm} is the firm's discount factor and the flow profit is $mc \cdot Z - w$. The flow profit is positive because retailers operate under monopolistic competition: $mc_t < 1$ is the real marginal cost (price of intermediate goods relative to the final good price), and the competitive labor-services price $mc_t \cdot Z_t$ exceeds the Nash-bargained wage w_t by the DMP surplus share. In the frictionless limit ($\kappa_v \to 0$, $s \to 0$), $mc \cdot Z \to w$ and the match surplus vanishes.

Firm-side insulation from searcher composition. The free entry condition (13) and the Nash wage (12) are both aggregate objects: neither depends on the human capital of the searching pool. Individual human capital enters household income through $w(h) \cdot e \cdot h$ (Eq. (4)), but firms create vacancies against the aggregate surplus. This is standard in the HANK+SAM literature (Ravn and Sterk, 2021; McKay and Reis, 2021) and is driven by tractability, since making surplus depend on searcher h would require heterogeneous vacancy posting, endogenous separation rates, and a fundamentally different solution method. It also makes the distributional results conservative: if low-h matches were less profitable, vacancy creation would fall as the search pool deteriorates during recessions, disproportionately harming low-h (and therefore low-liquid) workers through deeper unemployment traps and slower vacancy recovery. The distributional gradient documented in Section 6.2 is therefore likely a lower bound on the true incidence of scarring through the vacancy-creation channel.

3.4. Firms and Pricing

Production. Aggregate output is produced using effective labor:

$$Y_t = Z_t \cdot H_{\text{eff},t}, \tag{14}$$

where $H_{\text{eff}} = \int h \cdot \mathbf{1}\{s = E\}\, d\Gamma$ is aggregate effective labor, computed from the cross-sectional distribution Γ over all household states, with the indicator selecting employed workers weighted by their human capital.

New Keynesian Phillips Curve. Monopolistically competitive retailers face Rotemberg quadratic price adjustment costs, generating the NKPC:

$$\log(1 + \pi_t) = \kappa_p \left(mc_t - \frac{1}{\mu} \right) + \frac{Y_{t+1}}{Y_t} \cdot \frac{\log(1 + \pi_{t+1})}{1 + r_{t+1}}, \tag{15}$$

where π is inflation, $\mu = \epsilon/(\epsilon - 1)$ is the gross markup, ϵ is the elasticity of substitution, and θ_p is the Rotemberg adjustment cost parameter. The forward-looking term uses the real rate $1/(1 + r_{t+1})$ as the stochastic discount factor, which equals β_{firm} in steady state; using the market rate rather than a fixed β ensures consistency with the general equilibrium interest rate path during transitions. The slope κ_p governs the sensitivity of inflation to marginal cost deviations. In the standard Rotemberg derivation with quadratic adjustment costs $\frac{\theta_p}{2}\pi_t^2 Y_t$, the coefficient on the level marginal cost gap $(mc - 1/\mu)$ is $\kappa_p = (\epsilon - 1)/\theta_p$. With $\epsilon = 10$ and $\theta_p = 100$, this gives $\kappa_p = 0.09$.[2] In steady state with $\pi = 0$, this pins $mc = 1/\mu$.

3.5. Government

Monetary policy. The central bank sets the real interest rate according to a Taylor rule:

$$r_t = r^* + \phi_\pi \, \pi_t + \phi_y \left(\frac{Y_t}{Y^{\text{ref}}} - 1 \right), \tag{16}$$

where r^* is the natural rate of interest, ϕ_π is the inflation response coefficient, ϕ_y is the output gap response coefficient, and Y^{ref} is steady-state output. Eq. (16) is written directly for the real rate r_t. The coefficient $\phi_\pi = 1.5$ on the real rate implies that a one-percentage-point increase in inflation raises the real rate by 1.5pp (and the nominal rate by 2.5pp), satisfying the Taylor principle.

Implementation note. To avoid a DAG cycle, in which r depends on Y but Y depends on household decisions that depend on r, I implement the Taylor rule as a residual equation $\text{taylor\_res} = r - r^* - \phi_\pi \pi - \phi_y (Y/Y^{\text{ref}} - 1) = 0$, treating r as an additional unknown. When $\phi_y = 0$, this reduces to the standard monetary rule $r = r^* + \phi_\pi \pi$.

Fiscal policy. The government issues one-period liquid bonds B^g and levies a lump-sum tax to finance interest payments:

$$T_t = r_t \cdot B^g. \tag{17}$$

Government debt B^g is held constant. The lump-sum transfer T_t is a net object: it implicitly finances UI payments ($b_{ui} \cdot e$ for unemployed workers), absorbs firm profits from monopolistic markups (which are not rebated to households as dividends), and services government debt.[3]

[2] Some derivations write $\kappa_p = \epsilon/\theta_p$, yielding 0.10. The difference reflects the normalization of the marginal cost gap term and is absorbed into the calibrated θ_p, which is not independently identified. Results are invariant to the convention given the same κ_p.

[3] Specifically, the household budget constraint treats UI as direct income ($y = b_{ui} \cdot e$ when unemployed) and the tax T as a single net transfer. Firm profits from the markup ($Y_t - mc_t \cdot Y_t = (1 - mc_t)Y_t$) are not distributed

Asset returns. The liquid return equals the policy rate, $r^b = r$. The illiquid return includes an exogenous premium: $r^a = r + \xi$, where ξ captures the higher average return on illiquid assets (housing, retirement accounts).

Bond market clearing. Household aggregate liquid bond holdings equal government bond supply:

$$B_t = B^g, \tag{18}$$

where $B = \int b \, d\Gamma$ is aggregate liquid savings. The illiquid asset has an exogenous return and does not require a separate market-clearing condition.

3.6. Equilibrium

Definition 1. A **recursive competitive equilibrium** consists of household policy functions $\{c(b,a,h,e,s),\ b'(b,a,h,e,s),\ a'(b,a,h,e,s)\}$, firm decisions $\{\theta,\ \text{pricing}\}$, and aggregates $\{Y, C, r, \pi, w, \theta\}$ such that:

(i) Households optimize: policy functions satisfy the Bellman equation.

(ii) Firms optimize: free entry (13) holds, and prices satisfy the NKPC (15).

(iii) The monetary rule (16) and fiscal rule (17) hold.

(iv) The liquid bond market clears: $B = B^g$.

(v) The distribution $\Gamma(b,a,h,e,s)$ is stationary and consistent with policy functions and the transition matrix.

3.7. Solution Method

The model is solved using the Sequence-Space Jacobian (SSJ) method of Auclert et al. (2021).

Steady state. The steady-state unknowns are market tightness θ and the real interest rate r. These are pinned by two targets: the free entry condition (13) and bond market clearing (18). The NKPC pins $mc = 1/\mu$ in steady state (with $\pi = 0$), and production gives $Y = Z \cdot H_{\text{eff}}$. The household problem is solved via a 7-step EGM procedure that handles the constrained (liquid borrowing limit binding) and unconstrained cases separately. The steady state is found by a standard nonlinear solver.

to households and are absorbed by this wedge. The aggregate resource constraint $Y_t = C_t + \kappa_v \cdot V_t + \Psi_{\text{adj},t}$ (output equals consumption plus vacancy posting costs plus portfolio adjustment costs) is verified numerically at the steady state and holds as an equilibrium condition via Walras' law given the bond market clearing condition (18) and the household and firm budget constraints.

Human capital discretization. Since human capital evolves deterministically via (8), the implied next-period value h' will generically fall between grid points. To handle this on a discrete grid, I use a lottery method: for each current grid point h_i, I compute h' from (8) and find the two adjacent grid points $h_j \leq h' \leq h_{j+1}$. Transition probability is then split between these points with weights proportional to proximity, i.e. $\omega = (h_{j+1} - h')/(h_{j+1} - h_j)$ is placed on h_j and $1 - \omega$ on h_{j+1}. This preserves the conditional expectation of h' exactly and avoids the need for interpolation over the human capital dimension during backward iteration. The resulting transition probabilities are embedded in the joint (employment status $\times$ human capital) Markov matrix Π_{sh}.

Constrained EGM. In the unconstrained case, the two-asset EGM inverts both Euler equations (6)–(7) simultaneously: for each candidate b', the ratio of discounted marginal values pins the optimal illiquid choice a', and the budget constraint recovers the endogenous current liquid wealth b. When the liquid borrowing constraint binds ($b' = 0$), however, the liquid Euler equation holds with strict inequality and cannot be inverted. Instead, I introduce a Lagrange multiplier $\lambda \geq 0$ on the liquid constraint and use it as the endogenous grid variable: for each value of λ on an auxiliary grid, the modified first-order condition $W^a/[(1+\lambda)W^b] = 1 + \Psi_1$ pins a', and the budget constraint (with $b' = 0$) recovers b. Inverting the resulting $\lambda \to b$ mapping yields the constrained illiquid policy $a'(b)$. The final policy combines the two cases: if the unconstrained solution implies $b' < 0$, the constrained policy is used instead.

Dynamics. The economy is decomposed into blocks: (i) the household heterogeneous-agent block, (ii) the labor market block (vacancy posting and Nash wage), (iii) the NKPC, (iv) the Taylor rule, (v) fiscal policy, (vi) production, and (vii) bond market clearing. These blocks are assembled into a directed acyclic graph (DAG), and Jacobians of each block's outputs with respect to its inputs (as sequences over time) are computed numerically. General equilibrium impulse responses are computed by solving the system:

$$\mathbf{H}(\mathbf{X}) = \mathbf{0}, \tag{19}$$

where $\mathbf{X}$ is the sequence of endogenous unknowns and $\mathbf{H}$ stacks all equilibrium conditions.

For the baseline model (without the output gap term in the Taylor rule), the dynamic unknowns are $\{\theta_t, mc_t, \pi_t\}$ and the target residuals are free entry (13), the NKPC (15), and bond market clearing (18). When $\phi_y > 0$, r_t is added as a fourth unknown with the Taylor rule (16) as the corresponding target.

MIT shock assumption. The SSJ method computes impulse responses to unanticipated one-time (MIT) shocks around a deterministic steady state. Households' savings decisions in the steady state reflect optimal self-insurance against *idiosyncratic* risk (employment transitions, productivity shocks) but do not price *aggregate* risk; they behave as if aggregate recessions have zero probability. If agents could form expectations over the ergodic distribution of aggregate shocks (including severe state-dependent depreciation episodes with $\kappa > 0$), standard buffer-stock motives would increase precautionary savings at the liquid margin, reducing the HtM share. This is a maintained assumption shared by the broader HANK literature using the SSJ method (Auclert et al., 2021; Kaplan et al., 2018). The 27% HtM share is calibrated to match the empirical SCF range (24–31%, Kaplan et al., 2014), so the steady-state wealth distribution is consistent with the portfolio behavior of households who do in fact face aggregate recession risk. To the extent that the observed SCF distribution already reflects precautionary responses to aggregate risk, the calibration implicitly accounts for it.

3.8. State-Dependent Scarring and Training Subsidies

The baseline depreciation rate δ_h is constant: an unemployed worker loses human capital at the same rate regardless of the aggregate state. However, there is strong evidence that displacement is *more damaging* during severe recessions. Davis and von Wachter (2011) show that earnings losses from displacement are roughly twice as large for workers separated during high-unemployment periods. Huckfeldt (2022) documents that occupational downgrading, a proxy for skill mismatch and human capital loss, is concentrated in severe downturns. The mechanisms include destruction of firm-specific capital (more firm closures, not just layoffs), worse skill mismatch when many workers search simultaneously, and fewer retraining opportunities when vacancies are scarce.

I capture this state-dependence by allowing depreciation to intensify during recessions. Define the effective depreciation rate:

$$\delta_{h,t}^{\text{eff}} = \max\Big[\delta_h,\ \ \delta_h\,(1+\kappa\,\hat{s}_t)\,(1-\tau\,\hat{s}_t)\,\Big], \qquad \hat{s}_t \equiv \frac{s_t - s^{\text{ss}}}{s^{\text{ss}}}, \tag{20}$$

where $\hat{s}_t$ is the proportional deviation of the separation rate from steady state, $\kappa \geq 0$ is the scarring amplification parameter, and $\tau \geq 0$ is the training subsidy intensity. When $\kappa = 0$, depreciation is constant (the baseline). When $\kappa > 0$, a recession ($s_t > s^{\text{ss}}$) accelerates human capital loss, capturing the empirical finding that displacement is more damaging during severe downturns. The training subsidy τ directly offsets the accelerated depreciation, representing a reduced-form stand-in for active labor market programs (retraining,

apprenticeships, on-the-job subsidies). I impose a floor $\delta_{h,t}^{\text{eff}} \geq \delta_h$: the subsidy can reduce recession-amplified depreciation but cannot push depreciation below its normal-times level.[4]

A key property is that *the steady state is invariant to* (κ, τ): since $\hat{s}_t = 0$ at the steady state, the effective depreciation rate equals δ_h regardless of the policy parameters. This allows me to solve the steady state once and vary only the dynamic Jacobians across policy scenarios.

The fiscal cost of the training subsidy is:

$$T_t = r_t \cdot B^g + p_{\text{train}} \cdot \tau \cdot \hat{s}_t, \tag{21}$$

where p_{train} is the fiscal cost scaling parameter. The subsidy cost is proportional to the exogenous shock $\hat{s}_t$, not to endogenous unemployment, avoiding feedback loops in the model's DAG structure.

4. Calibration

All parameters are quarterly. The calibration strategy is sequential: preferences and technology are set to standard values, labor market parameters target US flow rates, and human capital parameters are chosen to generate meaningful scarring dynamics.

4.1. Externally Set Parameters

Table 1 reports parameters set externally.

[4]Without the floor, the multiplicative interaction $(1 + \kappa\hat{s})(1 - \tau\hat{s})$ can imply $\delta_{h,t}^{\text{eff}} < \delta_h$ for sufficiently large $\hat{s}$ when $\tau > 0$. For example, with $\kappa = 2$, $\tau = 1$, and $\hat{s} = 0.6$ (a 3pp shock around $s^{\text{ss}} = 0.05$), the raw multiplier is $0.88 < 1$. The floor ensures the subsidy is a pure offset of recession-amplified depreciation, not a technology subsidy that accelerates skill accumulation beyond normal times.

Table 1: *Externally Calibrated Parameters*

Parameter	Symbol	Value	Source/Target
Preferences			
Discount factor	β	0.99	Standard quarterly
EIS	σ^{-1}	1.0	Log utility
Production and pricing			
Aggregate TFP	Z	1.0	Normalization
Elasticity of substitution	ϵ	10	$\sim$11% markup
Rotemberg cost	θ_p	100	Standard
NKPC slope	κ_p	0.09	$= (\epsilon - 1)/\theta_p$
Labor market			
Matching elasticity	α_m	0.5	Petrongolo–Pissarides
Bargaining power	ν	0.5	Hosios condition
UI replacement	b_{ui}	0.4	US average
Quarterly separation rate	s	0.05	JOLTS
Firm discount factor	β_{firm}	0.99	Standard
Human capital			
Depreciation (unemployed)	δ_h	0.04	GE-consistent calibration
Learning rate	ϕ_0	0.01	Matched to recovery
Learning curvature	ϕ_2	0.5	Diminishing returns
Scarring elasticity	ζ	0.5	Duration dependence
Wage-HC elasticity	ϕ_w	0.5	GE-consistent calibration
HC range	$[h_{\min}, h_{\max}]$	$[0.1, 2.0]$	–
Reference HC	h_{ref}	1.8	$\approx$ mean h of employed
HC grid points	n_H	20	–
Two-asset structure			
Adj. cost level	χ_0	0.25	KMV (2018)
Adj. cost scale	χ_1	6.5	HtM share in KVW range
Adj. cost curvature	χ_2	2.0	KMV (2018)
Illiquidity premium (quarterly)	ξ	0.007	$\sim$2.8% annual
Monetary and fiscal			
Taylor rule: inflation	ϕ_π	1.5	Standard
Taylor rule: output gap	ϕ_y	0.125	Standard quarterly
Government liquid debt	B^g	0.75	$B^g/Y \approx 0.45$
Scarring policy (Section 3.8)			
Scarring amplification	κ	2	HC share of recession/expansion differential
Training subsidy intensity	τ	$\{0, 0.5, 1\}$	Policy scenarios
Subsidy fiscal cost	p_{train}	0.5	–

Notes: Parameters are set prior to solving the steady state. The vacancy posting cost κ_v and matching efficiency χ_m are determined endogenously (Section 4). Within-employment productivity e follows an AR(1) with persistence $\rho_e = 0.966$. The scarring policy parameters κ and τ affect only dynamic depreciation and are inactive at the steady state ($\hat{s} = 0$). $p_{\text{train}} = 0.5$ is the baseline fiscal scaling; $p_{\text{train}} = 0.10$ (realistic) is also reported.

4.2. Endogenously Determined Parameters

The vacancy posting cost κ_v and matching efficiency χ_m are derived from the target unemployment rate and flow rates. To illustrate the logic, consider the simplifying normalization $\theta = 1$: given targets $u = 0.06$ and $s = 0.05$ quarterly, the implied job-finding rate is $f = s(1-u)/u \approx 0.783$, which pins $\chi_m = f/\theta^{1-\alpha_m} \approx 0.783$. Free entry then determines $\kappa_v \approx 0.531$, and Nash bargaining gives the steady-state wage $w = \nu(mc \cdot Z + \kappa_v\theta) + (1-\nu)b_{ui}$. In the full model, the NK markup wedge ($mc = 1/\mu \approx 0.9$) and human capital heterogeneity cause θ to deviate from unity. The matching efficiency χ_m is therefore calibrated via bisection on the full steady state to deliver $U = 6\%$ (Section 5), which endogenously determines $\theta = 0.824$ and the remaining labor market aggregates reported in Table 2.

4.3. Within-Employment Productivity

Within-employment productivity e follows a 5-state Markov chain approximating an AR(1) process with persistence $\rho_e = 0.966$ and innovation standard deviation $\sigma_e = 0.5$, discretized using the Rouwenhorst method. The liquid asset grid has 50 points from 0 to 50, the illiquid asset grid has 30 points from 0 to 500, and the Lagrange multiplier grid (for the constrained EGM step) has 4 points.

4.4. Human Capital Parameter Calibration

The human capital parameters $(\delta_h, \zeta, \phi_w, \kappa)$ are calibrated using a GE-consistent procedure that requires parameter vectors to produce a well-defined steady state. I simulate the expected earnings path of a worker displaced at the mean employed human capital level ($h_0 = h_{\text{ref}} = 1.8$), tracking their probability of re-employment and the evolution of their human capital and wages relative to a control worker who remains employed throughout.

Figure 2 shows the model-implied earnings loss profile. Panel (a) compares the model to the empirical benchmarks for different values of ϕ_w. The wage-HC channel ($\phi_w > 0$) is essential for generating persistent earnings losses: with $\phi_w = 0$, the wage rate $w(h) = w$ is independent of h, so the wage-channel component of earnings losses disappears immediately upon re-employment. Residual losses from the h-multiplier in the income expression $w \cdot e \cdot h$ attenuate within 2–3 quarters as h recovers through learning-by-doing. With $\phi_w = 0.5$, the model generates persistent losses of approximately 2–3% at 5 years post-displacement.

I calibrate the state-dependent scarring parameter κ to target the human capital component of the recession-expansion displacement differential. Davis and von Wachter (2011)

document that earnings losses from recession displacement are roughly twice as large as those from expansion displacement. Huckfeldt (2022) decomposes persistent earnings losses into human capital depreciation (∼50–60%) and other channels (occupational downgrading, firm-specific premia, employer stigma). Since the model captures only the HC channel, I calibrate κ so that the model's recession/expansion HC-loss ratio falls within the range implied by these estimates, subject to GE feasibility constraints. Panel (d) shows that $\kappa = 2$ generates recession displacement losses 1.3–1.5$\times$ expansion losses. The multiplicative specification $\delta_h(1 + \kappa\hat{s})$ is the simplest functional form linking depreciation to the aggregate state; an additive alternative $\delta_h + \kappa'\hat{s}$ would produce qualitatively similar results. Appendix I.2 reports distributional welfare for $\kappa \in [0, 5]$, showing that the distributional gradient widens monotonically with κ, so the qualitative findings do not depend on a precise calibration.

The omitted channels (occupational downgrading, firm-specific wage premia, employer stigma) are likely to *reinforce* the distributional ranking. All three are amplified by illiquidity: liquidity-constrained workers cannot sustain a prolonged search for a within-occupation match (Huckfeldt, 2022), accept the first available offer at lower wages (Krolikowski, 2017), and face longer unemployment duration that compounds employer stigma (Kroft, Lange, and Notowidigdo, 2013). The distributional gradient documented in Section 6.2 is therefore a conservative estimate of the true incidence of scarring.

A key constraint on the calibration is general equilibrium consistency. Pure partial-equilibrium matching of the full ∼20% earnings drop from Jacobson et al. (1993) requires $\delta_h \approx 0.17$/quarter, which in general equilibrium collapses output by half and pushes the real rate negative. The GE-feasible bound is $\delta_h \leq 0.05$, with the binding constraint being ζ: high finding-rate elasticity traps workers in long-term unemployment and pushes U above 12%. I set $\delta_h = 0.04$/quarter and $\phi_w = 0.5$ as the baseline, which generates a healthy steady state ($r = 0.15\%$, $Y = 1.652$, $U = 6.0\%$) while maximizing the HC scarring channel within GE bounds.

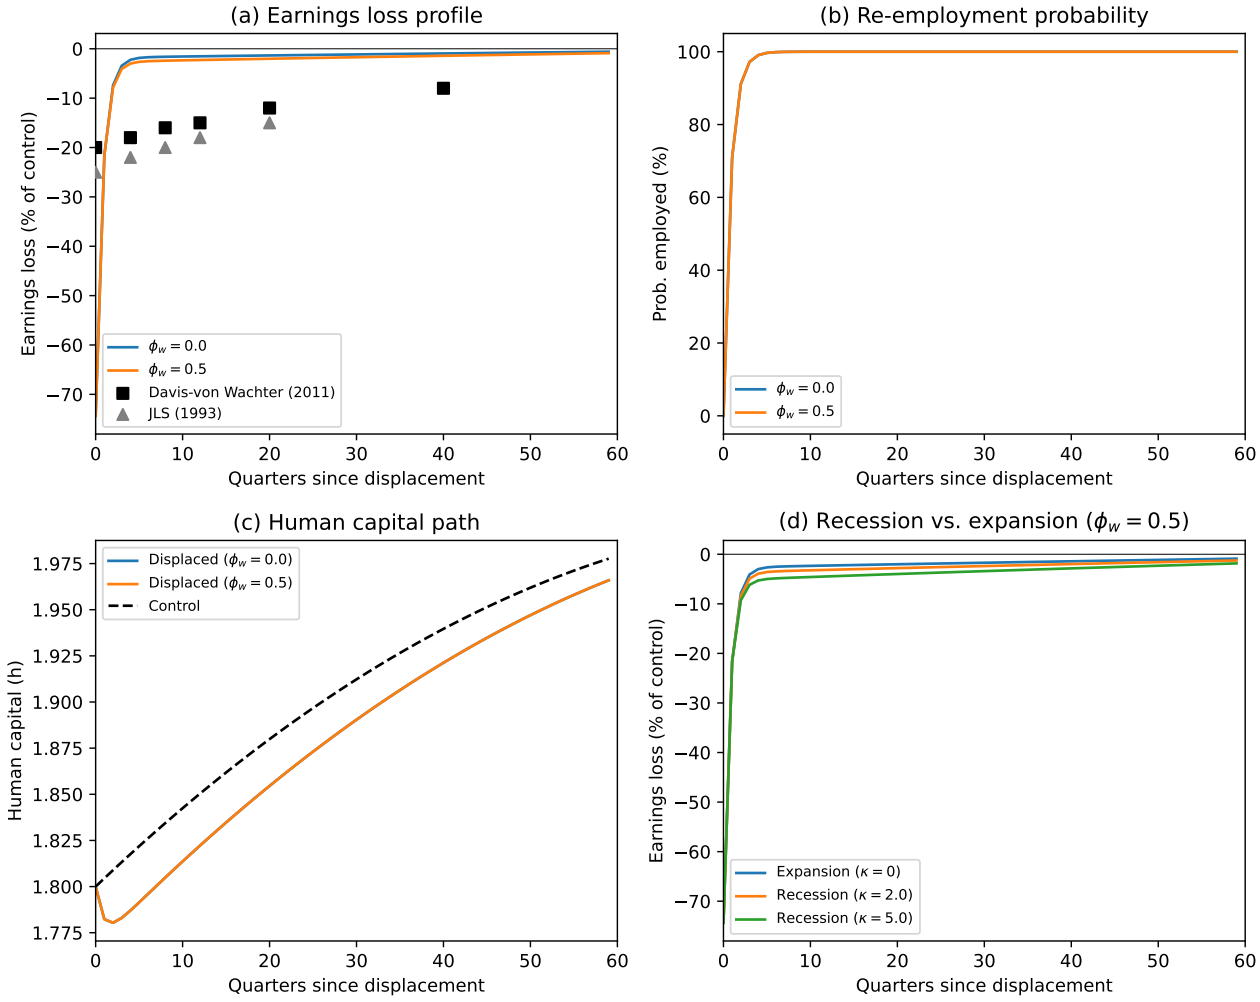

Figure 2: *Model-implied earnings loss profile for a worker displaced at $h_0 = 1.8$. (a) Earnings loss as percent of control earnings, comparing $\phi_w = 0$ (no wage-HC channel) and $\phi_w = 0.5$ (moderate wage scarring), with empirical benchmarks from Davis and von Wachter (2011) and Jacobson et al. (1993). (b) Re-employment probability. (c) Human capital paths for displaced and control workers. (d) Recession vs. expansion displacement losses under state-dependent scarring with $\phi_w = 0.5$: the calibrated $\kappa = 2$ generates recession losses 1.3–1.5× expansion losses, targeting the HC component of the ∼2× differential documented by Davis and von Wachter (2011), adjusted for Huckfeldt's (2022) decomposition of persistent losses into HC depreciation (∼50–60%) and other channels.*

5. Steady State

Table 2 reports key steady-state statistics for the baseline calibration.

Table 2: *Steady-State Equilibrium ($\zeta = 0.5$, $\phi_w = 0.5$, $\delta_h = 0.04$)*

Variable	Symbol	Value
Labor market		
Market tightness	θ	0.824
Wage	w	0.869
Job-finding rate	f	0.826
Unemployment rate	U	6.0%
Real marginal cost	mc	0.900
Production and aggregates		
Effective labor	H_{eff}	1.652
Output	Y	1.652
Asset markets		
Real interest rate (quarterly)	r	0.0015
Real interest rate (annual)	$4r$	0.61%
Liquid return	r^b	0.0015
Illiquid return	r^a	0.0085
Government liquid debt	B^g	0.75
Liquid wealth / output	B/Y	0.45
Illiquid wealth / output	A/Y	19.7
Lump-sum tax	T	0.001
Hand-to-mouth shares		
Total HtM		27.2%
Wealthy HtM		27.2%
HtM (employed)		26.7%
HtM (unemployed)		35.2%

Notes: Deterministic steady state. HtM defined as $b = 0$ (first liquid grid point). Interest rates are quarterly. The real rate $r = 0.15\%$ quarterly ($\approx$ 0.61% annualized) clears the liquid bond market against government debt $B^g = 0.75$. The steady state is invariant to the scarring policy parameters (κ, τ).

The steady state features a 6.0% unemployment rate, matching the calibration target. Market tightness $\theta = 0.824$ is below the normalization target of 1.0 because the NK markup wedge ($mc = 1/\mu \approx 0.9$) reduces effective firm productivity in bargaining, lowering the value of posting vacancies. The real interest rate of 0.15% quarterly (0.61% annually) clears the liquid bond market against government debt $B^g = 0.75$.

Wealth distribution and hand-to-mouth households. Approximately 27% of households hold near-zero liquid assets ($b \approx 0$), within the 24–31% range documented by Kaplan et al. (2014). Nearly all are "wealthy hand-to-mouth": households with positive illiquid wealth but no liquid buffer. The HtM share is higher among the unemployed (35.2%) than the employed (26.7%), reflecting the rapid drawdown of liquid savings during joblessness. Figure 3 shows the joint wealth distribution.

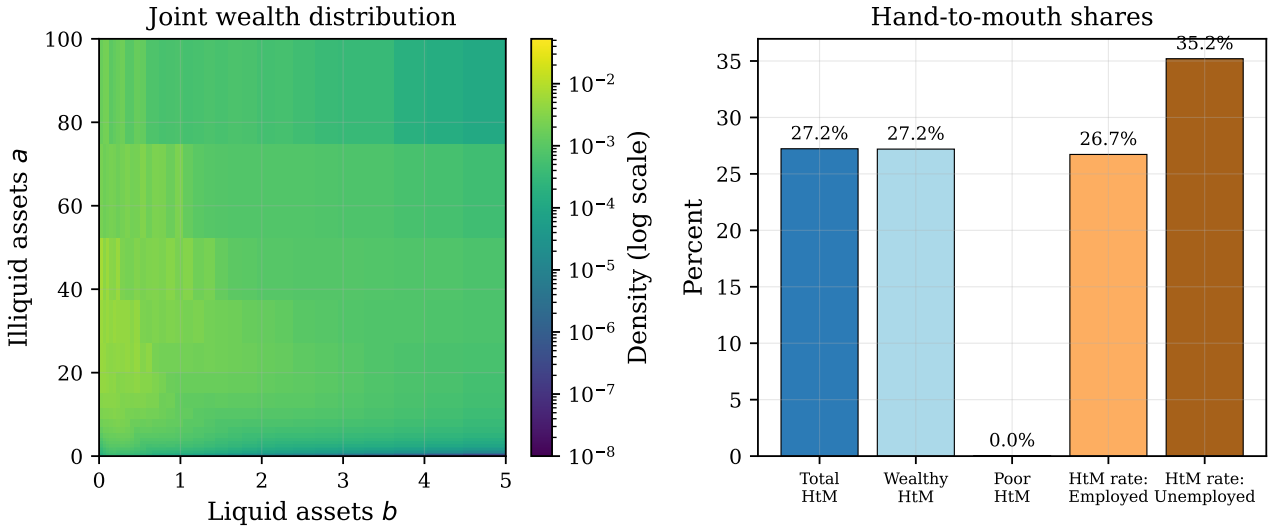

Figure 3: *Left: Joint distribution of liquid (b) and illiquid (a) wealth in steady state (log density scale). Mass is concentrated along the $b = 0$ edge across a range of illiquid wealth, the "wealthy hand-to-mouth" pattern. Right: Hand-to-mouth shares by employment status. Nearly all HtM households are wealthy HtM. The HtM rate is higher among the unemployed (35.2%) than the employed (26.7%).*

6. Aggregate and Distributional Dynamics

I study the model's response to a separation rate shock, a labor market disturbance that directly triggers the scarring mechanism. Section 6.1 establishes baseline aggregate dynamics. Section 6.2 introduces the distributional methodology and shows that constant depreciation produces only a mild distributional divide. Section 6.3 introduces state-dependent scarring, which sharpens this divide dramatically. Section 6.4 presents the paper's main contribution: the distributional incidence of recession costs at Great Recession severity and the convexity of distributional costs in shock size.

6.1. Separation Shock

The baseline experiment is a 1 percentage point increase in the quarterly separation rate, following an AR(1) process with persistence $\rho_s = 0.9$:

$$s_t - s^{ss} = \varepsilon \cdot \rho_s^t, \qquad \varepsilon = 0.01. \tag{22}$$

This generates a recession through the labor market channel: higher separations raise unemployment, reduce effective labor, lower output, and depress consumption through both income and precautionary savings effects. Output drops approximately 1.1% on impact, reaching a peak decline of 1.2%, consumption falls 1.0%, and unemployment rises 1.1 percentage points, recovering over 30+ quarters.[5] Figure 4 shows the impulse responses.

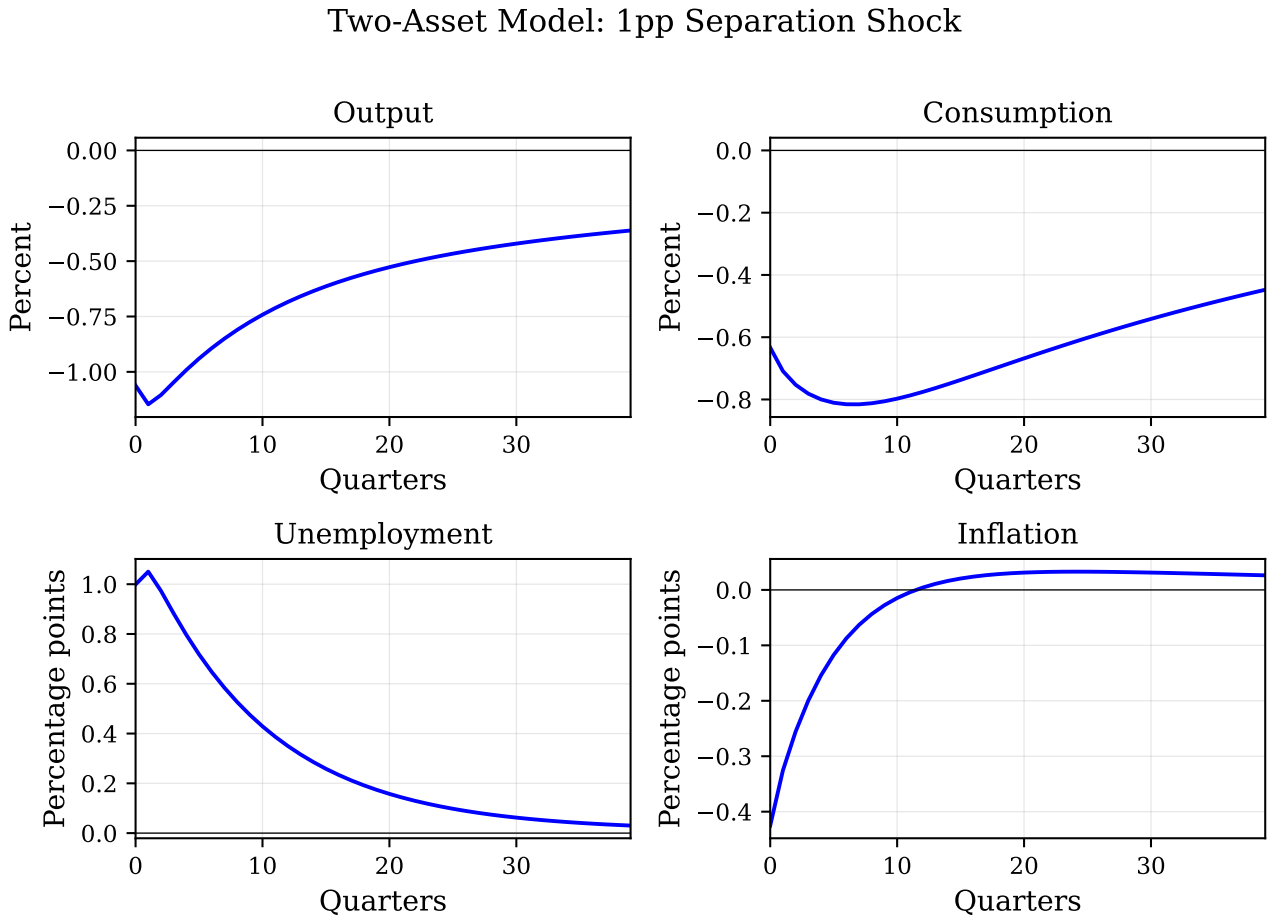

Figure 4: *Impulse responses to a 1pp separation shock (AR(1), $\rho = 0.9$) under the baseline Taylor rule ($\phi_\pi = 1.5$, $\phi_y = 0.125$) with constant depreciation ($\kappa = 0$). Panels show output, consumption, unemployment, and inflation over 60 quarters.*

[5] A comparison with the one-asset model is in Appendix D. Output and unemployment paths are nearly identical across model variants, but consumption and inflation dynamics differ due to the liquid-illiquid wealth composition.

6.2. Distributional Incidence at Constant Depreciation

The model's comparative advantage is not aggregate amplification but its ability to trace how recession costs are allocated across the wealth distribution. I exploit the nonlinear simulation's ability to track the full joint distribution $\Gamma_t(b, a, h, e, s)$ to compute group-specific consumption paths, unemployment rates, and welfare losses.

Methodology. I use the SSJ toolkit's nonlinear solver to extract the distribution path $\{D_t\}$ and consumption policy $\{c_t\}$ at each period. I define four groups by liquid wealth position: (i) hand-to-mouth households ($b = 0$), approximately 27% of the population; and among non-HtM households, (ii) low liquid wealth (bottom tercile of $b > 0$), (iii) mid liquid (middle tercile), and (iv) high liquid (top tercile). Group-specific average consumption $\bar{C}_{g,t}$ and unemployment rates $U_{g,t}$ are computed by integrating over the time-varying distribution within each group. The consumption-equivalent (CE) welfare variation for group g is:

$$\text{CE}_g = \exp\left(\frac{\sum_{t=0}^{H} \beta^t (\bar{C}_{g,t} - \bar{C}_g^{\text{ss}})/\bar{C}_g^{\text{ss}}}{\sum_{t=0}^{H} \beta^t}\right) - 1. \tag{23}$$

where $H = 60$ quarters is the welfare horizon. The CE formula uses the linear consumption ratio $(\bar{C}_{g,t} - \bar{C}_g^{\text{ss}})/\bar{C}_g^{\text{ss}}$ inside the exponent rather than the exact $\log(\bar{C}_{g,t}/\bar{C}_g^{\text{ss}})$; at the 3pp shock's peak consumption drop of $\sim$6%, the approximation error is $\sim$0.18pp, which is small relative to the distributional gradient of -1.5pp. This measure uses group-average consumption rather than integrating expected utility across individual household states within each group.[6]

Table 3 reports steady-state characteristics and recession welfare costs for each liquid wealth group.

[6]With log utility and within-group heterogeneity in (b, a, h, e), Jensen's inequality implies $u(\mathbb{E}[c]) \geq \mathbb{E}[u(c)]$, so group-average CE may understate the welfare cost for groups with higher within-group consumption dispersion. Group-average CE is the appropriate welfare metric in this setting because it correctly accounts for both channels through which recessions reduce group welfare: (i) changes in the consumption policy function $c_t(b, a, h, e, s)$, and (ii) compositional shifts in the distribution Γ_t within each group (e.g., more households becoming unemployed or losing human capital). Both channels are captured by computing $\bar{C}_{g,t}$ from actual distribution-weighted consumption $\int c_t \, d\Gamma_t$ at each t, using the full nonlinear distribution path. An alternative approach, computing individual-level CE at each grid point and then averaging, would capture policy function changes but miss the compositional channel, since it evaluates welfare changes at fixed states without tracking how households transition between states over time. The potential Jensen's bias is a second-order concern relative to this first-order compositional effect, particularly for a 1pp shock where within-group consumption dispersion changes are modest.

Table 3: *Distributional Incidence: CE Welfare (%) by Liquid Wealth Group for a 1pp Separation Shock (AR(1), $\rho = 0.9$) with Constant Depreciation ($\kappa = 0$)*

	Aggregate	HtM ($b = 0$)	Low liquid	Mid liquid	High liquid
Steady-state characteristics					
Population share (%)	100	27.2	24.0	22.7	26.1
Avg. consumption	1.70	1.83	1.34	1.56	2.02
Unemployment rate (%)	6.0	7.7	5.4	5.2	5.3
Recession costs (1pp separation shock, $\kappa = 0$)					
CE welfare (%)	−0.57	−0.53	−0.56	−0.42	−0.46
Peak C drop (%)	−1.0	−0.7	−1.0	−0.8	−0.6

Notes: 1pp separation shock, AR(1) with $\rho = 0.9$. CE variation computed over 60 quarters. HtM defined as $b = 0$ (first liquid grid point). Non-HtM terciles defined by CDF of b conditional on $b > 0$. The aggregate CE is computed from total consumption, not as a weighted average of subgroup CEs; with nonlinear CE (exponential mapping) and time-varying group composition, the aggregate need not lie within the range of subgroup values.

With the wage-human capital channel ($\phi_w = 0.5$), the distributional pattern at constant depreciation is relatively uniform. HtM households bear a CE loss of -0.53%, comparable to low-liquid non-HtM workers (-0.56%), while mid- and high-liquid groups face somewhat smaller losses (-0.42% and -0.46%). The mild gradient reflects two forces. First, HtM households face the highest steady-state unemployment rate (7.7% vs. 5.2–5.4% for non-HtM groups), increasing their exposure to the depreciation channel. Second, the wage-human capital channel ($w_{\text{eff}} = w \cdot (h/h_{\text{ref}})^{\phi_w}$) means that upon re-employment, workers with depreciated h earn lower wages, and HtM households, who cannot buffer income losses during unemployment, are most exposed to this penalty. However, at constant depreciation rates the quantitative impact is modest, leaving costs roughly evenly distributed. The h-dependent finding rate channel ($\zeta = 0.5$) contributes only modestly at constant depreciation: removing it ($\zeta = 0$) changes group-level CE by at most -0.08pp (Appendix B).

Figure 5 shows the dynamics. Panel (a) displays the CE welfare cost by group. Panel (b) shows consumption paths by group: low-liquid workers experience the sharpest consumption decline ($\sim$1.0% peak), reflecting their low average consumption level, while HtM workers experience a peak drop of $\sim$0.7%. Panel (c) shows unemployment responses are relatively uniform across groups; the heterogeneity in welfare costs comes from consumption insurance and the wage channel, not differential unemployment exposure. Panel (d) shows that the aggregate HtM share decreases during the recession as the precautionary savings motive moves marginal households above the liquid constraint.

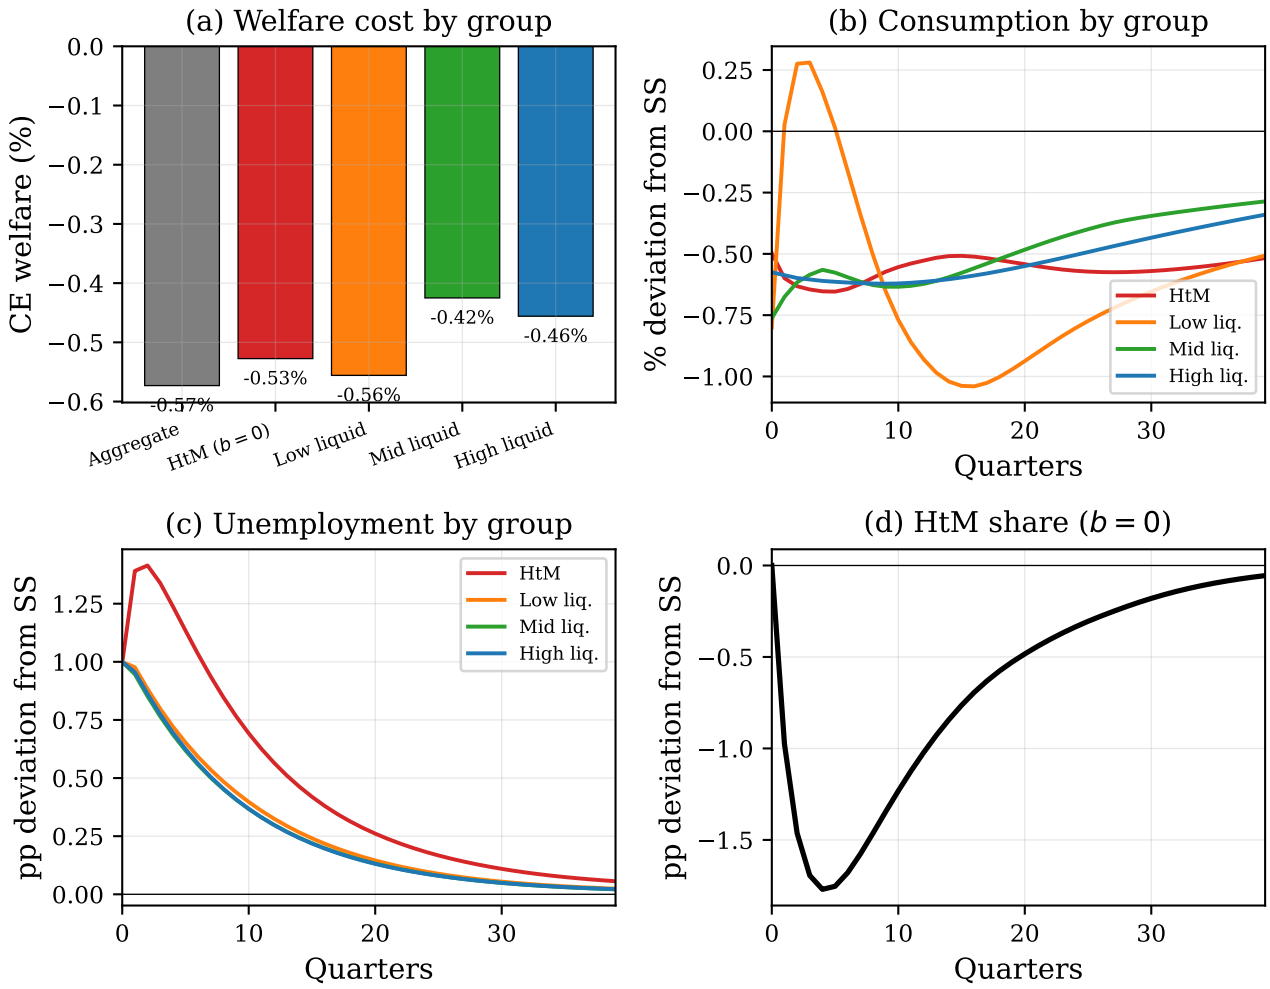

Figure 5: *Distributional incidence of a 1pp separation shock (AR(1), $\rho = 0.9$) with wage-human capital channel ($\phi_w = 0.5$) and constant depreciation ($\kappa = 0$). (a) CE welfare cost by liquid wealth group. (b) Consumption paths (percent deviation from SS). (c) Unemployment by group (percentage point deviation). (d) HtM share dynamics. At constant depreciation, recession costs are distributed relatively evenly across liquid wealth groups; state-dependent scarring ($\kappa > 0$, Table 6) generates a sharper distributional divide.*

6.3. State-Dependent Scarring

State-dependent scarring (Section 3.8) intensifies human capital depreciation during recessions. With $\kappa = 2$, the effective depreciation rate increases by 40% at the peak of a 1pp separation shock ($\hat{s} = 0.2$), to $\delta_h^{\text{eff}} \approx 0.056$/quarter. This accelerated skill loss pushes more workers into the low-h region *during* the recession, when the scarring trap is most damaging.

The primary aggregate consequence is persistence, not amplification. The peak output loss increases modestly from -1.2% to -1.4%, but the shape of the recovery changes fundamentally. Under constant depreciation, output recovers to -0.5% by quarter 20 and

-0.3% by quarter 40. Under $\kappa = 2$, output remains 1.3% below steady state at quarter 20, when the exogenous shock has dissipated to 12% of its initial value, and is still 1.1% below at quarter 40. The CE welfare cost nearly doubles from -0.57% to -1.02%, driven almost entirely by this persistence rather than the peak response.[7]

State-dependent scarring also sharpens the distributional gradient considerably. Recall that at constant depreciation ($\kappa = 0$, Table 3), recession costs are distributed relatively evenly across liquid wealth groups. Under $\kappa = 2$ (Table 6), HtM welfare deteriorates by -0.60pp from the baseline, the largest worsening of any group, making them clearly the worst-off group at -1.13%.

6.4. Great Recession Calibration and Nonlinearity

The distributional costs of scarring become dramatically larger at Great Recession severity. I repeat the analysis with a 3pp separation shock, calibrated to the approximately 5 percentage point rise in the US unemployment rate during 2007–2009.[8] Table 4 reports the results. Under constant depreciation ($\kappa = 0$), the 3pp shock produces aggregate CE of -1.64% with a modest distributional gradient of -0.23pp. Under amplified scarring ($\kappa = 2$), aggregate CE reaches -3.24% and the distributional gradient widens to -1.49pp: HtM households suffer CE losses of -4.07% while high-liquid households lose -2.57%. The HtM loss is nearly 60% larger than the high-liquid group's, a concentration of costs that is far sharper than in the 1pp case (-1.13% vs. -0.84% high-liquid, a 35% excess).

[7] This persistent output effect from a temporary shock is consistent with the empirical findings of Cerra et al. (2023), who document that output losses from recessions are highly persistent. In the model, the economy eventually returns to the same steady state (which is invariant to κ), but recovery takes more than 40 quarters under $\kappa = 2$. Quantitatively, a 40% increase in skill depreciation during the worst quarter of a recession is empirically moderate; results with higher values ($\kappa = 5$, implying a doubling of depreciation) show qualitatively similar but larger effects.

[8] A 3pp separation shock in the model generates a peak unemployment increase of approximately 3.3pp, broadly consistent with the magnitude of the Great Recession. This is a much larger departure from steady state than the 1pp benchmark, putting the nonlinear solver to a harder test.

Table 4: *Distributional Welfare at Great Recession Severity: CE (%) by Liquid Wealth Group for a 3pp Separation Shock (AR(1), $\rho = 0.9$)*

	Aggregate	HtM ($b = 0$)	Low liquid	Mid liquid	High liquid
CE welfare (%)					
Baseline ($\kappa = 0$)	-1.64	-1.59	-1.13	-1.36	-1.36
Amplified ($\kappa = 2$)	-3.24	-4.07	-2.95	-2.27	-2.57
Peak consumption drop (%)					
Baseline ($\kappa = 0$)	—	-2.14	-2.33	-2.26	-1.94
Amplified ($\kappa = 2$)	—	-6.16	-4.05	-3.84	-4.21
Distributional gradient (HtM − High liquid, pp)					
1pp amplified			-0.29		
3pp amplified			-1.49		
Linear prediction (3× 1pp)			-0.87		
Nonlinear excess			-0.62		

Notes: 3pp separation shock, AR(1) with $\rho = 0.9$, $\kappa = 2$ (amplified scarring). CE variation computed over 60 quarters. The distributional gradient is HtM CE minus high-liquid CE; more negative values indicate greater distributional inequality. The linear prediction is the 1pp gradient scaled by 3. Horizon sensitivity: the gradient is -1.49pp at 60Q, -1.51pp at 80Q, and -1.68pp at 100Q, confirming the finding is robust to the CE horizon and, if anything, strengthens at longer horizons.

Figure 6 shows the distributional dynamics. HtM workers suffer a peak consumption drop of -6.16%, compared to -4.21% for high-liquid workers, a 1.46:1 ratio, capturing roughly three-quarters of the 2:1 empirical benchmark of Farrell, Ganong, Greig, and Noel (2019). The contrast with the 1pp case (Figure 5) is dramatic: the distributional divide is modest in average recessions but severe in deep ones.

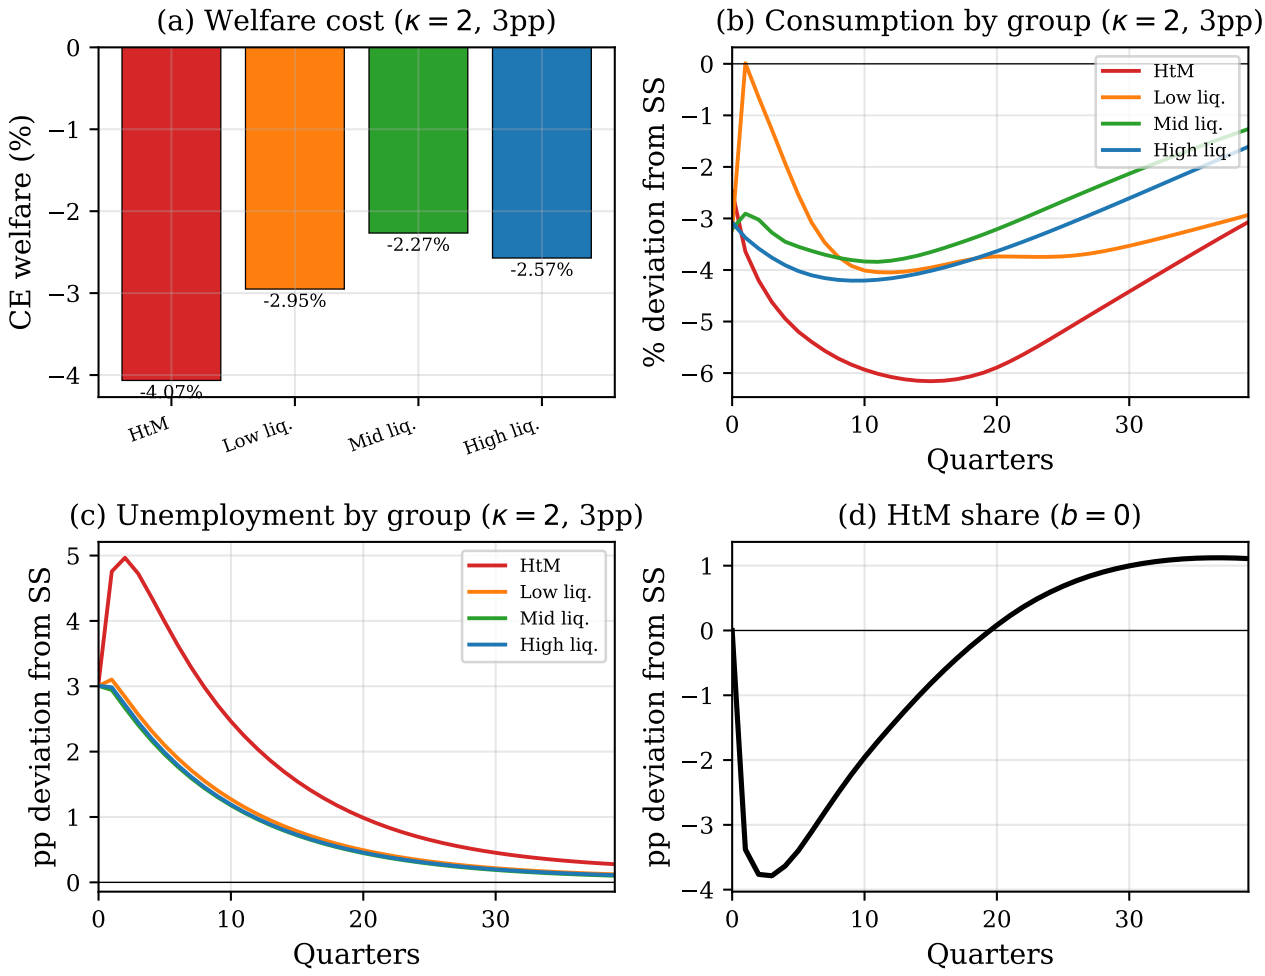

Figure 6: *Distributional incidence of a 3pp separation shock (AR(1), $\rho = 0.9$; Great Recession calibration) with amplified scarring ($\kappa = 2$). (a) CE welfare cost by group. (b) Consumption paths by group. (c) Unemployment paths by group. (d) HtM share dynamics. Compare with Figure 5 (1pp, $\kappa = 0$): the distributional divide widens dramatically under both larger shocks and state-dependent scarring.*

Nonlinear amplification. The most striking finding emerges from comparing shock sizes. If scarring effects scaled linearly, the 3pp distributional gradient would be $3 \times (-0.29) = -0.87$pp. The actual gradient of -1.49pp is approximately five times the 1pp value ($5.1\times$) and exceeds the linear prediction by -0.62pp. Figure 7 shows the result of a shock-size sweep over $[0.5, 4]$pp: the distributional gradient is *convex* in shock size, accelerating as the scarring trap binds more tightly. At 4pp, the gradient reaches -2.48pp, nearly nine times the 1pp value. This convexity implies that severe recessions generate disproportionately larger distributional costs than mild ones. Figure 8 compares the 1pp and 3pp impulse responses and distributional outcomes directly.

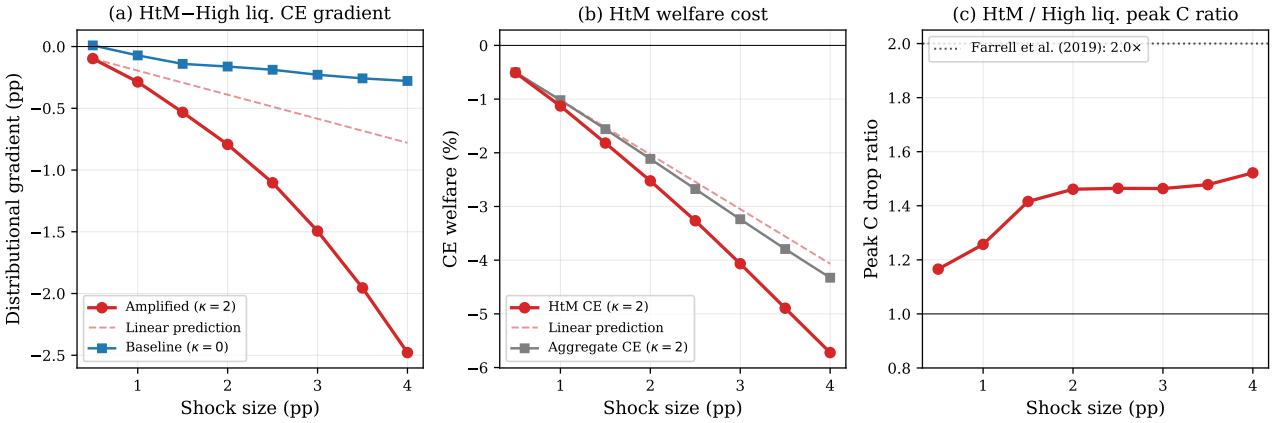

Figure 7: *Nonlinear amplification of distributional costs. Separation shocks of* $[0.5, 4]$*pp (AR(1),* $\rho = 0.9$*) under amplified scarring (*$\kappa = 2$*). (a) Distributional gradient (HtM minus high-liquid CE) as a function of shock size, with linear prediction from smallest shock. The gradient is convex: severe recessions produce disproportionately larger distributional divides. (b) HtM CE welfare as a function of shock size. (c) Peak consumption drop ratio (HtM/high-liquid) as a function of shock size, with the Farrell et al. (2019) 2:1 empirical benchmark.*

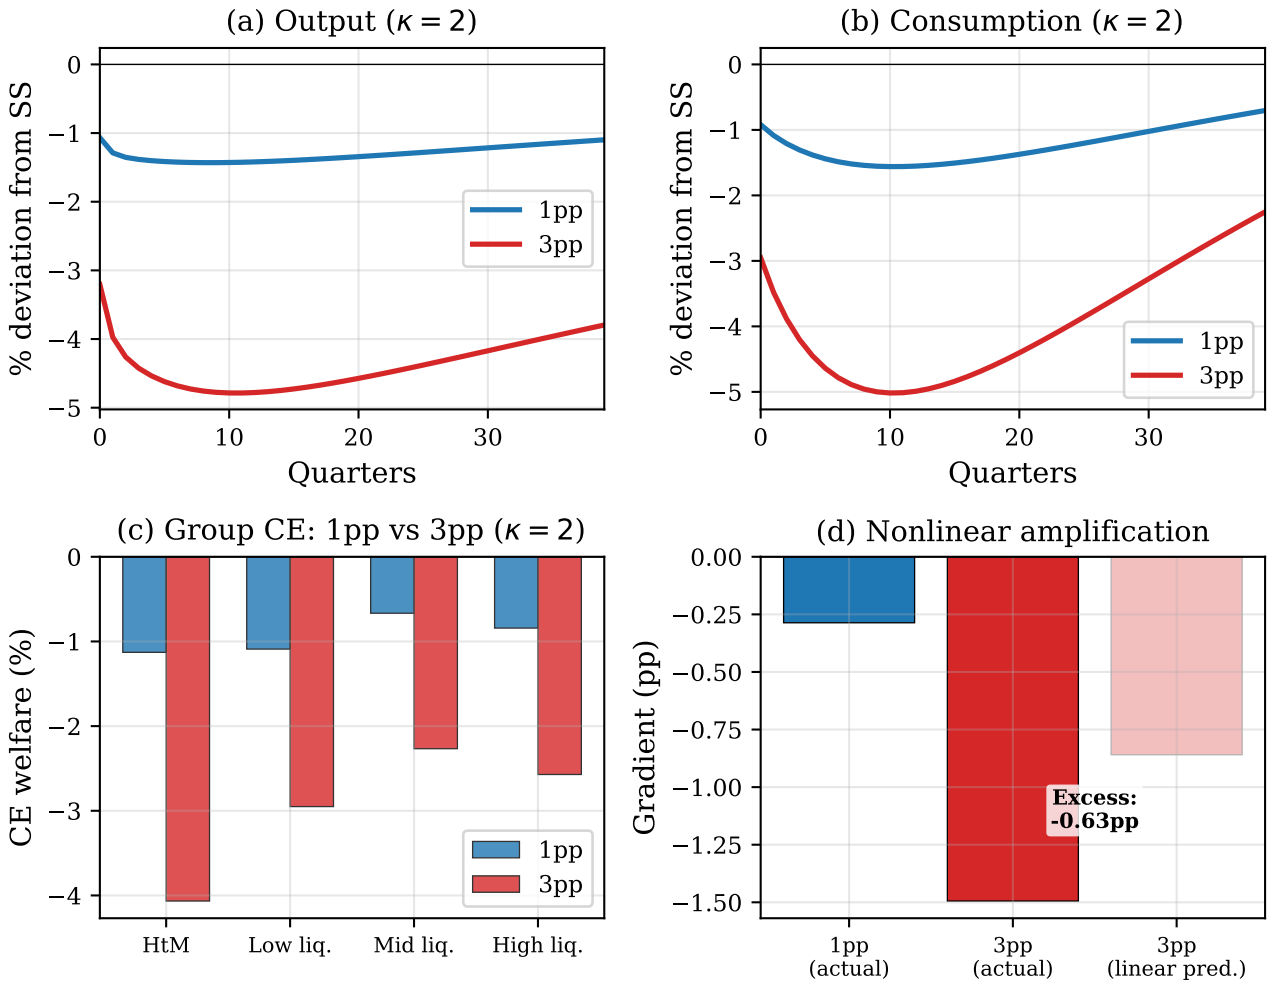

Figure 8: *Comparison of 1pp (average recession) and 3pp (Great Recession) separation shocks (AR(1), $\rho = 0.9$) under amplified scarring ($\kappa = 2$). (a) Output paths. (b) Consumption paths. (c) Group-level CE welfare by liquid wealth group. (d) Nonlinear amplification of distributional gradient: the 3pp gradient exceeds the linear prediction by -0.62pp.*

Welfare decomposition. A potential concern with contemporaneous group membership is that compositional reclassification, rather than within-group welfare changes, drives the distributional gradient. I decompose aggregate CE (-3.24%) into intensive and extensive margins using a (wealth group $\times$ employment status) cross-section (Table 5, Figure 9). The **intensive margin**, welfare changes within each cell weighted by steady-state mass, accounts for 88% of aggregate welfare loss (-2.85%). The **extensive margin** accounts for only 12% (-0.39%). Within employment cells, the HtM-to-high-liquid gradient persists: -1.2pp among the unemployed, -1.3pp among the employed. A fixed-cohort check (assigning households to groups based on steady-state positions) reverses the gradient sign ($+0.18$pp), but this reflects the well-understood limitation of evaluating welfare at unchanged state positions rather than tracking household transitions. The dominance of the intensive margin

confirms that the distributional finding is not driven by compositional shifts.

Table 5: *Welfare Decomposition: Wealth Group × Employment Status for a 3pp Separation Shock (AR(1), $\rho = 0.9$) with Amplified Scarring ($\kappa = 2$)*

	HtM ($b = 0$)	Low liquid	Mid liquid	High liquid
CE welfare (%)				
Employed	−3.76	−2.82	−2.16	−2.45
Unemployed	−4.36	−2.47	−1.94	−3.19
SS population share (%)				
Employed	25.1	22.7	21.5	24.7
Unemployed	2.1	1.3	1.2	1.4

Distributional gradient (HtM − High liquid)		
Among unemployed	−1.17pp	
Among employed	−1.31pp	
Aggregate decomposition		
Intensive (within-cell)	−2.85%	(88% of aggregate −3.24%)
Extensive (mass flow)	−0.39%	(12%)

Notes: 3pp separation shock, AR(1) with $\rho = 0.9$, amplified scarring ($\kappa = 2$). CE variation computed over 60 quarters. Intensive margin = $\sum_{g,e} M^{ss}_{g,e} \cdot \text{CE}_{g,e} / \sum_{g,e} M^{ss}_{g,e}$; extensive margin = aggregate CE minus intensive. The intensive margin measures welfare changes within each cell at steady-state composition; the extensive margin captures the welfare effect of mass redistribution across cells during the recession.

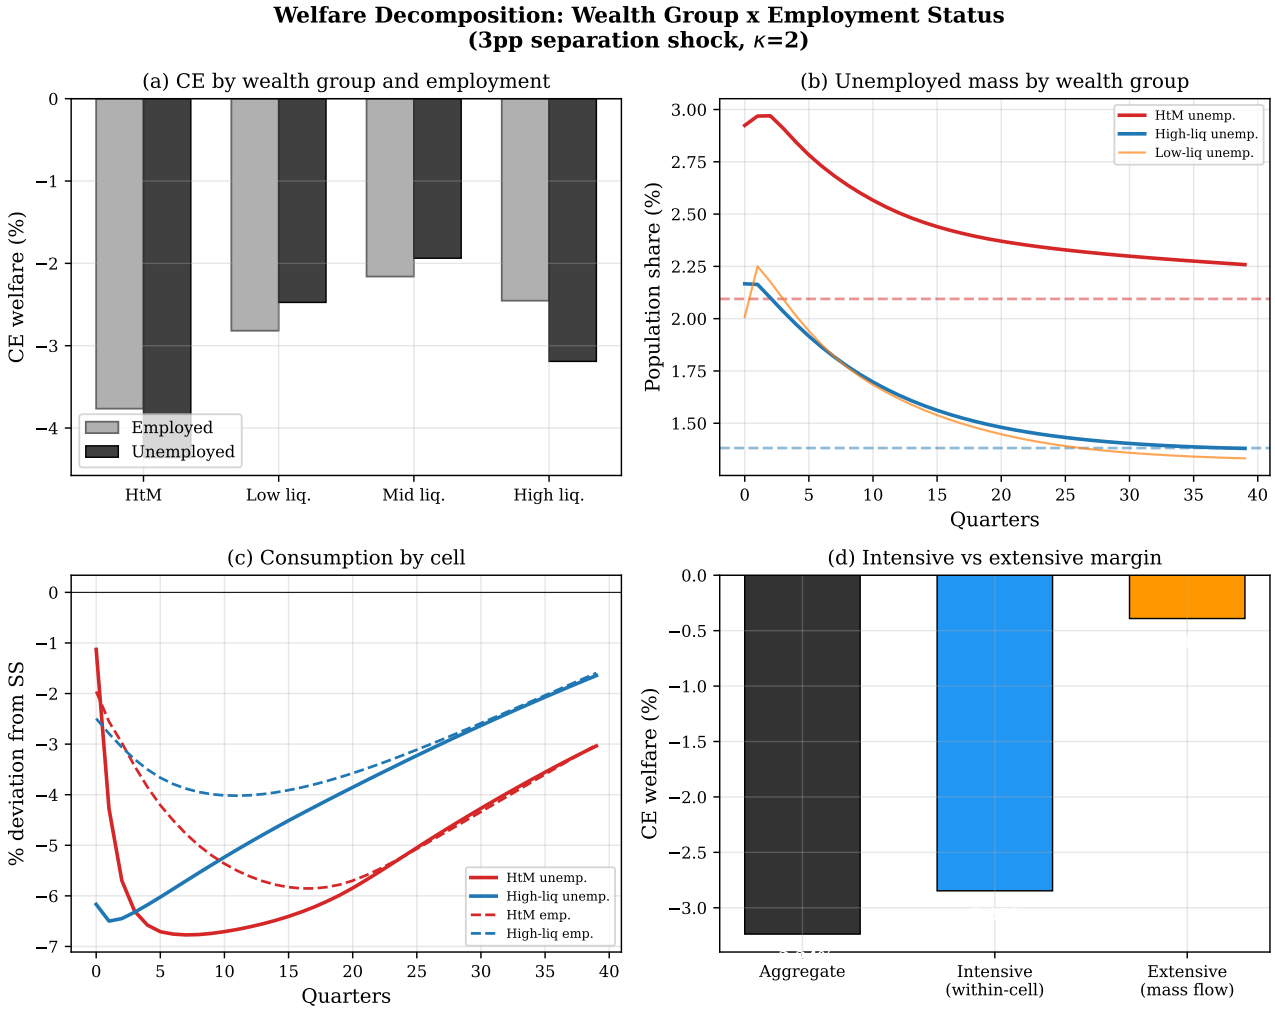

Figure 9: *Welfare decomposition for a 3pp separation shock (AR(1), $\rho = 0.9$) under amplified scarring ($\kappa = 2$). (a) CE welfare by liquid wealth group and employment status: within each group, unemployed workers fare worse, and among both employed and unemployed, HtM workers bear the largest losses. (b) Population share of unemployed by wealth group over time: HtM unemployment peaks highest and is most persistent. (c) Consumption paths by cell: HtM unemployed suffer the sharpest initial drop while high-liquid unemployed experience a slower but deeper trough from scarring. (d) Intensive vs. extensive margin: 88% of aggregate welfare loss is within-cell (intensive), confirming the distributional finding is not driven by compositional shifts.*

Empirical validation. The distributional predictions are broadly consistent with related empirical evidence. Farrell et al. (2019) document a 2:1 spending drop ratio between low-liquid and high-liquid households at unemployment onset; the model produces a 1.5:1 ratio at Great Recession severity (Table 4). Fontaine, Jensen, and Vejlin (2024) show that liquid assets significantly reduce unemployment duration while illiquid assets have a precisely estimated zero effect, exactly the mechanism embedded in the portfolio structure. Schmieder, von Wachter, and Heining (2023) document that displacement costs are sharply

countercyclical, consistent with the state-dependent depreciation that amplifies skill loss during recessions. Table 20 (Appendix I) provides a comprehensive comparison across 12 empirical moments.

Role of the two-asset structure. The liquid-illiquid distinction is essential for isolating the scarring vulnerability channel. A one-asset model generates a distributional gradient across wealth quartiles, but it cannot produce the wealthy hand-to-mouth phenomenon: households with positive illiquid wealth but zero liquid assets who face binding consumption constraints despite positive net worth. Under amplified scarring ($\kappa = 2$), the HtM grouping isolates the vulnerability channel more precisely than wealth quartiles, producing a sharper distributional gradient (-0.29pp vs. -0.22pp in the one-asset model). The net worth quartile analysis (Table 19, Appendix I) confirms that liquid wealth, not net worth, is the operative constraint. See Appendix D for the full one-asset comparison and Table 13.

The model also accommodates match quality heterogeneity, where re-employed workers draw match quality m that affects their learning rate. This channel has negligible aggregate effects but reshapes the cross-sectional distribution of human capital recovery; see Appendix C for details.

7. Policy Experiments

7.1. Monetary Policy

A Taylor rule grid search over (ϕ_π, ϕ_y) reveals that monetary policy has limited leverage over separation shocks: the quadratic stabilization loss is nearly invariant across the parameter space (Appendix F, Figure 15). The mechanism is straightforward: separation shocks operate through the matching function, which is largely orthogonal to the interest rate. More accommodative policy reduces the real rate and stimulates demand, but the effect is quantitatively small because the labor market wedge dominates the interest-sensitive consumption margin. The distributional results below provide a deeper explanation for this finding: the households most affected by scarring (HtM) are precisely those with no savings, so they do not benefit from interest rate channels. A caveat is that this near-irrelevance may be specific to separation shocks; under demand-driven recessions (e.g., discount factor shocks or a binding zero lower bound), monetary policy would have greater leverage. The limited effectiveness of demand management for the labor market shocks studied here motivates the training subsidy as a policy that targets the scarring mechanism directly.

7.2. Training Subsidies

The state-dependent scarring mechanism (Section 3.8) provides a natural policy target: if accelerated human capital depreciation during recessions is the source of amplified persistence and distributional costs, a training subsidy that directly offsets this depreciation should reduce both. I first evaluate four scenarios using nonlinear simulations of a 1pp separation shock, then extend the analysis to Great Recession severity (3pp) below.

1. **Baseline** ($\kappa = 0$, $\tau = 0$): constant depreciation, no subsidy.
2. **Amplified** ($\kappa = 2$, $\tau = 0$): state-dependent scarring, no subsidy.
3. **Partial subsidy** ($\kappa = 2$, $\tau = 0.5$): scarring with half offset.
4. **Full subsidy** ($\kappa = 2$, $\tau = 1$): scarring fully offset.

Because the steady state is invariant to (κ, τ) (Section 3.8), all four scenarios share the same steady state; only the dynamic Jacobians differ.

ALMP context. The training subsidy represents a reduced-form stand-in for active labor market programs (ALMPs) that preserve or rebuild human capital during downturns, including employer-side training grants, apprenticeship subsidies, and worker retraining programs. The meta-analysis of Card, Kluve, and Weber (2018) finds that training programs are among the most effective ALMPs, with positive employment effects that grow over time (in contrast to job search assistance, which has only short-run effects). OECD countries typically spend 0.5–1.5% of GDP on ALMPs, with spending reaching 1.5–2% in Nordic countries during severe recessions. In the model, the peak fiscal cost of a full subsidy ($\tau = 1$) for a 1pp separation shock is $p_{\text{train}} \cdot \tau \cdot \hat{s}_{\text{peak}} = 0.5 \cdot 1 \cdot 0.2 = 0.1$ per quarter, or approximately 6.1% of steady-state output, well above observed ALMP spending. This suggests that $p_{\text{train}} = 0.5$ overstates the fiscal cost, or alternatively, that the full subsidy ($\tau = 1$) is unrealistically generous. The partial subsidy ($\tau = 0.5$) at a lower fiscal scaling would be more comparable to observed programs. In Appendix I, I report sensitivity to p_{train}.

Distributional results. Table 6 reports the CE welfare cost by liquid wealth group across the four scenarios. Table 7 decomposes the subsidy effect into its human capital preservation and fiscal cost components using a zero-cost subsidy variant ($\tau = 1$, $p_{\text{train}} = 0$) that preserves human capital without imposing any fiscal burden.

Table 6: *Training Subsidies: CE Welfare (%) by Liquid Wealth Group for a 1pp Separation Shock (AR(1), $\rho = 0.9$). Panel A: baseline fiscal cost ($p_{train} = 0.5$); Panel B: realistic fiscal cost ($p_{train} = 0.10$); Panel C: difference from constant-depreciation baseline; Panel D: improvement from amplified scarring*

	Aggregate	HtM ($b = 0$)	Low liquid	Mid liquid	High liquid
Panel A: CE welfare (%), baseline fiscal cost ($p_{train} = 0.5$)					
Baseline	-0.57	-0.53	-0.56	-0.42	-0.46
Amplified	-1.02	-1.13	-1.09	-0.67	-0.84
Partial subsidy	-0.69	-0.48	-0.97	-0.59	-0.73
Full subsidy	-0.36	$+0.21$	-0.86	-0.52	-0.63
Panel B: CE welfare (%), realistic fiscal cost ($p_{train} = 0.10$)					
Realistic full ($\tau = 1$)	-0.67	-0.59	-0.79	-0.52	-0.61
Realistic partial ($\tau = 0.5$)	-0.85	-0.85	-0.94	-0.59	-0.72
Panel C: Difference from Baseline ($\kappa = 0$, pp)					
Amplified	-0.45	-0.60	-0.53	-0.24	-0.39
Full sub ($p = 0.5$)	$+0.22$	$+0.74$	-0.31	-0.09	-0.17
Realistic full ($p = 0.10$)	-0.10	-0.06	-0.23	-0.09	-0.15
Panel D: Improvement from Amplified ($\kappa = 2$, pp)					
Full sub ($p = 0.5$)	$+0.66$	$+1.34$	$+0.23$	$+0.15$	$+0.21$
Realistic full ($p = 0.10$)	$+0.34$	$+0.54$	$+0.30$	$+0.15$	$+0.23$

Notes: 1pp separation shock, AR(1) with $\rho = 0.9$, nonlinear simulation. CE variation computed over 60 quarters. The subsidy costs $p_{\text{train}} \cdot \tau \cdot \hat{s}_t$ in additional lump-sum taxes (Eq. 21). Panel A uses the baseline fiscal scaling ($p_{\text{train}} = 0.5$, peak cost $\approx 6.1\%$ of output). Panel B uses $p_{\text{train}} = 0.10$ (peak cost $\approx 1.2\%$ of output), consistent with observed ALMP spending in OECD countries during severe recessions. Panel C reports differences from the constant-depreciation baseline ($\kappa = 0$). Panel D reports improvements relative to the amplified scenario ($\kappa = 2$, no subsidy).

Table 7: *Decomposition of Training Subsidy Effect: HC Preservation vs. Fiscal/GE Channel for a 1pp Separation Shock (AR(1), $\rho = 0.9$) with $\kappa = 2$*

	Aggregate	HtM ($b = 0$)	Low liquid	Mid liquid	High liquid
CE welfare (%)					
Amplified ($\kappa = 2$, no sub)	-1.02	-1.13	-1.09	-0.67	-0.84
Zero-cost sub ($\tau = 1$, $p = 0$)	-0.75	-0.77	-0.76	-0.52	-0.60
Full sub ($\tau = 1$, $p = 0.5$)	-0.36	$+0.21$	-0.86	-0.52	-0.63
Decomposition of full subsidy effect (pp, relative to Amplified)					
HC preservation	$+0.27$	$+0.36$	$+0.33$	$+0.15$	$+0.24$
Fiscal/GE channel	$+0.39$	$+0.98$	-0.10	-0.00	-0.03
Total	$+0.66$	$+1.34$	$+0.23$	$+0.15$	$+0.21$

Notes: HC preservation = CE(zero-cost sub) − CE(amplified): effect of reduced depreciation without fiscal burden. Fiscal/GE channel = CE(full sub) − CE(zero-cost sub): effect of adding the fiscal cost and its general equilibrium incidence. The zero-cost subsidy ($p_{\text{train}} = 0$) preserves human capital identically to the full subsidy but imposes no taxes.

The results reveal a sharp distributional divide in the efficacy of training subsidies.

HtM households are most affected by amplified scarring. Under amplified scarring without a subsidy, HtM welfare worsens from -0.53% to -1.13%, the largest deterioration of any group (-0.60pp). This reflects the triple vulnerability of HtM households: they face the highest unemployment rate, the most cumulative human capital depreciation (from spending more time unemployed), and upon re-employment earn lower wages through the $w(h)$ channel. The full training subsidy more than reverses this: HtM CE improves from -1.13% to $+0.21\%$, a gain of 1.34pp that leaves HtM households with a *net welfare gain* relative to the constant-depreciation baseline.

Decomposition. Table 7 decomposes this gain into two channels using a zero-cost subsidy variant ($\tau = 1$, $p_{\text{train}} = 0$) that preserves human capital without any fiscal burden. The HC preservation channel accounts for $+0.36$pp of the HtM gain: by reducing the *amplified* depreciation (from $\sim 1.4\delta_h$ at the shock's peak back toward $\sim 1.12\delta_h$), the subsidy preserves human capital that translates directly into higher re-employment wages. This channel benefits all groups ($+0.15$ to $+0.36$pp), with HtM benefiting most because they spend more time unemployed. The remaining $+0.98$pp for HtM comes from the fiscal/GE channel: lump-sum taxes to finance the subsidy (Eq. 21) tighten the liquid bond market, raise the real

rate, and redistribute consumption toward income-dependent households.[9] This fiscal/GE channel is mildly negative for other groups (-0.10pp for low-liquid, near zero for mid and high), confirming that the +0.21% HtM gain is driven by progressive GE incidence of the lump-sum financing, not by below-baseline depreciation.

Fiscal incidence and the role of lump-sum taxation. The decomposition raises an important question: is the headline HtM result primarily about training or about fiscal redistribution? Under the baseline fiscal scaling ($p_{\text{train}} = 0.5$), the fiscal/GE channel accounts for $+0.98$pp of the $+1.34$pp total HtM improvement, approximately 73%. This channel operates through the lump-sum tax that finances the subsidy: higher taxes tighten the liquid bond market, raise the equilibrium real rate, and redistribute consumption toward income-dependent (and therefore liquidity-constrained) households. Importantly, this fiscal incidence depends on the assumption of lump-sum financing. Under proportional income taxation, the fiscal burden would fall more heavily on high-income households, potentially reducing or reversing the progressive GE incidence. The training subsidy should therefore be understood as operating through *two* complementary channels: a direct HC preservation channel that benefits all groups but especially HtM workers ($+0.36$pp), and a fiscal/GE channel whose distributional incidence depends on the tax instrument. The HC preservation channel is robust to the financing assumption; the fiscal/GE channel is not. At the realistic fiscal scaling ($p_{\text{train}} = 0.10$), the fiscal burden is five times smaller and the relative contribution of the two channels shifts toward HC preservation, making the result less sensitive to the tax assumption.

Training vs. untargeted transfers. To test whether the training subsidy's distributional benefit is specific to HC preservation or could be replicated by generic fiscal spending, I compare it against a pure lump-sum transfer of equal fiscal magnitude. The transfer scenario keeps amplified scarring at full strength ($\kappa = 2$, $\tau = 0$) but reduces lump-sum taxes by the same amount the training subsidy costs, giving every household extra disposable income during the recession ($T_t = r \cdot B^g - p \cdot \hat{s}_t$ with $p = 0.10$). Table 8 reports the results at both 1pp and 3pp shock sizes.

The pure transfer *widens* the distributional gradient rather than narrowing it. At 3pp, the gradient moves from -1.49pp (Amplified) to -1.96pp (Pure transfer), compared to

[9] The full subsidy does not reduce effective depreciation below the constant baseline δ_h. With $\kappa = 2$ and $\tau = 1$, $\delta^{\text{eff}}_{h,t} = \delta_h(1 + 2\hat{s})(1 - \hat{s}) = \delta_h(1 + \hat{s} - 2\hat{s}^2)$. This expression exceeds δ_h for $\hat{s} \in [0, 0.5]$; for $\hat{s} > 0.5$ (separation rate shocks exceeding 50% of steady state) it could fall below δ_h, but the max operator in Eq. (20) ensures $\delta^{\text{eff}}_{h,t} \geq \delta_h$ regardless. For all shock sizes considered in this paper ($\hat{s} \leq 0.8$), the floor binds only at the largest shocks. The net welfare gain relative to the $\kappa = 0$ baseline therefore requires the GE channel documented in Table 7.

+0.10pp under the training subsidy. HtM households are worse off with the transfer (−4.57%) than with no intervention at all (−4.07%). The mechanism is general equilibrium: the transfer boosts aggregate consumption demand, tightening the bond market and raising the real rate. Higher real rates benefit households with savings (mid- and high-liquid groups earn more on their assets) but provide no relief to HtM households, who hold zero liquid wealth. Meanwhile, HtM households continue to suffer full amplified scarring, losing human capital that depresses their job-finding rates and re-employment wages. The transfer addresses the symptom (low consumption) through a channel that bypasses the most vulnerable, while leaving the underlying cause (skill loss) untreated.

This result demonstrates that the training subsidy's distributional benefit is entirely specific to its HC preservation mechanism. The 73% "fiscal/GE channel" identified in the decomposition above is not generic fiscal redistribution; it reflects the general equilibrium consequences of preserving human capital specifically. When the same fiscal resources are deployed as an untargeted transfer instead, the GE effects flow to asset-holders rather than to the liquidity-constrained. This also explains the limited effectiveness of monetary policy (Section 7.1): interest rate channels primarily affect interest-sensitive, asset-holding households, not the HtM households who bear the largest scarring costs.

Table 8: *Training Subsidy vs. Pure Lump-Sum Transfer: CE Welfare (%) by Liquid Wealth Group*

	Aggregate	HtM ($b = 0$)	Low liquid	Mid liquid	High liquid
Panel A: 1pp separation shock					
Amplified ($\kappa = 2$)	−1.02	−1.13	−1.09	−0.67	−0.84
Training subsidy ($\tau = 1$, $p = 0.10$)	−0.67	−0.59	−0.79	−0.52	−0.61
Pure transfer ($p = 0.10$)	−1.09	−1.31	−1.06	−0.67	−0.84
Panel B: 3pp separation shock (Great Recession)					
Amplified ($\kappa = 2$)	−3.24	−4.07	−2.95	−2.27	−2.57
Training subsidy ($\tau = 1$, $p = 0.10$)	−1.71	−1.49	−1.52	−1.52	−1.59
Pure transfer ($p = 0.10$)	−3.48	−4.57	−2.95	−2.33	−2.61
Distributional gradient (HtM − High liquid, pp)					
	1pp			3pp	
Amplified	−0.29			−1.49	
Training subsidy	+0.02			+0.10	
Pure transfer	−0.47			−1.96	

Notes: Two-asset model with amplified scarring ($\kappa = 2$). Training subsidy: $\tau = 1$, $p_{\text{train}} = 0.10$, taxes rise during recession ($T = rB^g + p \cdot \hat{s}$), HC depreciation offset. Pure transfer: $\tau = 0$, $p = 0.10$, taxes *fall* during recession ($T = rB^g - p \cdot \hat{s}$), no HC effect. Both have the same fiscal magnitude ($0.10 \cdot \hat{s}$). AR(1) shock with $\rho = 0.9$. CE computed over 60 quarters.

Realistic fiscal cost. Panel B of Table 6 reports results under an empirically grounded fiscal scaling ($p_{\text{train}} = 0.10$), which implies a peak fiscal cost of approximately 1.2% of output, within the range of OECD ALMP spending during severe recessions. At this realistic cost, a full subsidy improves HtM welfare from -1.13% to -0.59% ($+0.54$pp), the largest gain of any group (Panel D). Relative to the constant-depreciation baseline, the remaining HtM cost is only -0.06pp (Panel C), meaning the realistic subsidy offsets 90% of the amplified scarring cost (-0.60pp) for the most vulnerable group at an empirically feasible fiscal scale.

The subsidy improvement is overwhelmingly concentrated in HtM households ($+1.34$pp vs. $+0.15$ to $+0.23$pp for other groups), revealing the complementarity between liquidity constraints and the scarring channel. Low-liquid workers face the second-largest losses under amplified scarring (-1.09%) and benefit from HC preservation ($+0.33$pp), but the fiscal/GE channel is slightly negative for non-HtM groups. Figure 10 visualizes these results.

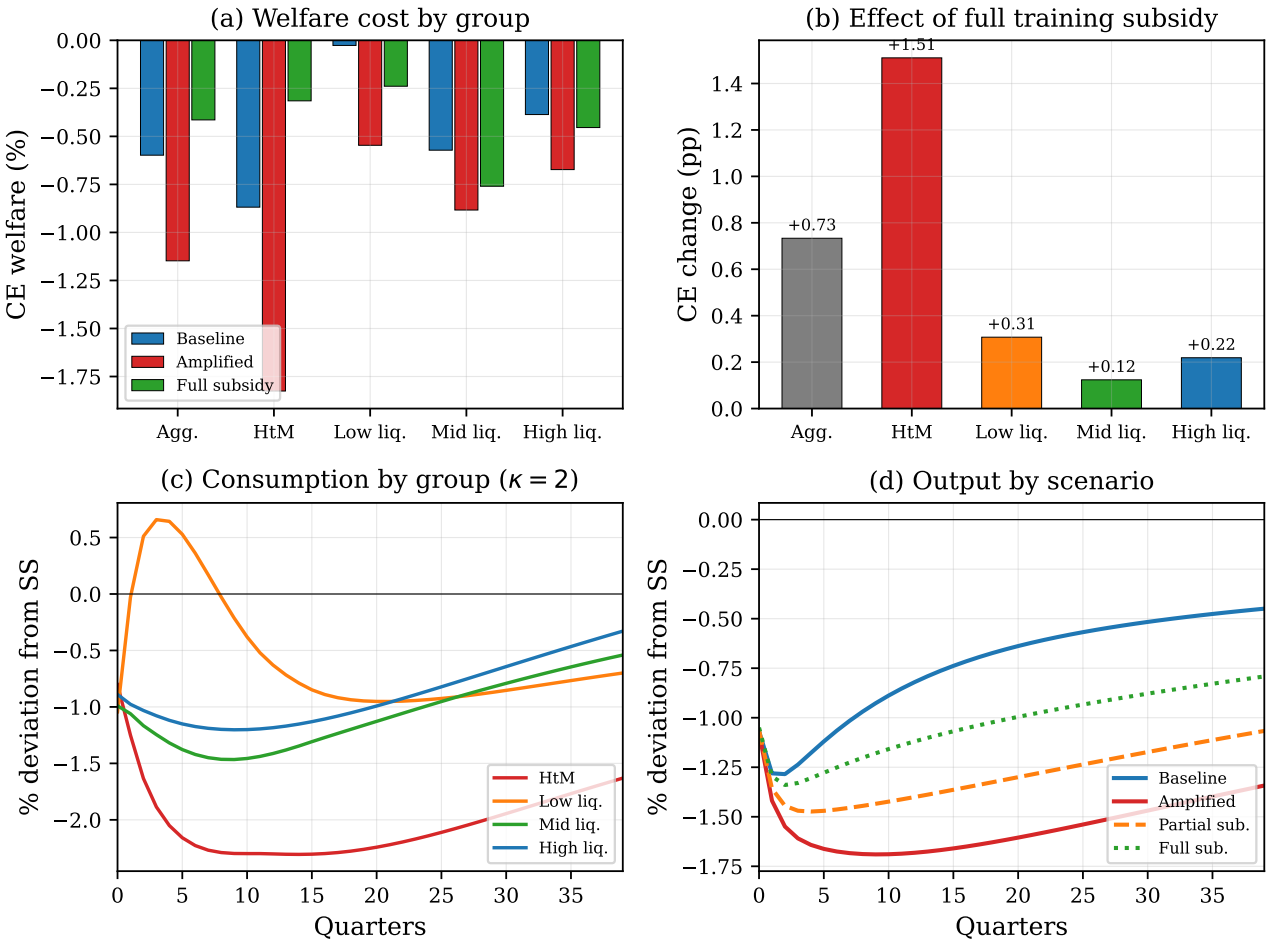

Figure 10: *Training subsidies: distributional welfare for a 1pp separation shock (AR(1), $\rho = 0.9$). (a) CE welfare cost by liquid wealth group across three key scenarios: Baseline ($\kappa = 0$), Amplified ($\kappa = 2$), and Full subsidy ($\kappa = 2$, $\tau = 1$). (b) Effect of full training subsidy: CE change from Amplified to Full subsidy by group. (c) Consumption paths by group under amplified scarring ($\kappa = 2$). (d) Aggregate output paths across all four scenarios. The full training subsidy improves HtM welfare from -1.13% to $+0.21\%$, the largest improvement of any group.*

Implications. The central finding is that training subsidies are most effective for the most vulnerable group. The decomposition in Table 7 shows that this operates through two reinforcing channels: HC preservation (which helps all groups but HtM most, at $+0.36$pp) and progressive GE incidence of the fiscal transfer ($+0.98$pp for HtM). The wage-human capital channel ($\phi_w > 0$) is critical: without it, HtM households' vulnerability stems only from consumption smoothing; with it, their vulnerability compounds through lower re-employment wages. At realistic fiscal cost ($p_{\text{train}} = 0.10$), the subsidy still delivers a $+0.54$pp improvement for HtM, confirming that the mechanism is quantitatively relevant at empirically feasible spending levels. The subsidy addresses a primary driver of the distributional gradient, skill loss, rather than the symptom, making it a natural complement

to liquidity provision. By contrast, UI generosity has limited aggregate leverage: sweeping b_{ui} over $[0.35, 0.60]$ varies the CE welfare cost by only 5 basis points while substantially worsening steady-state unemployment (Table 14 in Appendix G).

Training subsidies at Great Recession severity. The policy relevance of training subsidies is sharpest at 3pp shocks, where the distributional gradient is large enough to be a first-order policy concern. Table 9 reports the results. Under amplified scarring ($\kappa = 2$) without subsidy, the distributional gradient is -1.49pp: HtM CE $= -4.07\%$ versus high-liquid CE $= -2.57\%$. An empirically scaled training subsidy ($p_{\text{train}} = 0.10$, $\tau = 1$, peak fiscal cost $\sim$3.6% of output) reduces HtM losses from -4.07% to -1.49% and nearly eliminates the distributional gradient, reducing it from -1.49pp to $+0.10$pp. The zero-cost subsidy decomposition ($\tau = 1$, $p_{\text{train}} = 0$) yields intermediate results: HtM CE $= -2.06\%$, gradient $= -0.48$pp. This confirms that at Great Recession severity, approximately two-thirds of the gradient reduction comes from the HC preservation channel, with the remainder from fiscal/GE redistribution.[10]

Table 9: *Training Subsidies at Great Recession Severity: CE Welfare (%) by Liquid Wealth Group for a 3pp Separation Shock (AR(1), $\rho = 0.9$) with $\kappa = 2$*

	Aggregate	HtM ($b = 0$)	Low liquid	Mid liquid	High liquid
CE welfare (%)					
Baseline ($\kappa = 0$)	-1.64	-1.59	-1.13	-1.36	-1.36
Amplified ($\kappa = 2$)	-3.24	-4.07	-2.95	-2.27	-2.57
Zero-cost sub ($\tau = 1$, $p = 0$)	-1.95	-2.06	-1.47	-1.52	-1.58
Realistic sub ($\tau = 1$, $p = 0.10$)	-1.71	-1.49	-1.52	-1.52	-1.59
Distributional gradient (HtM − High liquid, pp)					
Amplified ($\kappa = 2$)			-1.49		
Zero-cost subsidy			-0.48		
Realistic subsidy			$+0.10$		

Notes: 3pp separation shock. Realistic subsidy uses $p_{\text{train}} = 0.10$, $\tau = 1$ (peak fiscal cost $\approx$ 3.6% of output at 3pp). The full subsidy ($p_{\text{train}} = 0.5$) does not converge at 3pp.

These welfare gains for HtM households are a conservative estimate. The training subsidy is financed by lump-sum taxes (Eq. 21), which are regressive in consumption-

[10]The full subsidy ($\tau = 1$, $p_{\text{train}} = 0.5$) fails to converge at 3pp shock sizes, as the combination of a large separation shock and the fiscal cost ($\sim$18% of output at peak) pushes the equilibrium outside the nonlinear solver's convergence region. This further supports the view that $p_{\text{train}} = 0.5$ is unrealistically large; the realistic calibration ($p_{\text{train}} = 0.10$) converges without difficulty and is within the range of observed ALMP spending.

equivalent terms; low-consumption HtM households bear a proportionally larger tax burden. That HtM households benefit substantially from the subsidy *despite* this regressive financing implies that the subsidy's gross human capital benefit to HtM households is even larger than the net CE improvement suggests. Under progressive financing (e.g., proportional income taxes), the distributional advantage of training subsidies would be more pronounced.

8. Conclusion

This paper has built a two-asset HANK model with search-and-matching unemployment, endogenous human capital dynamics, wage scarring, and state-dependent depreciation to study the distributional incidence of recessions and evaluate training subsidies as a targeted policy response. The scarring mechanism, through which human capital depreciation reduces both job-finding rates and re-employment wages, is calibrated to match the empirical earnings loss profile of displaced workers. The title, "HANKs for the Memories," captures the central insight: the individual labor market *memories* that recessions leave behind are distributed unevenly across the wealth distribution, and this distributional incidence is where both the economic action and the policy leverage reside.

The distributional costs of scarring are *convex* in shock size: the HtM-to-high-liquid welfare gradient is -0.3pp in an average recession (1pp) but -1.5pp in a Great Recession (3pp), approximately five times larger and far exceeding the linear prediction. HtM households suffer CE losses of -4.1% at 3pp, nearly 60% larger than the high-liquid group's -2.6%, and a cross-sectional decomposition confirms this is driven by the intensive (within-cell) margin, not compositional reclassification. An empirically scaled training subsidy nearly eliminates this gradient (from -1.5pp to $+0.1$pp) at feasible fiscal cost, while an untargeted lump-sum transfer of equal magnitude *widens* the gradient, confirming that the benefit is specific to human capital preservation. Monetary policy has limited leverage over these labor market shocks: interest rate channels primarily benefit asset-holding households, not the HtM households who bear the largest scarring costs. The distributional incidence of training subsidies suggests that combining them with liquidity provision could offer more uniform recession insurance, though this remains a conjecture for future work.

Extensions. Several directions remain. First, extending the distributional analysis to UI generosity, decomposing the welfare effects of higher b_{ui} across liquid wealth groups, would complete the policy comparison. Second, adding endogenous search intensity, where liquidity-constrained workers search less effectively, would compound the scarring vulnerability of low-liquid workers. Third, finite UI duration and the interaction of benefit

expiry with human capital depreciation are natural extensions of the training subsidy analysis. Fourth, the meta-analytic evidence of Card et al. (2018) suggests that different types of ALMPs have different time profiles of effectiveness; modeling the subsidy with an explicit human capital production function (rather than as a depreciation offset) would allow richer policy comparisons.

References

Acabbi, E. M., A. Alati, and L. Mazzone (2024). Human capital ladders, cyclical sorting, and hysteresis. Bank of England Staff Working Paper No. 1077; R&R at AER.

Acharya, S., J. Bengui, K. Dogra, and S. L. Wee (2022). Slow recoveries and unemployment traps: Monetary policy in a time of hysteresis. *Economic Journal 132*(646), 2007–2047.

Alves, F. and G. L. Violante (2024). From micro to macro hysteresis: Long-run effects of monetary policy. Bank of Canada Staff Working Paper 2024-39.

Alves, F. and G. L. Violante (2025). Monetary policy under Okun's hypothesis. NBER Working Paper No. 33488.

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

Auclert, A., M. Rognlie, and L. Straub (2024). The intertemporal Keynesian cross. *Journal of Political Economy 132*(12).

Auclert, A., M. Rognlie, and L. Straub (2025). Fiscal and monetary policy with heterogeneous agents. *Annual Review of Economics*.

Baley, I., A. Figueiredo, and R. Ulbricht (2022). Mismatch cycles. *Journal of Political Economy 130*(11), 2970–3012.

Bayer, C., B. Born, and R. Luetticke (2024). Shocks, frictions, and inequality in US business cycles. *American Economic Review 114*(5), 1211–1247.

Bilbiie, F. O. (2020). The new Keynesian cross. *Journal of Monetary Economics 114*, 90–108.

Bilbiie, F. O. (2024). Monetary policy and heterogeneity: An analytical framework. *Review of Economic Studies 92*(4), 2398–2436.

Blanchard, O. J. and L. H. Summers (1986). Hysteresis and the European unemployment problem. *NBER Macroeconomics Annual 1*, 15–78.

Broer, T., J. Druedahl, K. Harmenberg, and E. Öberg (2025). The unemployment-risk channel in business-cycle fluctuations. *International Economic Review*.

Browning, M. and T. F. Crossley (2001). Unemployment insurance benefit levels and consumption changes. *Journal of Public Economics 80*(1), 1–23.

Card, D., R. Chetty, and A. Weber (2007). Cash-on-hand and competing models of intertemporal behavior: New evidence from the labor market. *Quarterly Journal of Economics 122*(4), 1511–1560.

Card, D., J. Kluve, and A. Weber (2018). What works? A meta analysis of recent active labor market program evaluations. *Journal of the European Economic Association 16*(3), 894–931.

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

Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy 116*(2), 173–234.

Consolo, A. and M. Hansel (2024). HANK faces unemployment. ECB Working Paper No. 2953.

Davis, S. J. and T. von Wachter (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity 2011*(2), 1–72.

Fallick, B. and P. Krolikowski (2018). Hysteresis in employment among disadvantaged workers. Federal Reserve Bank of Cleveland Working Paper No. 18-01.

Farrell, D., P. Ganong, F. Greig, and P. Noel (2019). Recovering from job loss: The role of unemployment insurance. Report, JPMorgan Chase Institute.

Fontaine, F., J. N. Jensen, and R. Vejlin (2024). Wealth, portfolios, and nonemployment duration. *Journal of Money, Credit and Banking 56*(7), 1861–1886.

Fornaro, L. and M. Wolf (2023). The scars of supply shocks: Implications for monetary policy. *Journal of Monetary Economics 140*, S18–S36.

Galí, J. (2022). Insider-outsider labor markets, hysteresis and monetary policy. *Journal of Money, Credit and Banking*.

Ganong, P., D. Jones, P. J. Noel, F. E. Greig, D. Farrell, and C. Wheat (2025). Liquid wealth, debt, and the consumption response to typical labor income shocks. NBER Working Paper 27552, R&R at American Economic Review.

Ganong, P. and P. Noel (2019). Consumer spending during unemployment: Positive and normative implications. *American Economic Review 109*(7), 2383–2424.

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

Gornemann, N., K. Kuester, and M. Nakajima (2021). Doves for the rich, hawks for the poor? Distributional consequences of systematic monetary policy. ECB Working Paper No. 2588.

Graves, S. (2025). Does unemployment risk affect business cycle dynamics? *American Economic Journal: Macroeconomics 17*(2), 65–100.

Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review 87*(1), 192–205.

Guerrieri, V. and G. Lorenzoni (2017). Credit crises, precautionary savings, and the liquidity trap. *Quarterly Journal of Economics 132*(3), 1427–1467.

Heathcote, J., K. Storesletten, and G. L. Violante (2009). Quantitative macroeconomics with heterogeneous households. *Annual Review of Economics 1*, 319–354.

Heathcote, J., K. Storesletten, and G. L. Violante (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review 104*(7), 2075–2126.

Huckfeldt, C. (2022). Understanding the scarring effect of recessions. *American Economic Review 112*(4), 1273–1310.

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan (1993). Earnings losses of displaced workers. *American Economic Review 83*(4), 685–709.

Jarosch, G. (2023). Searching for job security and the consequences of job loss. *Econometrica 91*(3), 855–898.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to HANK. *American Economic Review 108*(3), 697–743.

Kaplan, G., G. L. Violante, and J. Weidner (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity 2014*(1), 77–138.

Kekre, R. (2023). Unemployment insurance in macroeconomic stabilization. *Review of Economic Studies 90*(5), 2439–2480.

Kolsrud, J., C. Landais, P. Nilsson, and J. Spinnewijn (2018). The optimal timing of unemployment benefits: Theory and evidence from Sweden. *American Economic Review 108*(4–5), 985–1033.

Kroft, K., F. Lange, and M. J. Notowidigdo (2013). Duration dependence and labor market conditions: Evidence from a field experiment. *Quarterly Journal of Economics 128*(3), 1123–1167.

Krolikowski, P. (2017). Job ladders and earnings of displaced workers. *American Economic Journal: Macroeconomics 9*(2), 1–31.

Krueger, D., K. Mitman, and F. Perri (2016). Macroeconomics and household heterogeneity. *Handbook of Macroeconomics 2*, 843–921.

Lentz, R. (2009). Optimal unemployment insurance in an estimated job search model with savings. *Review of Economic Dynamics 12*(1), 37–57.

Lise, J. and J.-M. Robin (2017). The macrodynamics of sorting between workers and firms. *American Economic Review 107*(4), 1104–1135.

Ljungqvist, L. and T. J. Sargent (1998). The European unemployment dilemma. *Journal of Political Economy 106*(3), 514–550.

McKay, A. and R. Reis (2021). Optimal automatic stabilizers. *Review of Economic Studies 88*(5), 2375–2406.

Mercan, Y. and J. Nichols (2025). Understanding the role of wealth in worker flows. *Federal Reserve Bank of Kansas City Economic Review*.

Oreopoulos, P., T. von Wachter, and A. Heisz (2012). The short- and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics 4*(1), 1–29.

Pissarides, C. A. (1992). Loss of skill during unemployment and the persistence of employment shocks. *Quarterly Journal of Economics 107*(4), 1371–1391.

Ravn, M. O. and V. Sterk (2021). Macroeconomic fluctuations with HANK & SAM: An analytical approach. *Journal of the European Economic Association 19*(2), 1162–1202.

Schmieder, J. F., T. von Wachter, and S. Bender (2016). The effect of unemployment benefits and nonemployment durations on wages. *American Economic Review 106*(3), 739–777.

Schmieder, J. F., T. von Wachter, and J. Heining (2023). The costs of job displacement over the business cycle and its sources: Evidence from Germany. *American Economic Review 113*(5), 1208–1254.

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

A. Baseline Scarring Comparison

To isolate the role of constant-rate scarring, I solve the one-asset model at $\zeta = 0$ (no h-dependent finding rates) and $\zeta = 0.5$ (baseline scarring), both at $B = 7$ for comparability. The $\zeta = 0$ model does not converge at $B = 9$ because removing h-dependent finding rates changes the ergodic distribution of human capital and savings demand.

The aggregate effects are small: $\zeta = 0.5$ adds modestly to cumulative output loss and CE welfare cost relative to $\zeta = 0$. Figure 11 overlays the impulse responses; the differences are visible but modest. This near-invariance at constant depreciation rates motivates the state-dependent scarring extension in Section 6.3.

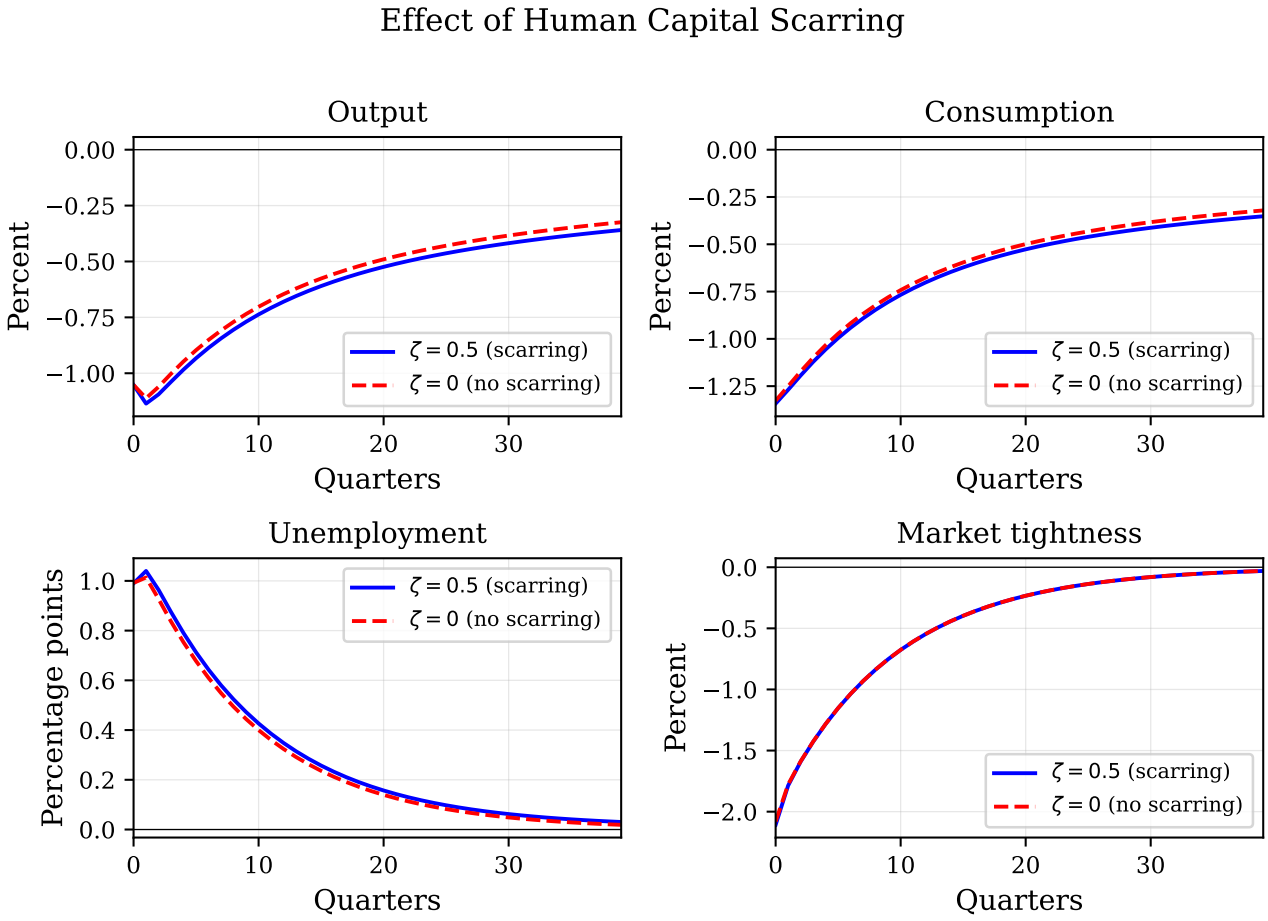

Figure 11: *Scarring comparison for a 1pp separation shock (AR(1), $\rho = 0.9$) in the one-asset model: $\zeta = 0.5$ (solid blue) vs. $\zeta = 0$ (dashed red), both at $B = 7$ with constant depreciation ($\kappa = 0$). The h-dependent finding rate generates modest additional persistence in all variables.*

B. Scarring Heterogeneity at Constant Depreciation

Table 10 and Figure 12 compare the CE welfare cost by liquid wealth group with and without h-dependent finding rates ($\zeta = 0.5$ vs. $\zeta = 0$). The aggregate difference is -0.025pp. Scarring concentrates its costs on HtM workers (-0.083pp) and low-liquid workers (-0.025pp), while the additional cost for mid- and high-liquid groups is smaller (-0.018pp and -0.015pp). Even at constant depreciation, the scarring channel widens the distributional gradient.

Table 10: *Distributional Cost of the Scarring Channel ($\zeta = 0.5$ vs. $\zeta = 0$): CE Welfare (%) by Liquid Wealth Group for a 1pp Separation Shock (AR(1), $\rho = 0.9$) in the Two-Asset Model at Constant Depreciation ($\kappa = 0$)*

	Aggregate	HtM ($b = 0$)	Low liquid	Mid liquid	High liquid
$\zeta = 0.5$ (scarring)	-0.57	-0.53	-0.56	-0.42	-0.46
$\zeta = 0$ (no scarring)	-0.55	-0.44	-0.53	-0.41	-0.44
Difference (pp)	-0.025	-0.083	-0.025	-0.018	-0.015

Notes: CE welfare variation (%) from a 1pp separation shock in the two-asset model. Difference is the additional cost attributable to the scarring channel ($\zeta = 0.5$ minus $\zeta = 0$); negative values indicate scarring makes the group worse off.

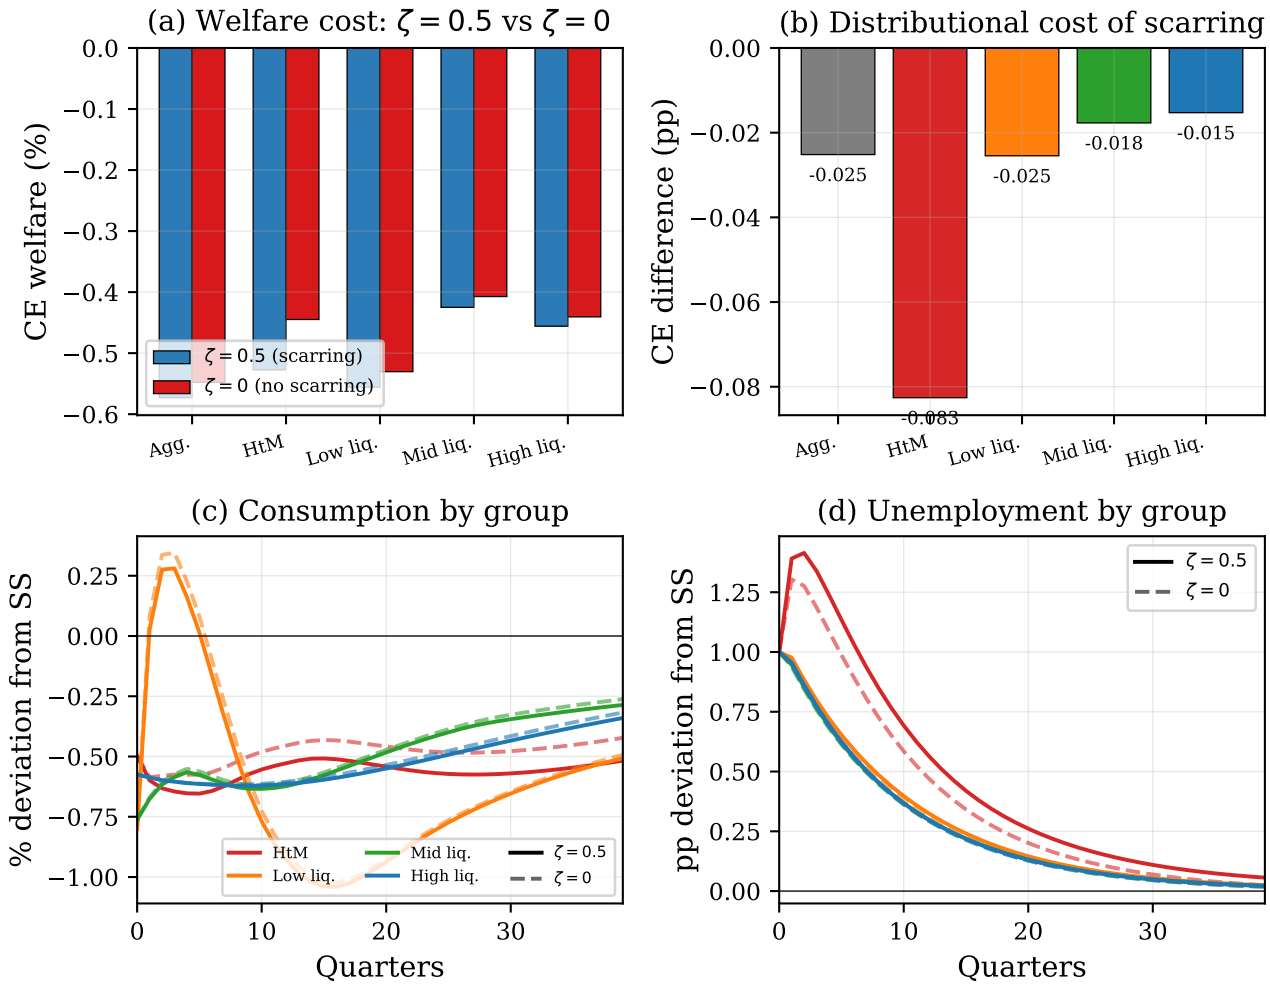

Figure 12: *Distributional cost of the scarring channel for a 1pp separation shock (AR(1), $\rho = 0.9$) in the two-asset model at constant depreciation ($\kappa = 0$). (a) CE welfare cost by liquid wealth group: $\zeta = 0.5$ (blue) vs. $\zeta = 0$ (red). (b) Difference (additional scarring cost). (c)–(d) Consumption and unemployment paths by group: solid lines are $\zeta = 0.5$, dashed are $\zeta = 0$. The differences are quantitatively small.*

C. Match Quality

I extend the model to include match quality m as a discrete state variable. Upon job finding, workers draw match quality $m \sim G(m)$ from a uniform distribution over $[m_{\min}, m_{\max}]$. Match quality persists while employed and is lost upon separation. The learning rate depends on match quality:

$$\phi_0(m) = \phi_0 \cdot \left(\frac{m}{m_{\text{ref}}} \right)^{\phi_1}, \tag{24}$$

where ϕ_1 governs the elasticity of learning with respect to match quality. The state space expands from $2n_H$ joint (status, h) states to $n_H \cdot n_M + n_H$ states. With $n_H = 10$ and $n_M = 4$,

this yields 50 states.

Table 11 reports the results. Even with an aggressive elasticity ($\phi_1 = 3.0$), cumulative output loss changes by only $\sim$5%. The match quality channel operates at the distributional level, reshaping which workers accumulate human capital quickly after re-employment, but washing out in aggregation.

Table 11: *Match Quality Extension: Aggregate Effects of a 1pp Separation Shock (AR(1), $\rho = 0.9$) Comparing Baseline ($n_M = 1$) to Match Quality ($n_M = 4$, $\phi_1 = 0.5$) in the One-Asset Model with NK Pricing*

	Peak response		$t = 12$	
Variable	Baseline	Match quality	Baseline	Match quality
Y (%)	-2.10	-2.10	-1.01	-1.01
C (%)	-2.25	-2.25	-0.93	-0.92
U (pp)	$+1.12$	$+1.12$	$+0.39$	$+0.39$
Cum. Y (40Q)	-0.351	-0.350		

Notes: 1pp separation shock (AR(1), $\rho = 0.9$). Both models use NK pricing, $B = 7$, $\zeta = 0.5$. Baseline: $n_M = 1$. Match quality: $n_M = 4$, $m \in [0.5, 1.5]$, $\phi_1 = 0.5$. Even $\phi_1 = 3.0$ changes cumulative output loss by only $\sim$5%.

D. One-Asset Model Comparison

Aggregate dynamics. Figure 13 compares the one-asset and two-asset model responses to a 1pp separation shock. Output and unemployment paths are nearly identical, but the one-asset consumption drop is roughly 2.5 times as large ($\sim$1.5% vs. 0.6%) because all savings are liquid, and inflation responses have opposite signs. These differences reflect how the liquid-illiquid wealth composition shapes the transmission of labor market shocks.

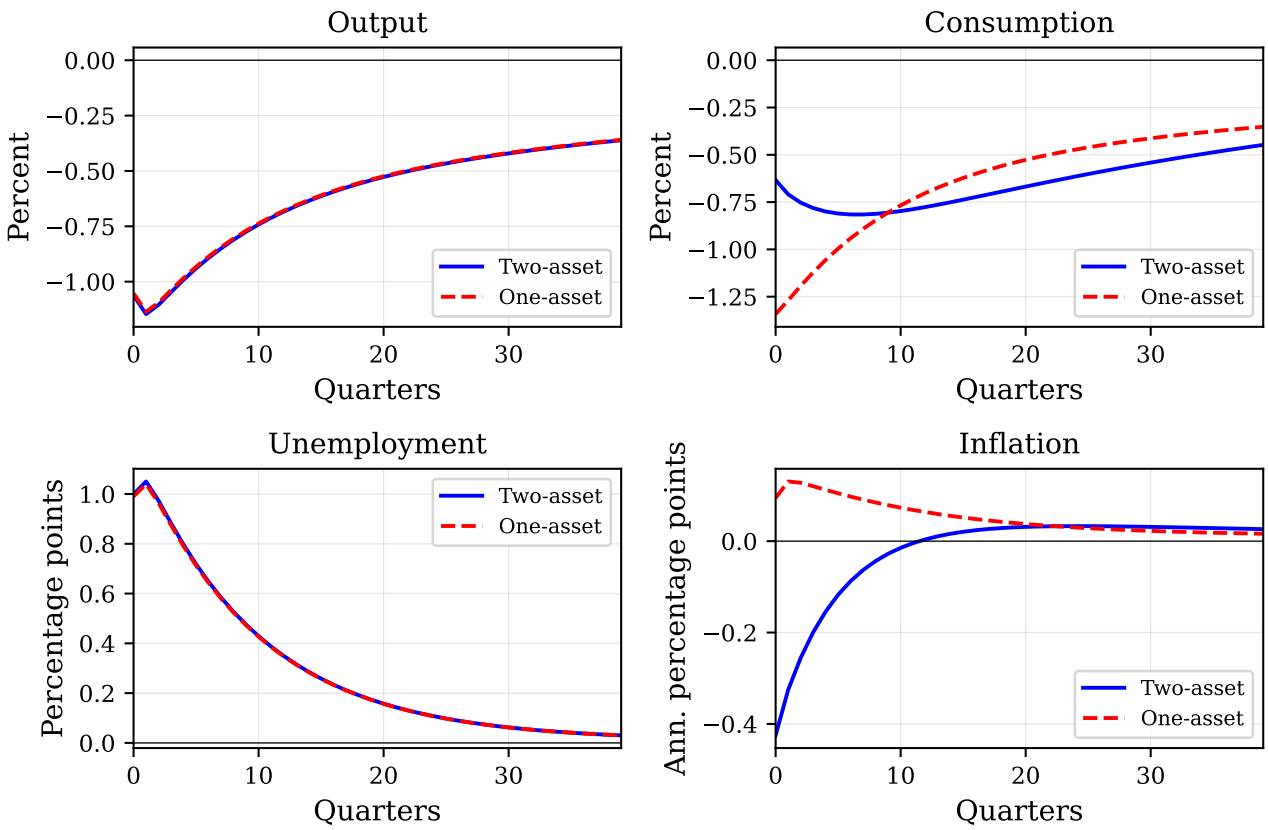

Figure 13: *Comparison of one-asset and two-asset model responses to a 1pp separation shock (AR(1), $\rho = 0.9$) with constant depreciation ($\kappa = 0$). Output and unemployment are nearly identical, but the consumption response is roughly 2.5 times as large in the one-asset model and inflation responses have opposite signs. Both models use $\zeta = 0.5$.*

Distributional results. Table 12 reports the distributional incidence in the one-asset model using wealth quartiles. The distributional gradient is present but more moderate than in the two-asset model: the bottom quartile suffers a CE loss of -0.66% (1.4 times the aggregate), while the top quartile bears -0.35%. The bottom-to-top ratio of $1.9\times$ indicates meaningful but less extreme heterogeneity, precisely because the one-asset framework cannot generate wealthy hand-to-mouth households.

Table 12: *Distributional Incidence: CE Welfare (%) by Wealth Quartile for a 1pp Separation Shock (AR(1), $\rho = 0.9$) in the One-Asset Model ($\zeta = 0.5$, $\kappa = 0$, $B = 7$)*

	Aggregate	Q1 (poorest)	Q2	Q3	Q4 (wealthiest)
CE welfare (%)	−0.47	−0.66	−0.52	−0.42	−0.35
Q/Agg ratio	1.0	1.4	1.1	0.9	0.7
SS U rate (%)	6.6	10.2	6.8	5.4	4.1

Notes: One-asset model at $B = 7$, $\zeta = 0.5$, 1pp separation shock. CE computed over 60 quarters. Quartiles defined by steady-state asset distribution.

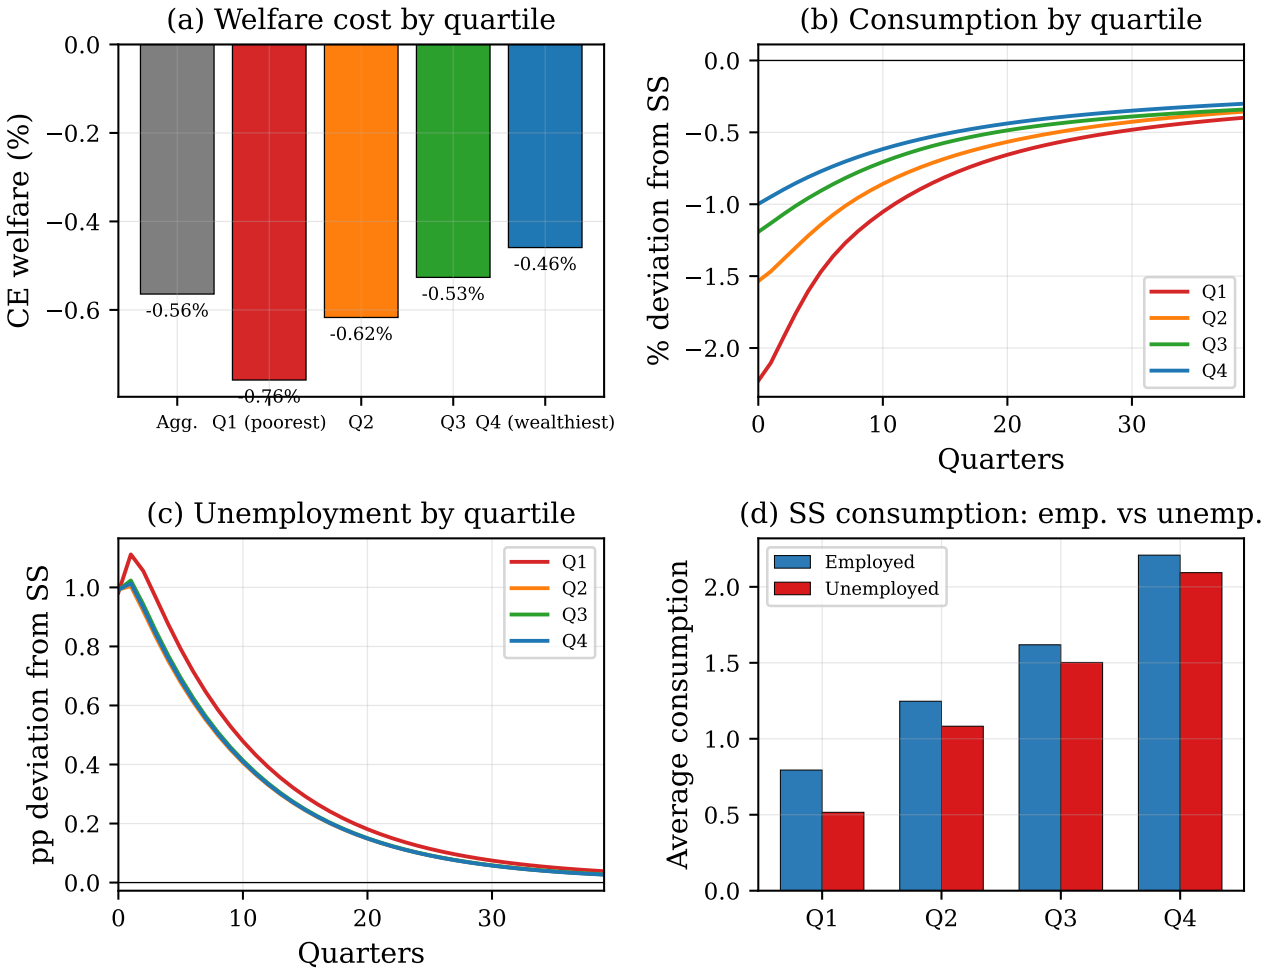

Figure 14: *Distributional incidence of a 1pp separation shock (AR(1), $\rho = 0.9$) in the one-asset model ($\zeta = 0.5$, $\kappa = 0$, $B = 7$). (a) CE welfare cost by wealth quartile. (b) Consumption paths by quartile (percent deviation from SS). (c) Unemployment by quartile (pp deviation). (d) SS consumption levels for employed vs. unemployed workers by quartile.*

Direct comparison. Table 13 compares distributional welfare across the one-asset and two-asset models under identical shocks. Under constant depreciation ($\kappa = 0$), the one-asset gradient across wealth quartiles (-0.31pp) is larger because Q1 bundles all low-wealth households together, while the HtM group captures only the liquidity margin. Under amplified scarring ($\kappa = 2$), the ranking reverses: the HtM grouping isolates the scarring vulnerability channel more precisely than wealth quartiles.

Table 13: *One-Asset vs. Two-Asset Distributional Welfare for a 1pp Separation Shock (AR(1), $\rho = 0.9$)*

	One-asset (quartiles)		Two-asset (liquid wealth)	
	Q1 (poorest)	Q4 (wealthiest)	HtM ($b = 0$)	High liquid
Constant depreciation ($\kappa = 0$)				
CE welfare (%)	-0.66	-0.35	-0.53	-0.46
Gradient (pp)	-0.31[a]		-0.07	
Ratio (worst/best)	1.9×		1.2×	
Amplified scarring ($\kappa = 2$)				
CE welfare (%)	-0.92	-0.70	-1.13	-0.84
Gradient (pp)	-0.22		-0.29	
Ratio (worst/best)	1.3×		1.3×	

[a] The one-asset gradient under $\kappa = 0$ is larger because Q1 includes all low-wealth households (poor and liquidity-constrained alike), while the HtM group captures only the liquidity margin.

Notes: 1pp separation shock, AR(1) with $\rho = 0.9$. One-asset model uses $\phi_w = 0$, $B = 7$; two-asset model uses $\phi_w = 0.5$, $B_g = 0.75$. CE variation computed over 60 quarters. Gradient = worst group CE minus best group CE.

E. Linearity

A nonlinear perfect-foresight analysis confirms that the linearized model provides an excellent approximation of *aggregate* impulse responses under *constant depreciation* ($\kappa = 0$). For separation shocks ranging from 0.5 to 8 percentage points, the nonlinear aggregate IRFs (output, unemployment, consumption) are virtually identical to the linear ones. This reflects the slow h depreciation rate: even in a severe recession, very few workers reach the region of the state space where finding rates drop sharply.

This linearity result applies to aggregate dynamics under constant depreciation, not to distributional welfare or state-dependent scarring. The distributional analysis in Section 6.2 and the policy experiments in Section 7 use full nonlinear simulations precisely because the distributional gradient is convex in shock size, a nonlinearity that operates through the

cross-sectional distribution rather than through aggregate pricing or matching equations. The nonlinear simulations in Section 7 solve for the full distribution path at each shock size, capturing the interaction between state-dependent depreciation and endogenous compositional shifts across the wealth distribution.

F. Taylor Rule Grid Search

I search over Taylor rule coefficients (ϕ_π, ϕ_y) to find the rule that minimizes a standard New Keynesian quadratic loss function:

$$\mathcal{L} = \sum_{t=0}^{H-1} \left[w_\pi \, \pi_t^2 + w_y \left(\frac{Y_t - Y^{ss}}{Y^{ss}} \right)^2 \right], \tag{25}$$

where $w_\pi = w_y = 1$ and $H = 60$ quarters. The grid covers $\phi_\pi \in \{1.1, 1.25, 1.5, 2.0, 3.0\}$ and $\phi_y \in \{0.0, 0.0625, 0.125, 0.25, 0.5\}$. Since the steady state is independent of (ϕ_π, ϕ_y), I solve it once and vary only the Taylor rule coefficients.

For separation shocks, the quadratic loss surface is extremely flat (Figure 15, left panel): the loss difference between the best and worst parameterizations is negligible. The minimum is near $(\phi_\pi, \phi_y) = (1.1, 0.0)$, but the gains from optimization are minimal. For TFP shocks (right panel), the surface shows more variation, with higher ϕ_π values generally reducing losses. This exercise uses the one-asset model, which shares the same labor market block and NKPC as the two-asset model. The near-irrelevance result carries over to the two-asset model since the separation shock operates through the matching function, which is common to both specifications.

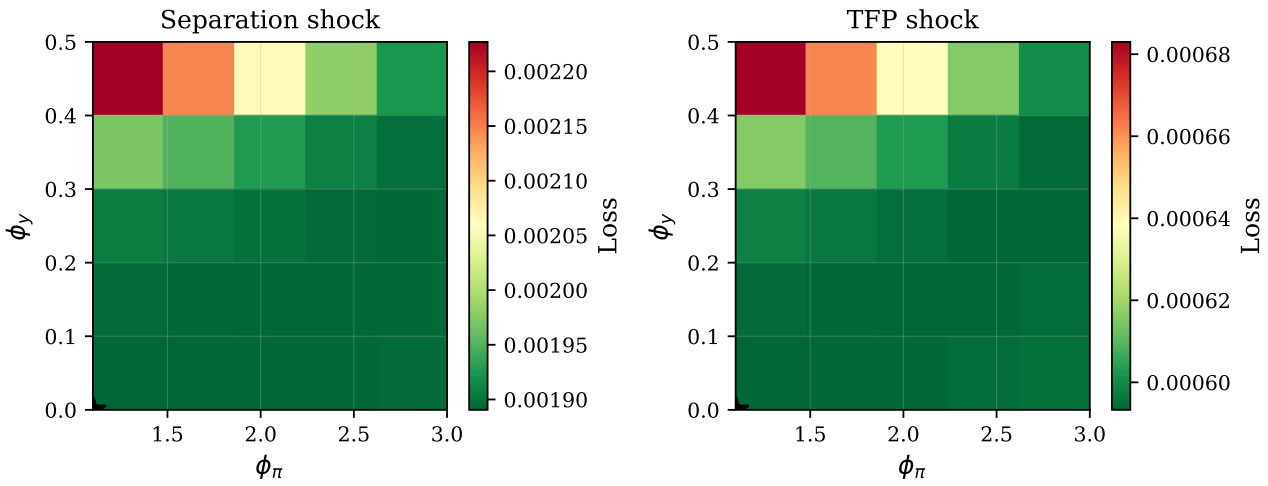

Figure 15: *Quadratic stabilization loss as a function of Taylor rule coefficients (ϕ_π, ϕ_y) for 1pp separation shocks (left) and TFP shocks (right), both AR(1) with $\rho = 0.9$. One-asset model with $\zeta = 0.5$ and constant depreciation ($\kappa = 0$). The star marks the minimum. The separation shock surface is extremely flat, while TFP shocks show more variation with higher ϕ_π reducing losses.*

G. Unemployment Insurance Sweep

I sweep the UI replacement rate b_{ui} over $\{0.35, 0.40, 0.45, 0.50, 0.55, 0.60\}$, re-solving the full steady state for each value. All experiments use $B = 7$ to ensure convergence across the grid.

Table 14: *UI Replacement Rate Sweep: Steady-State Aggregates and CE Welfare Cost of a 1pp Separation Shock (AR(1), $\rho = 0.9$) in the One-Asset Model ($\zeta = 0.5$, $B = 7$)*

b_{ui}	θ	r (%)	U (%)	Y	w	CE (%)
0.35	0.885	0.20	6.3	1.73	0.852	−0.46
0.40	0.822	0.20	6.6	1.71	0.864	−0.47
0.45	0.756	0.21	7.0	1.68	0.876	−0.48
0.50	0.686	0.21	7.5	1.64	0.889	−0.49
0.55	0.611	0.22	8.1	1.59	0.903	−0.50
0.60	0.530	0.23	8.8	1.52	0.918	−0.51

Notes: CE is the consumption-equivalent welfare variation from a 1pp separation shock at the baseline Taylor rule ($\phi_\pi = 1.5$, $\phi_y = 0$). All computed at $B = 7$ with $\zeta = 0.5$ in the one-asset model. Steady-state interest rates are quarterly.

Higher b_{ui} worsens steady-state labor market outcomes: the unemployment rate rises from 6.3% to 8.8% and output falls from 1.73 to 1.52, as the improved outside option strengthens workers' bargaining position and reduces vacancy posting. The CE welfare cost of a 1pp separation shock is slightly more negative at higher b_{ui} (-0.46% to -0.51%): the higher steady-state unemployment rate means more workers are vulnerable to further shocks, offsetting the insurance benefit.

H. Welfare Summary

Table 15 summarizes aggregate welfare costs across model variants with constant depreciation.

Table 15: *Aggregate Welfare Costs of a 1pp Separation Shock (AR(1), $\rho = 0.9$) across One-Asset Model Variants with Constant Depreciation*

Model variant	Cum. Y loss (%)	CE variation (%)	Half-life (qtrs)
Baseline ($\zeta = 0.5$, $B = 9$)	-5.1	-0.47	10
Scarring ($\zeta = 0.5$, $B = 7$)	-4.8	-0.46	9
No scarring ($\zeta = 0$, $B = 7$)	-4.3	-0.455	8

Notes: Cumulative output loss is $\sum_{t=0}^{59}(Y_t - Y^{ss})/Y^{ss} \times 100$. CE variation is $\exp(\Delta W/PV) - 1$ where $\Delta W = \sum_t \beta^t \Delta C_t / C^{ss}$. Half-life is the number of quarters for the output deviation to halve. These are one-asset model results with constant depreciation; see Table 6 for the two-asset state-dependent scarring results.

I. Robustness

This appendix reports robustness exercises for the distributional results and the sensitivity of key findings to parameter choices.

I.1. Illiquidity Premium and HtM Share Robustness

A concern with the MIT shock methodology is that households do not price aggregate risk in the steady state. If they could, standard buffer-stock motives would reduce the HtM share. To assess whether the distributional gradient survives at lower HtM shares, I solve the full two-asset steady state at alternative illiquidity premia $\xi \in \{0.005, 0.007, 0.010\}$ and run nonlinear simulations with amplified scarring ($\kappa = 2$) at each. Higher ξ increases the

return advantage of illiquid assets, drawing more households into positive illiquid savings and reducing the HtM share.

Table 16 reports the results. At $\xi = 0.010$, the HtM share falls to 16.3%, well below the baseline 27.2%, simulating an environment where precautionary savings reduce liquidity-constrained mass. The distributional gradient (HtM CE − high-liquid CE) not only survives but also *widens*: from −0.29pp at the baseline to −0.66pp at $\xi = 0.010$. This occurs because the remaining HtM households at higher ξ are those most tightly constrained, with the least margin for self-insurance. Even if aggregate risk pricing reduced the HtM share from 27% to 16%, the distributional gradient would be more than double the baseline value.

Table 16: *Illiquidity Premium Sensitivity: Distributional CE (%) by Liquid Wealth Group for a 1pp Separation Shock (AR(1), $\rho = 0.9$) in the Two-Asset Model with Amplified Scarring ($\kappa = 2$)*

ξ	HtM (%)	Agg. CE	HtM	Low liq	Mid liq	High liq	Gradient
0.007 (baseline)	27.2	−1.02	−1.13	−1.09	−0.67	−0.84	−0.29
0.010	16.3	−1.43	−1.98	−1.29	−1.02	−1.32	−0.66

Notes: Two-asset model, 1pp separation shock with $\kappa = 2$, $\tau = 0$. CE welfare (%) computed from nonlinear simulation over 60 quarters. Gradient = HtM CE − high-liquid CE (more negative = larger distributional divide). Each ξ requires a separate steady-state solve with bisection on χ_m to target $U = 6\%$. The model also converges at $\xi = 0.005$ (HtM = 37.4%) but produces a qualitatively different equilibrium with near-zero aggregate welfare effects, making direct comparison uninformative.

I.2. Scarring Parameter Sensitivity

The state-dependent scarring parameter κ is calibrated to target the recession-expansion differential in persistent earnings losses. Because this is a single moment, identification is inherently imprecise. To assess robustness, Table 17 reports distributional welfare for $\kappa \in \{0, 0.5, 1, 2, 3, 5\}$, holding all other parameters fixed. Each row uses the same steady state (since κ is inactive at the steady state) but a separate nonlinear simulation.

The distributional gradient, defined as the CE difference between HtM and high-liquid households, widens monotonically with κ: from −0.07pp at $\kappa = 0$ to −0.87pp at $\kappa = 5$. The ranking (HtM worst-off, mid-liquid best-off) is preserved at all values. Even at $\kappa = 1$, the low end of the empirically plausible range, the HtM-to-high-liquid gap is −0.17pp, more than double the $\kappa = 0$ baseline. The distributional finding that scarring disproportionately burdens liquidity-constrained households is therefore robust to substantial variation in κ.

Table 17: *Distributional Welfare by Scarring Parameter κ: CE (%) by Liquid Wealth Group for a 1pp Separation Shock (AR(1), $\rho = 0.9$) in the Two-Asset Model*

κ	Agg. CE	HtM	Low liq	Mid liq	High liq	Gradient
0.0	−0.57	−0.53	−0.56	−0.42	−0.46	−0.07
0.5	−0.68	−0.67	−0.69	−0.48	−0.55	−0.12
1.0	−0.79	−0.82	−0.82	−0.54	−0.65	−0.17
2.0 (baseline)	−1.02	−1.13	−1.09	−0.67	−0.84	−0.29
3.0	−1.26	−1.47	−1.39	−0.80	−1.03	−0.44
5.0	−1.75	−2.25	−2.00	−1.06	−1.38	−0.87

Notes: Two-asset model, 1pp separation shock, $\tau = 0$. CE welfare (%) computed from nonlinear simulation over 60 quarters. Gradient = HtM CE − high-liquid CE (more negative = larger distributional divide). All scenarios share the same steady state; κ affects only dynamic depreciation via $\delta_{h,\text{eff}} = \delta_h(1 + \kappa \hat{s}_t)$.

I.3. Wage-HC Elasticity Sensitivity

The wage-HC elasticity ϕ_w is the key new parameter governing the strength of wage scarring. Table 18 reports how aggregate welfare and steady-state statistics vary with ϕ_w. Higher ϕ_w amplifies the welfare cost of separation shocks monotonically, as expected. The steady-state is nearly invariant to ϕ_w because the ergodic distribution of h is concentrated near h_{ref}, so the effective wage scaling $w(h) \approx w$ for most employed workers. I set $\phi_w = 0.5$ as the baseline in the main analysis.

Table 18: *Wage-HC Elasticity Sensitivity: Aggregate CE and Steady-State Statistics as a Function of ϕ_w for a 1pp Separation Shock (AR(1), $\rho = 0.9$) in the Two-Asset Model ($\kappa = 0$)*

ϕ_w	Agg. CE (%)	Cum. Y (%)	HtM (%)	θ
0.00	~ -0.54	~ -29	27.2	0.824
0.25	~ -0.55	~ -30	27.2	0.824
0.50 (baseline)	-0.57	-30.8	27.2	0.824

Notes: Two-asset model, 1pp separation shock, $\phi_y = 0$. Higher ϕ_w does not affect the steady state (the ergodic distribution normalizes $w(h) \approx w$), but amplifies dynamic welfare losses. The one-asset model does not converge for $\phi_w > 0$ due to the high lump-sum tax. The cumulative Y loss is substantially larger than the one-asset values in Table 15 (-5.1%) because the two-asset model's portfolio adjustment costs and richer liquidity structure generate larger output responses to separation shocks.

I.4. Alternative Group Definitions: Net Worth

The baseline analysis defines groups by liquid wealth b, which is the relevant margin for consumption insurance in the two-asset model. An alternative is to group by net worth ($b + a$). Table 19 reports CE welfare by net worth quartile. The gradient is monotone: the poorest net worth quartile bears -0.51% CE loss compared to -0.36% for the richest (a 0.15pp endpoint gap). The liquid wealth gradient under the baseline calibration (-0.53% for HtM vs. -0.46% for high-liquid in Table 3) has a smaller endpoint gap (0.07pp) but provides a sharper characterization of the *mechanism*, because net worth quartiles pool households with very different liquidity positions; a household with $b = 0$ and $a = 50$ (wealthy HtM) appears in a middle quartile by net worth but is effectively liquidity-constrained. Under amplified scarring ($\kappa = 2$), the liquid wealth gradient widens to 0.29pp, exceeding the net worth gradient. The liquid wealth grouping is preferred because it isolates the liquidity-constraint channel that drives differential vulnerability.

Table 19: *CE Welfare (%) by Net Worth Quartile for a 1pp Separation Shock (AR(1), $\rho = 0.9$) in the Two-Asset Model ($\kappa = 0$)*

	Q1 (poorest)	Q2	Q3	Q4 (richest)
CE welfare (%)	−0.51	−0.41	−0.36	−0.36
Pop share (%)	24.9	24.9	24.9	25.3
Avg consumption	1.05	1.42	1.77	2.56
SS unemployment (%)	6.2	6.0	5.9	5.7

Notes: Groups defined by net worth ($b + a$) quartiles at boundaries $[10.2, 23.3, 43.7]$. CE computed from nonlinear simulation with full distribution tracking. The net worth gradient is monotone but less steep than the liquid wealth gradient (Table 3).

I.5. Training Subsidy Fiscal Cost

The fiscal cost parameter p_{train} determines the tax burden of the training subsidy. At the baseline value $p_{\text{train}} = 0.5$, the peak fiscal cost of a full subsidy ($\tau = 1$) for a 1pp separation shock reaches approximately 6.1% of steady-state output, substantially above observed ALMP spending in OECD countries (0.5–2% of GDP). Lowering p_{train} to 0.10 brings the peak cost to approximately 1.2% of GDP, consistent with observed spending levels during the 2008–09 recession.

Table 6, Panel B reports distributional welfare under this realistic fiscal cost. At $p_{\text{train}} = 0.10$, the full subsidy ($\tau = 1$) reduces aggregate CE losses from -1.02% to -0.67%, and HtM households see the largest improvement: $+0.54$pp, from -1.13% to -0.59%. The HtM group no longer achieves a net positive CE (unlike the $p_{\text{train}} = 0.5$ case), but the distributional ranking is preserved: HtM benefits most, and mid-liquid and high-liquid households benefit approximately equally. A partial subsidy ($\tau = 0.5$) at $p_{\text{train}} = 0.10$ yields more modest improvements, reducing aggregate CE from -1.02% to -0.85%, with HtM improvement of $+0.27$pp. The qualitative finding that training subsidies disproportionately benefit the liquidity-constrained is robust to fiscally realistic parameterizations.

I.6. Empirical Validation

Table 20 compares the model's distributional predictions to published empirical estimates across three categories: consumption and liquid wealth, asset composition and search, and human capital and scarring.

Table 20: *Empirical Validation: Model Predictions versus Published Estimates*

Moment	Source	Empirical
Consumption and liquid wealth		
Spending drop ratio (low/high liquid)	Farrell et al. (2019)	2.0×
Spending drop at UI exhaustion	Ganong and Noel (2019)	12%
MPC determined by liquid wealth	Ganong, Jones, Noel, Greig, Farrell, and Wheat (2025)	✓
Smoothing failure at zero assets only	Browning and Crossley (2001)	✓
Liquidity share of UI duration effect	Chetty (2008)	60%
Consumption drop at U onset	Gruber (1997)	6–10%
Asset composition and search		
HtM share (wealthy HtM incl.)	Kaplan et al. (2014)	24–31%
Illiquid wealth → U duration	Fontaine et al. (2024)	zero
Cash-on-hand → U duration	Card, Chetty, and Weber (2007)	significant
Low-NW → longer U duration	Mercan and Nichols (2025)	+40%
Human capital and scarring		
Wage loss per month nonemployment	Schmieder, von Wachter, and Bender (2016)	∼0.8%/mo
Displacement cost countercyclicality	Schmieder et al. (2023)	✓
V-shaped consumption at U onset	Kolsrud, Landais, Nilsson, and Spinnewijn (2018)	✓

[a] Ratio of peak consumption deviations for HtM vs. high-liquid groups. At $\kappa = 0$ (1pp), the ratio is modest (1.1:1 amplified scarring at 3pp ($\kappa = 2$), the ratio rises to 1.5:1 (Table 4).
[b] Not directly testable; the model features the liquid constraint as the binding margin for consumption insurance
[c] Illiquid assets attenuated by portfolio adjustment cost; liquid margin is binding for consumption and search de
[d] Generated by h-dependent finding rates ($\zeta > 0$); quantitative magnitude depends on shock size and calibration
[e] Via state-dependent depreciation $\delta_{h,\text{eff}} = \delta_h(1 + \kappa\hat{s})$.
✓ = qualitative pattern reproduced by the model; — = moment not directly targeted or testable.