Local Methods for Large Transfers*

Christian Bayer† Luigi-Maria Briglia‡

Ralph Luetticke§ Maximilian Weiss¶

Yannik Winkelmann‖

August 26, 2026

Abstract

Large fiscal transfers shift households between regions with high and low marginal propensities to consume (MPCs), resulting in nonlinear aggregate responses. We develop a Nonlinear DEGM Update (NDU) method that shifts the wealth distribution nonlinearly over a short time window while solving the aggregate economy using a fast, first-order state-space approximation. A 10 percent transfer of annual GDP increases output by 4.5 percent in the nonlinear and NDU solutions, but by 12.0 percent in the linear solution. The empirical liquid-wealth distribution around zero disciplines this nonlinearity and exhibits a strong asymmetry. A stochastic debt-entry cost closely reproduces this empirical pattern.

Keywords: heterogeneous agents, fiscal transfers, nonlinearities, state space, distribution dynamics.
JEL codes: C46, C63, D15, E21, E62.

*Christian Bayer gratefully acknowledges support by the DFG through CRC-TR224 EPoS, project C05 and in the framework of the German Excellence Strategy–EXC2126/1–390838866. Luigi-Maria Briglia, Ralph Luetticke, Maximilian Weiss, and Yannik Winkelmann gratefully acknowledge funding by the European Union (ERC, AIRMAC, 101114991). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

†Universität Bonn, CEPR, and IZA, *email:* christian.bayer@uni-bonn.de
‡Universität Tübingen, *email:* luigi.briglia@uni-tuebingen.de
§Universität Tübingen and CEPR, *email:* ralph.luetticke@uni-tuebingen.de
¶Universität Tübingen, *email:* maximilian.weiss@uni-tuebingen.de
‖Universität Tübingen, *email:* yannik.winkelmann@uni-tuebingen.de

1 Introduction

Large fiscal transfers create a general challenge for local solution methods in heterogeneous-agent models. In models with high and dispersed marginal propensities to consume (MPCs), a transfer moves households across regions of the asset distribution where MPCs differ sharply, making the aggregate response strongly concave in transfer size. Bianchi and Kaplan (2026) show that linear methods can therefore substantially overstate the efficacy of transfers. The source of this nonlinearity is not the curvature of individual policy functions by itself, but its interaction with the asset distribution. Aggregate consumption is an integral over that distribution, and a transfer moves the borrowing-constraint threshold through it. The relevant nonlinear term is the change in the fraction of high-MPC households. We make three contributions. First, we provide an efficient method for handling this nonlinearity in macroeconomic models. Second, we measure how quickly this fraction changes in the data. Third, we develop a model of household debt that closely matches this statistic.

For the method, we build on the distributional endogenous-grid method (DEGM) of Bayer et al. (2026). DEGM represents the distribution function in a way that captures how transfers shift it. It writes the Kolmogorov Forward Equation (KFE) as a distribution-function update at the cash-on-hand points that map into each grid point for future asset holdings. A transfer shifts cash-on-hand. Holding policies fixed, the DEGM update then gives the distribution of end-of-period assets for any transfer size, including changes in the mass of households that choose the borrowing constraint. More technically speaking, DEGM provides a fully nonlinear and differentiable distribution update based on the preimages of households' policy functions. Bayer et al. (2026) show that this approach is both computationally efficient and precise. It removes the linearity built into the widely used method of Young (2010), which represents continuous choices as lotteries over a discrete state space and misses nonlinear terms. This makes DEGM well suited to the concern raised by Bianchi and Kaplan (2026).

While Bayer et al. (2026) use DEGM to compute a third-order perturbation solution, we first derive how policy-boundary crossings from high- to low-MPC regions affect aggregate consumption and then take a different computational route. With DEGM, the initial distribution after a transfer can be updated without approximation before solving for equilibrium prices. Feeding this new distribution into the linearized state-space representation already gives a response that is nonlinear in shock size. Accuracy improves further when we also account for second-round income effects in the first few periods. Inside each updated period, aggregate output, and thus income, is calculated from a scalar fixed point. The rest of the computation remains the linear state-space solution. We call this method the Nonlinear DEGM Update (NDU). It matches the fully nonlinear solution and is almost as fast as the linear solution.

Keeping the distribution as an explicit state is key. In the sequence-space implementation considered here, a one-period savings correction fixes the impact response, but it does not track the later dynamics of the corrected distribution; Appendix E reports this version. In state space, the distribution remains an integral part of the description of the economy. The linear distribution update can therefore be replaced by the exact DEGM update for one or a few periods, and the correction affects the entire future path. Extending the exact-update window improves the non-impact part of the response.

Quantitatively, NDU captures large-transfer nonlinearities at a cost close to that of a linear solution. With a transfer of 10 percent of annual GDP, the first-quarter general-equilibrium output response is 4.48 percent under the fully nonlinear solution, 11.96 percent under the linearized solution, and 4.52 percent under NDU. Updating only the first-quarter distribution matches the impact response, though small deviations remain later. Extending the window by a few periods largely eliminates them.

Our second contribution is empirical: we introduce the zero-crossing curve, which combines observed liquid wealth with model-implied consumption policies. For each transfer size, it measures the fraction of households that begins with weakly negative liquid wealth and ends with strictly positive liquid wealth after its consumption response. The curve directly disciplines the distributional margin behind the aggregate nonlinearity. In the SCF, 17.6 percent of households have negative liquid wealth, another 6.7 percent have zero liquid wealth, and the distribution is asymmetric around zero: shallow debt is sparse, while small positive wealth positions are dense.

Two workhorse specifications do not reproduce this pattern: a model with a zero borrowing limit (for example, Auclert et al., 2025; Bianchi and Kaplan, 2026) and a model with only a borrowing wedge (for example, Kaplan and Violante, 2022; Bayer et al., 2019). Neither produces the long, flat debt distribution observed in the data. As a result, both generate a flow across zero that is too large and too abrupt. We show that a parsimonious debt-access friction delivers the empirically observed asymmetry around zero liquid wealth. Incumbent borrowers retain access, while non-borrowers either receive a waiver or pay a fixed cost to enter debt.[1]

Matching the crossing curve matters for the aggregate nonlinearity of transfers. Because a large and dispersed share of households has negative or zero liquid wealth in the data, the crossing of households depends on the size of the transfer. This produces a gap between nonlinear and linearized solutions. The gap, however, is smaller than in the zero-debt setup of Bianchi and Kaplan (2026). A transfer equal to 5 percent of annual GDP moves about 18 percent of households from weakly negative to strictly positive liquid wealth. Doubling the transfer to 10 percent raises the crossing share to only about 21 percent, so the curve becomes flatter as transfers grow. NDU uses the exact CDF update to capture this distributional nonlinearity.

The transfer-size effect has direct precedents in household evidence and quantitative consumption models. Kaplan and Violante (2014b) find that the consumption response to the 2008 stimulus payments was about one-third smaller than in 2001, mainly because the 2008 payments were larger. Using Norwegian lottery prizes, Fagereng et al. (2021) show that MPCs decline with prize size and liquid wealth. Survey evidence likewise shows that larger windfalls have smaller impact MPCs but more persistent spending effects (Jappelli et al., 2026). Direct evidence also shows that households use transfers to deleverage. Agarwal et al. (2007) find that recipients initially used the 2001 tax rebates to reduce credit-card debt before increasing spending, while Koşar et al. (2025) document a large marginal propensity to repay debt among households with low net wealth relative to income and rationalize it with borrowing rates that increase with debt. This literature establishes that both shock size and the timing of household spending

[1]Related models use stochastic credit limits (Fulford, 2015) or explicit applications and persistent lending relationships (Herkenhoff, 2019). The fixed cost also provides a possible rationalization of coholding: with separate gross positions, households may keep borrowing access while holding liquid assets (Boutros and Mijakovic, 2024).

matter.[2] Bianchi and Kaplan (2026) show why this fact creates a problem for local solution methods in HANK models. Our contribution is to identify the distributional object behind this problem—the mass of households that crosses the high-MPC boundary—and discipline it with the liquid-wealth distribution.

Two contemporaneous papers clarify the scope of this problem. Auclert et al. (2026) show that third-order perturbation is accurate in a smooth HANK model for many shocks. They also find that its largest errors arise for large, one-time lump-sum transfers and their impact effects, which is precisely the experiment studied here. Guerreiro et al. (2026) provide suggestive aggregate evidence that larger temporary transfers have smaller effects per dollar. Our paper is complementary to both. We show how to compute this difficult transfer experiment at near-linear cost and how its size dependence changes when the liquid-wealth distribution near zero is matched to the data.

Our new method, NDU, is related to methods that improve local solutions away from the steady state. Evans and Phillips (2017) linearize around the current state, and Mennuni et al. (2025) compute time-varying first-order approximations along the equilibrium path. Andreasen and Kronborg (2022) instead combine a nonlinear extended-path calculation with perturbation corrections.[3] NDU takes a more directed approach. It does not recompute policy derivatives or solve the full nonlinear path. It keeps the steady-state first-order decision rule fixed, replaces only the distribution transition and the associated aggregate-savings equation by their exact nonlinear counterparts, and passes the corrected distribution to the original linear continuation.

The remainder of the paper proceeds as follows. Sections 2 and 3 develop the distributional argument: modeling the distribution as a CDF allows us to write the Kolmogorov Forward Equation update as a closed-form function of a transfer that moves households from high- to low-MPC policy branches. Section 4 gives the state-space representation, and Section 5 develops NDU. Section 6 presents the quantitative evidence in three steps. It first measures the zero-crossing flow, then studies the partial-equilibrium consumption response, and finally compares NDU with the linear and fully nonlinear general-equilibrium solutions. Appendix A states the full economic model. Appendix B documents the calibration and empirical mapping. Appendix C gives the complete state-space implementation, Appendix D extends the policy-boundary analysis to debt-entry costs, and Appendix E develops a sequence-space analogue to NDU but also shows that this fails to achieve the precision gain of the state-space version because it does not iterate the distributional adjustments forward.

2 The Distribution-Function Update

Consider an economy with a continuum of households that face uninsurable income risk and accumulate an asset a to self-insure. Households can also differ persistently in preferences or productivity. The rest of the economy is a standard New Keynesian environment: firms hire workers and produce, the government issues debt and levies taxes, and the central bank controls

[2] Additional survey evidence documents extensive-margin adjustment, size dependence, and strong heterogeneity in spending responses (Fuster et al., 2021), as well as sign and size asymmetries in responses to transitory income shocks (Christelis et al., 2019).

[3] Similarly, den Haan et al. (2016) solve the current equations nonlinearly while using a Taylor approximation for future behavior.

the nominal return on the asset. Appendix B gives the calibration, and Appendix C gives the complete implementation.

We focus first on the household side and consider a transfer from the government. Time is discrete, and the distribution carried into a period is the beginning-of-period asset CDF. Let $F^a_\theta(a) = \Pr(A \le a \mid \theta)$ denote the share of households in persistent state $\theta \in \Theta$ whose assets are weakly below a[4], and let $\mu(d\theta)$ be the population measure over these states. Cash-on-hand is obtained from assets through the monotone map

$$x = \phi_{\theta k}(a) \equiv R(a)a + \iota_{\theta k}, \qquad R(a) := \begin{cases} R, & a \ge 0, \\ R + \rho_0, & a < 0, \end{cases} \tag{1}$$

where $\iota_{\theta k}$ is current income conditional on transitory-income draw k, which occurs with probability ω_k, R is the gross asset return, and ρ_0 is the interest-rate wedge for negative asset holdings. Thus the beginning-of-period cash-on-hand CDF can be derived from the asset CDF as

$$F^x_\theta(q) = \sum_k \omega_k F^a_\theta\left(\phi^{-1}_{\theta k}(q)\right), \qquad \phi^{-1}_{\theta k}(q) = \begin{cases} \dfrac{q - \iota_{\theta k}}{R}, & q - \iota_{\theta k} \ge 0, \\ \dfrac{q - \iota_{\theta k}}{R + \rho_0}, & q - \iota_{\theta k} < 0. \end{cases} \tag{2}$$

For a common gross return R and abstracting from transitory-income draws, Equation (2) reduces to $F^x_\theta(q) = F^a_\theta((q - \iota_\theta)/R)$. For readability, we write $F_\theta := F^x_\theta$ and use f_θ for its density throughout this section. The state-space representation later returns to the asset CDF F^a_θ, which is the distributional state of the model.[5]

Let $g_\theta(x)$ denote the end-of-period asset policy at given prices. Conditional on persistent type θ[6], it maps current cash-on-hand into next period's assets. To describe a policy boundary generally, suppose

$$g_\theta(x) = \begin{cases} g^-_\theta(x), & x \le \bar{x}_\theta, \\ g^+_\theta(x), & x > \bar{x}_\theta, \end{cases} \tag{3}$$

where the two branches meet continuously at the type-specific cash-on-hand threshold $\bar{x}_\theta$, with

$$g^-_\theta(\bar{x}_\theta) = g^+_\theta(\bar{x}_\theta) = \bar{a}_\theta.$$

Away from flat segments, each branch is strictly increasing. Consumption is $c^j_\theta(x) = x - g^j_\theta(x)$, and the MPC on branch $j \in \{-, +\}$ is

$$m^j_\theta(x) = 1 - g^j_{\theta,x}(x). \tag{4}$$

[4]We use the left-limit notation $F^a_\theta(a^-) = \Pr(A < a \mid \theta)$ whenever the underlying choice inequality is strict.

[5]The conditioning on θ matters because a conditional distribution can have mass points, for example at a borrowing limit. The aggregate CDF integrates these conditional distributions over the type-income measure μ. As in Bayer et al. (2026), a continuous income component mixes them, so a cutoff that coincides with a mass point for one type is a measure-zero event in the aggregate integral.

[6]The notation in this and the next section treats θ as the index of the current policy that a household faces. In the full model with debt-entry cost κ, this index is augmented by beginning-of-period borrower status and an iid waiver draw. The measure μ then includes their probabilities. Furthermore, while the continuous-policy argument below uses $c = x - g$, for a non-waived entrant the corresponding identity is $c = x - \kappa - g$. See Appendix A.3.

The economically relevant case has $m_\theta^-(\bar{x}_\theta^-) > m_\theta^+(\bar{x}_\theta^+)$, where the superscripts $-$ and $+$ denote one-sided limits: the threshold separates a high-MPC region from a low-MPC region. A hard borrowing constraint is the special case $g_\theta^-(x) = \underline{a}$, so $\bar{a}_\theta = \underline{a}$ and $m_\theta^- = 1$.

Given this policy, a transfer of size Δ shifts cash-on-hand to $x + \Delta$. The distribution after the policy choice is the distribution of end-of-period assets,

$$\widetilde{F}^a_{\theta,\Delta}(a') = \Pr\left(g_\theta(x+\Delta) \le a'\right). \tag{5}$$

On either strictly increasing branch, the post-policy CDF is obtained by going back to the cash-on-hand level that chooses a':

$$\widetilde{F}^a_{\theta,\Delta}(a') = F_\theta\left((g_\theta^j)^{-1}(a') - \Delta\right), \qquad j = \begin{cases} - & a' \le \bar{a}_\theta, \\ + & a' > \bar{a}_\theta. \end{cases} \tag{6}$$

At the policy boundary itself,[7]

$$\widetilde{F}^a_{\theta,\Delta}(\bar{a}_\theta) = F_\theta\left(\bar{x}_\theta - \Delta\right). \tag{7}$$

Equations (6)–(7) are the DEGM forward step of Bayer et al. (2026): they produce the end-of-period asset CDF $\widetilde{F}^a_{\theta,\Delta}$ by evaluating the derived cash-on-hand CDF at the points that map into the asset grid. At a hard borrowing limit, this graph contains the point

$$(\underline{a},\ F_\theta(\bar{x}_\theta - \Delta))$$

and the slack-branch points

$$\left(g_\theta^+(X_i + \Delta),\ F_\theta(X_i)\right), \qquad X_i > \bar{x}_\theta - \Delta.$$

where X_i denotes a cash-on-hand grid node. A smooth monotone interpolant of the CDF over these points approximates the continuous KFE. Two features of this representation carry the rest of the paper. First, the update is *differentiable* in aggregates such as Δ, so Taylor expansions apply to the transition. Second, the update is *closed-form* in Δ: given the beginning-of-period CDF F_θ, the exact nonlinear transition requires only one set of CDF evaluations at the corresponding shifted cash-on-hand points.

3 Policy-Boundary Crossing Terms

The preceding section shows how a transfer moves the asset distribution. We now show how that movement affects aggregate consumption when households cross a boundary between policy

[7]The following formula also covers a hard borrowing limit with a slight adjustment. When $g_\theta^-(x) = \underline{a}$, the lower branch is flat rather than invertible, and (7) becomes the mass at the constraint:

$$\widetilde{F}^a_{\theta,\Delta}(\underline{a}) = F_\theta\left(\bar{x}_\theta - \Delta\right).$$

branches. Aggregate first-period consumption is[8]

$$C(\Delta) = \int_\Theta \int c_\theta(x+\Delta)\, dF_\theta(x)\, \mu(d\theta). \tag{9}$$

Differentiating this expression produces two kinds of terms: changes within each policy branch and changes in the mass of households on each branch. We refer to the latter as *boundary-flow terms*.[9]

Proposition 1 (Policy-boundary crossing terms)**.** *Let $C_\theta(\Delta)$ denote the conditional type-θ contribution to* (9)*. For μ-almost every θ, suppose the continuous part of F_θ has a differentiable density at the moved cutoff $q_\theta(\Delta) = \bar{x}_\theta - \Delta$, and any mass points are isolated. Suppose the set of θ for which $q_\theta(\Delta)$ coincides with such a mass point has μ-measure zero. Suppose the two consumption branches are continuous at $\bar{x}_\theta$ and sufficiently smooth away from it. Then, for regular θ where $f_\theta(\bar{x}_\theta - \Delta)$ exists,*

$$C'_\theta(\Delta) = \left[\int_{-\infty}^{\bar{x}_\theta-\Delta} m_\theta^-(x+\Delta) dF_\theta(x) + \int_{\bar{x}_\theta-\Delta}^{\infty} m_\theta^+(x+\Delta) dF_\theta(x)\right], \tag{10}$$

and

$$\begin{aligned} C''_\theta(\Delta) = &\left[m_\theta^+(\bar{x}_\theta^+) - m_\theta^-(\bar{x}_\theta^-)\right] f_\theta(\bar{x}_\theta - \Delta) \\ &+ \int_{-\infty}^{\bar{x}_\theta-\Delta} m_{\theta,x}^-(x+\Delta) dF_\theta(x) \\ &+ \int_{\bar{x}_\theta-\Delta}^{\infty} m_{\theta,x}^+(x+\Delta) dF_\theta(x). \end{aligned} \tag{11}$$

Under the corresponding integrability conditions, aggregate derivatives are obtained by integrating these expressions over $\mu(d\theta)$.

Proof. Let $q_\theta(\Delta) = \bar{x}_\theta - \Delta$. Split the type-$\theta$ contribution at the moving cutoff:

$$C_\theta(\Delta) = \int_{-\infty}^{q_\theta(\Delta)} c_\theta^-(x+\Delta)\, dF_\theta(x) + \int_{q_\theta(\Delta)}^{\infty} c_\theta^+(x+\Delta)\, dF_\theta(x). \tag{12}$$

When we differentiate, the moving limits produce boundary terms. They cancel in the first derivative because consumption is continuous at the threshold. What remains is the average MPC on each side, which gives (10). Differentiating once more with Leibniz' rule gives (11): the new boundary term is the change in the MPC across the threshold, multiplied by the density at the moved cutoff. Integrating over θ gives the aggregate derivative. The mass-point exceptions do not contribute because the set of affected types has μ-measure zero. □

[8] Equivalently, consumption equals aggregate resources minus aggregate end-of-period saving,

$$C(\Delta) = \mathcal{X}(\Delta) - \int_\Theta \int a'\, d\widetilde{F}^a_{\theta,\Delta}(a')\, \mu(d\theta) - EC(\Delta), \tag{8}$$

where $\mathcal{X}(\Delta) = \int_\Theta \int (x+\Delta)\, dF_\theta(x)\, \mu(d\theta)$ is aggregate cash-on-hand including the transfer and $EC(\Delta)$ denotes resources used for debt entry. See Section 6.2.

[9] Proposition 1 covers the model without a debt entry cost. A debt entry cost introduces jumps in consumption at the entry thresholds x^*_θ. The boundary terms then no longer cancel in the first derivative. Appendix D.1 states and proves the version for discontinuous policies in our stochastic entry-cost model.

For a regular type, the leading nonlinearity from movement across the policy boundary is therefore

$$\Big(m_\theta^+(\bar{x}_\theta^+) - m_\theta^-(\bar{x}_\theta^-)\Big)\, f_\theta(\bar{x}_\theta - \Delta). \tag{13}$$

This is the *boundary-flow term.* Since the MPC falls across the threshold, it is negative. It is generally nonzero even though the set exactly at the kink has measure zero, because the cutoff moves through a positive density of households. The hard borrowing-constraint result is nested as $m_\theta^- = 1$, in which case (13) reduces to $-[1 - m_\theta^+(\bar{x}_\theta^+)]f_\theta(\bar{x}_\theta - \Delta)$.

Remark 1. *Even if both individual consumption branches are linear, aggregate higher-order derivatives need not be zero. A transfer changes which households lie on each branch, and the CDF update evaluates the mass swept through the moved cutoff.*

3.1 A Minimal Example of Distributional Nonlinearity

The distributional origin of the higher-order terms is transparent when both individual policy branches are linear. Suppose

$$g_\theta(x) = \begin{cases} \bar{a}_\theta + \varsigma_\theta^-(x - \bar{x}_\theta), & x \le \bar{x}_\theta, \\ \bar{a}_\theta + \varsigma_\theta^+(x - \bar{x}_\theta), & x > \bar{x}_\theta, \end{cases} \qquad 0 \le \varsigma_\theta^- < \varsigma_\theta^+ \le 1. \tag{14}$$

The two branches meet continuously, and their constant MPCs satisfy $m_\theta^- = 1 - \varsigma_\theta^- > m_\theta^+ = 1 - \varsigma_\theta^+$. Nevertheless the aggregate response is

$$C_\theta(\Delta) - C_\theta(0) = m_\theta^+\Delta + (m_\theta^- - m_\theta^+)\int_0^\Delta F_\theta(\bar{x}_\theta - u)\, du. \tag{15}$$

The nonlinearity is entirely distributional. Expanding the CDF around $\bar{x}_\theta$ gives

$$\begin{aligned} C_\theta(\Delta) - C_\theta(0) = {}& \Big[m_\theta^+ + (m_\theta^- - m_\theta^+)F_\theta(\bar{x}_\theta)\Big]\,\Delta \\ & - \frac{1}{2}(m_\theta^- - m_\theta^+)f_\theta(\bar{x}_\theta)\Delta^2 + \frac{1}{6}(m_\theta^- - m_\theta^+)f_\theta'(\bar{x}_\theta)\Delta^3 + O(\Delta^4). \end{aligned} \tag{16}$$

Thus aggregate curvature requires neither individual policy curvature nor an actual borrowing constraint. It arises from the mass moved across a policy boundary and the MPC difference between the two branches. The borrowing-constraint example corresponds to $\varsigma_\theta^- = 0$, or $m_\theta^- = 1$. With heterogeneous types, the aggregate expansion integrates (16) over $\mu(d\theta)$.

The exact boundary $\bar{x}_\theta$ is nevertheless type-specific and model-dependent. It varies with preferences, income risk, and the household policy and is not directly observed in survey data. Taking the mechanism to the data therefore requires a common balance-sheet threshold that is measured consistently and captures, on average, movement from a high-MPC region to a lower-MPC region. Section 6.1 uses zero liquid wealth for this purpose and measures how quickly households cross it as transfers increase.

4 A State-Space Solution

We solve the sticky-wage HANK economy of Bianchi and Kaplan (2026) with the cross-sectional distribution as an explicit state (Reiter, 2009). Our baseline debt model permits negative assets and charges some non-borrowers a fixed cost when they enter debt. Modeling debt with a wedge creates a discrete mass at zero and requires us to track the mass strictly below zero separately. The fixed debt-entry cost additionally requires the tracking of control variables that govern the debt-entry decision for every type. We keep the description of the state space short here; Appendices A and C state the household problem, policy branches, distribution transition, resource costs, and equilibrium conditions needed to reconstruct the implementation in detail. Appendix B gives the calibration.

States and controls. Let $h =$ (discount-factor type, z) index the combined permanent-type and persistent-income state, with transition matrix $\Pi_{hh'}$, stationary population mass p_h, and N_H values. Unlike the conditional CDFs used in Sections 2–3, each numerical CDF column is joint: $F_{t,h}(a) = \Pr(a_t \leq a, h_t = h)$. It therefore ends at p_h, rather than at one. The state vector s_t stacks

- the distribution coordinates S_t: deviations of the beginning-of-period joint CDF columns $F_{t,h}(a)$ from their steady-state values, at every interior asset grid node—the distribution at *full* grid resolution—and, in debt models, the joint borrower masses $m^B_{t,h} = F_{t,h}(0^-)$.[10]
- the lagged real rate $r^{lag}_t = r_{t-1} - \bar{r}$, which converts beginning-of-period assets into cash-on-hand;
- a one-time MIT shock state ε_t that carries the transfer innovation into the fiscal rule.

The controls y_t stack a compressed representation dc_t of the consumption-policy deviations and the aggregates $(B_t, A_t, A^-_t, Y_t, C_t, \pi_t, i_t, r_t, T_t)$. Here B_t is government debt, A_t is aggregate beginning-of-period household assets, $A^-_t \leq 0$ is their negative part, Y_t is output, C_t is consumption, π_t is inflation, i_t and r_t are the net nominal and real rates, T_t is the transfer, and the constant G is government spending. Bars denote steady-state values, and the gross saving return satisfies $\bar{R} = 1 + \bar{r}$ and $R_t = \bar{R} + (r_{t-1} - \bar{r})$. With a fixed entry cost, the controls also include the boundary premium J_t and the left-limit marginal value $M^{0^-}_t$ for every state h. These boundary controls close the backward household problem; the endogenous entry threshold itself is recovered by value matching and substituted into the equilibrium system.

The DEGM transition as the distributional law of motion. The distribution block's law of motion is the discretized KFE in the CDF representation of Section 2. To show its basic logic, first consider a policy branch without an entry cost. The timing is as follows. $F_{t,h}$ is the beginning-of-period joint CDF column for assets carried into period t. A household in state h, with transitory draw k and beginning assets a, has cash-on-hand $x = R(a)a + \iota_{h,k}(Y_t, T_t)$ and saves $a' = x - c(x)$.

[10] Since we solve the model first-order, we do not have to use any dimensionality reduction for the distribution. Bayer and Luetticke (2020) propose reductions that make the state-space method as fast as sequence-space solution methods.

Current income is $\iota_{h,k}(Y,T) = (1-\tau)w_{h,k}Y + T$, where τ is the labor-income tax, $w_{h,k}$ is labor efficiency, and ω_k is the probability of draw k. Aggregate labor income is $(1-\tau)\omega Y$, with $\omega = \sum_h p_h \sum_k \omega_k w_{h,k} = 1$ under the income-process normalization. The gross return depends on the sign of beginning assets: $R(a) = R_t + \rho_0 \mathbf{1}\{a < 0\}$.

By the endogenous-grid construction (Carroll, 2006), the cash-on-hand threshold $x_{j,h}$ at which the household chooses exactly $a' = a_j$ is known. Because saving is monotone in cash-on-hand,

$$a' \leq a_j \iff a \leq a^{\text{pre}}_{j,h,k}, \qquad a^{\text{pre}}_{j,h,k} = \begin{cases} \dfrac{x_{j,h} - \iota_{h,k}}{R_t}, & x_{j,h} - \iota_{h,k} \geq 0, \\ \dfrac{x_{j,h} - \iota_{h,k}}{R_t + \rho_0}, & x_{j,h} - \iota_{h,k} < 0. \end{cases} \tag{17}$$

The sign of $x_{j,h} - \iota_{h,k}$ determines the return side because it is also the sign of beginning assets. The end-of-period savings CDF column for current state h therefore has a *closed-form, piecewise formula*: evaluate the interpolated beginning-of-period CDF column at $a^{\text{pre}}_{j,h,k}$,

$$\tilde{F}_{j,h} = \sum_k \omega_k \hat{F}_{t,h}\Big(a^{\text{pre}}_{j,h,k}\Big), \qquad F_{t+1,j,h'} = \sum_h \Pi_{hh'} \tilde{F}_{j,h} \equiv \mathcal{T}(F_t;\, x,\, \iota,\, R_t)_{j,h'}, \tag{18}$$

where $\hat{F}_{t,h}$ is a monotone-cubic interpolant of column h and ω_k are the transitory-draw weights. In the debt-entry model, this expression is evaluated separately on the negative and nonnegative sides. Incumbent borrowers, waived non-borrowers, and non-waived non-borrowers have different beginning-asset thresholds; the last group chooses among entering debt, remaining at zero, and saving. The transition therefore advances both the full CDF and m^B. Appendix C gives the complete mapping that treats the two sides separately. Aggregate total and negative assets, $A[F]$ and $A^-[F]$, follow from the same CDF interpolants by integration by parts.

The stacked equilibrium system. We next combine the distribution transition with household optimality and the aggregate equilibrium conditions. An equilibrium path $\{s_t, y_t\}$ satisfies

$$\mathbb{E}_t\, \mathcal{F}(s_t,\, y_t,\, s_{t+1},\, y_{t+1}) = 0, \tag{19}$$

with one row per condition. Grouping the conditions by economic block,

$$\mathcal{F} = \begin{pmatrix} (F_{t+1}, m^B_{t+1}) - \mathcal{T}\Big(F_t, m^B_t;\, \bar{x} + dc_t, J_t, M^{0^-}_t, Y_t, T_t, R_t\Big) & \text{(distribution: KFE rows)} \\ A_t - A[F_t, m^B_t], \qquad A^-_t - A^-[F_t, m^B_t], \qquad EC_t - \kappa e(F_t, m^B_t; \cdots) & \text{(distribution aggregates)} \\ dc_t - \mathcal{E}(dc_{t+1},\, Y_{t+1},\, T_{t+1},\, r_t), \qquad M^{0^-}_t - \mathcal{M}(M^{0^-}_{t+1}, \cdots) & \text{(household policy: Euler rows)} \\ J_t - \mathcal{J}\Big(dc_{t+1}, J_{t+1}, M^{0^-}_{t+1}, Y_{t+1}, T_{t+1}, r_t\Big) & \text{(boundary-value rows)} \\ C_t - \Big[R_t A_t + \rho_0 A^-_t + (1-\tau)\omega Y_t + T_t - EC_t - A_{t+1}\Big] & \text{(household budget)} \\ Y_t - C_t - G + \rho_0 A^-_t - EC_t, \qquad A_{t+1} - B_t & \text{(goods and bond clearing)} \\ \text{Phillips, Taylor, Fisher, fiscal rule} & \text{(macro rows)} \end{pmatrix}, \tag{20}$$

where $\mathcal{E}$ is the household block's backward Euler operator, while $\mathcal{J}$ and $\mathcal{M}$ are the boundary recursions for the debt entry-cost decision, which we define in Appendix C. Conditional on $(dc_t, J_t, M^{0^-}_t)$, value matching recovers the entry threshold x^*_t; the function $e(\cdot)$ is the resulting mass of non-waived entrants and $EC_t = \kappa e_t$. Since $A_t^- \leq 0$, intermediation costs are $-\rho_0 A_t^- \geq 0$; both intermediation and entry costs are wasted, i.e. they absorb goods. Combining the household budget with goods and bond clearing gives the aggregate-savings row on which NDU operates. The nonlinear operator $\mathcal{T}$ determines A_{t+1}, A_{t+1}^-, and the entry flow, so it affects both the distribution law and aggregate resource clearing.

Linearization. The linear solution replaces every row of (20) by its first-order expansion around the steady state and solves the resulting linear system, giving the decision rule and law of motion

$$y_t = g_s s_t, \qquad s_{t+1} = h_s s_t. \tag{21}$$

We write dZ_t for the first-order deviation of any variable Z_t. In particular, the linearized aggregate-savings row reads

$$dY_t \;=\; \bar{R}\, A'[F_{ss}]\, dS_t + \bar{A}\, dr_t^{lag} + (1-\tau)\omega\, dY_t + dT_t \;-\; d\, A[\mathcal{T}], \tag{22}$$

Here $A'[F_{ss}]dS_t$ is the derivative of the current-asset functional at the steady-state distribution. The term $d\,A[\mathcal{T}]$ is the first-order change in end-of-period saving. More precisely, it is the derivative of $A[\mathcal{T}(F;\, \bar{x} + dc,\, \iota,\, R)]$ with respect to each argument, evaluated at the steady state and applied to $(dS_t,\, dc_t,\, dY_t,\, dr_t^{lag},\, dT_t)$. The vector $\bar{x}$ collects the steady-state endogenous-grid cash-on-hand nodes.

This saving term is where the linear solution replaces the moved cutoff $F(\bar{x} - \Delta)$ by its tangent. The purely linear solution applies (21) from $t = 1$ with $s_1 = \Delta\, e_\varepsilon$, where e_ε is the unit vector selecting the one-time shock state. Section 6 shows how this approximation affects both the impact response and the subsequent transition path.

5 The Nonlinear DEGM Update

The nonlinear DEGM update (NDU) keeps the linear solution for future movements and for the household policy, but uses the exact equations where the shock causes a large shift in the distribution. For period 1—the impact period of the unanticipated transfer—it takes the system (20) and replaces two parts by their exact nonlinear counterparts:

- the *distribution rows*: instead of the linearized law of motion (the first block of (20) expanded to first order), the period-2 distribution is the exact image $(F_2, m_2^B) = \mathcal{T}(F_{ss}, m_{ss}^B;\, \bar{x} + dc_1, J_1, M_1^{0^-}, \iota(Y_1, T_1), \bar{R})$;

- the *aggregate-savings row*: instead of the linearized goods-market condition (22), in which the savings aggregate enters through its differential $d\,A[\mathcal{T}]$, period-1 output clears the market against the exact savings aggregate $A[\mathcal{T}(\cdot)]$ itself.

Every other row stays first-order. NDU therefore does not recompute the perturbation around the updated distribution: g_s and h_s remain the matrices obtained at the steady state. In particular, the Euler rows are *not* replaced: the period-1 household policy still comes from the linear decision rule g_s. Because the distribution is a state, the exact F_2 produced by the replaced distribution row can be passed directly to the linear continuation. This is how NDU propagates: the corrected distribution is carried forward by h_s and shapes every subsequent period. Operationally, the period-1 solve proceeds in three steps.

5.1 Impact NDU step by step

Step 1: First-order policy. The period-1 state is $s_1 = \Delta\, e_\varepsilon$: households start at the steady-state distribution, the predetermined return is fixed at $R_1 = \bar{R}$, and only the shock state changes. The period-1 consumption-policy and boundary-control deviations follow from the corresponding rows of the linear decision rule,

$$(dc_1, dJ_1, dM_1^{0^-}) = \Delta \cdot (g_s^{(c)}, g_s^{(J)}, g_s^{(M)}) e_\varepsilon. \tag{23}$$

Because g_s solves the full forward-looking system, dc_1 contains the linear path's response to future prices (i.e., we linearize substitution effects and expectations). However, we let today's income enter separately, through $\iota(Y_1, \bar{T} + \Delta)$, inside the distribution update, i.e., we keep contemporaneous income effects potentially nonlinear.

Step 2: Exact one-dimensional goods-market clearing. Given dc_1, period-1 end-of-period saving is the asset aggregate of the exact DEGM transition at the shifted thresholds $\bar{x} + dc_1$,

$$A_2(Y_1) = A\Big[\mathcal{T}\Big(F_{ss}, m_{ss}^B;\, \bar{x} + dc_1, J_1, M_1^{0^-}, \iota(Y_1, \bar{T} + \Delta), \bar{R}\Big)\Big], \tag{24}$$

consumption follows from the aggregate budget identity,

$$C_1(Y_1) = \bar{R}\bar{A} + \rho_0 \bar{A}^- + (1-\tau)\omega Y_1 + \bar{T} + \Delta - EC_1(Y_1) - A_2(Y_1),$$

and impact output solves the scalar goods-market condition

$$\boxed{Y_1 = C_1(Y_1) + G - \rho_0 \bar{A}^- + EC_1(Y_1).} \tag{25}$$

Equation (25) is the aggregate-savings row of (20), restored to its exact form: comparing with the linearized row (22), the first-order saving term $d\,A[\mathcal{T}]$ has been replaced by the full nonlinear savings aggregate $A_2(Y_1)$. The row's remaining terms—returns on beginning-of-period assets, labor income, and the transfer—are linear and therefore unchanged. Since prices and the policy are fixed inside this solve, the replacement turns the row into a one-dimensional equation in Y_1. The transition $\mathcal{T}$ then captures the full nonlinearity caused by households moving across the policy regions summarized by the zero-crossing curve in Panel (a) of Figure 1.

Proposition 2 (Monotonicity and uniqueness of scalar clearing)**.** *Let $H(Y_1) = Y_1 - \Phi(Y_1)$, where $\Phi(Y_1) = G - \rho_0 \bar{A}^- + EC_1(Y_1) + C_1(Y_1)$, and hold the policy thresholds and return fixed inside the scalar solve. Suppose every fixed saving policy is nondecreasing in cash-on-hand and $(1-\tau)\omega < 1$. Then H is strictly increasing. Hence* (25) *has at most one solution; if H is continuous on a bracket whose endpoints have opposite signs, it has exactly one solution in that bracket.*

Proof sketch. Entry and intermediation costs cancel between the household budget and goods clearing. Therefore $\Phi(Y_1) = K + (1-\tau)\omega Y_1 - A_2(Y_1)$, where K collects terms independent of Y_1. If $Y' > Y$, every household's cash-on-hand weakly increases. Monotone saving policies imply $A_2(Y') \geq A_2(Y)$, including when the saving policy jumps upward at the entry threshold. It follows that

$$H(Y') - H(Y) = [1 - (1-\tau)\omega](Y' - Y) + A_2(Y') - A_2(Y) > 0.$$

The existence statement then follows from continuity and the intermediate value theorem. □

In a model without debt-entry costs, $\Phi'(Y_1)$ equals $(1-\tau)$ times the income-weighted average MPC of the households whose saving decision moves with output. Each iteration of Φ therefore adds one round of induced spending, and the fixed point sums the within-period Keynesian-cross multiplier. The debt-entry policy instead has jumps at cash-on-hand thresholds x_h^*. The associated boundary-flow term can make $A_2'(Y_1)$ larger than the within-branch bound, so the Keynesian-cross map need not be a contraction even though the residual remains strictly increasing.[11]

Step 3: Prices and continuation. Given Y_1^*, the wage Phillips curve, Taylor rule, and Fisher equation determine the remaining period-1 prices. Expected inflation π_2 comes from the linear continuation evaluated at the continuation state. Because that state carries r_1 in its lagged-rate coordinate, a closed-form two-variable linear system determines (r_1, π_2).[12] Debt satisfies bond clearing, $B_1 = A_2$, equivalently the government budget identity implied by the other market-clearing conditions. The replaced distribution row delivers $(F_2, m_2^B) = \mathcal{T}(F_{ss}, m_{ss}^B; \bar{x} + dc_1, J_1, M_1^{0^-}, \iota(Y_1^*, \bar{T} + \Delta), \bar{R})$, and the state

$$s_2 = [\, S_2 - S_{ss}; \;\; r_1 - \bar{r}; \;\; 0 \,],$$

where S_2 contains both F_2 and m_2^B. This state is passed to the linear continuation (21) for all $t \geq 2$: from period 2 on, every row of the system is again its linearized version, but it operates on a state that contains the exact period-1 distribution shift.

The one-period update corrects the impact distribution and then hands that distribution to the linear continuation. NDU(K) extends the same idea to later quarters in which the distribution remains far from steady state.[13]

[11]The baseline implementation therefore solves $H(Y_1) = 0$ by safeguarded Newton/bisection rather than relying on Picard iteration.

[12]The 2×2 system is

$$r_1 = \bar{r} + \phi_\pi \kappa_w \widehat{g} + (\phi_\pi \bar{\beta} - \bar{R})\pi_2, \quad \pi_2 = g_s^{(\pi)} s_2(r_1),$$

Here $\widehat{g} = v'(Y_1) - (1-\tau)u'(C_1)$ is the wage-Phillips gap. The equations used are the Phillips curve $\pi_1 = \kappa_w \widehat{g} + \bar{\beta}\,\pi_2$, the Taylor rule $i_1 = \bar{r} + \phi_\pi \pi_1$, and the linearized Fisher equation $r_1 - \bar{r} = (i_1 - \bar{r}) - \bar{R}\,\pi_2$. The period-2 state $s_2(r_1)$ depends affinely on r_1 through its lagged-rate coordinate.

[13]We also experimented with a sequence space version of the algorithm, see Appendix E. Since the standard sequence space approach gains its feasibility from restricting itself to the explicit treatment of aggregates (and only propagates the distribution internally), there is no direct analogue to NDU, but only the possibility to solve some periods nonlinearly. This cannot propagate the distributional adjustment and thus creates stronger deviations to the nonlinear solution after the impact period.

5.2 NDU(K): Window Lengths

NDU(K) extends the exact treatment from one period to a window of K periods. The idea is simple: because the distribution is a state, the method moves it with the exact update for $t = 1, \ldots, K$, one quarter at a time. Equivalently, the two row replacements of Section 5—the distribution rows and the aggregate-savings row—are applied in each of the first K quarters rather than only on impact. Household policies remain first-order throughout the window. The nonlinear movement of the distribution captures the part of the aggregate saving response that the linear update misses.

The window pass. For $t = 1, \ldots, K$, repeat the same within-quarter calculation, but start from the quantities solved in the previous quarter rather than from the steady state:

1. *Inputs.* Beginning-of-period distribution coordinates $S_t = (F_t, m_t^B)$ (with $S_1 = S_{ss}$); the gross return $R_t = \bar{R} + (r_{t-1} - \bar{r})$, fixed by the previous quarter's price closure; and the transfer level from the fiscal rule on realized debt, $T_t = \bar{T} + \Delta\,\mathbb{I}[t{=}1] - \phi_B \bar{r}\,(B_{t-1} - \bar{B})$.

2. *First-order policy at the corrected state.* Evaluate $(dc_t, dJ_t, dM_t^{0^-}) = (g_s^{(c)}, g_s^{(J)}, g_s^{(M)}) s_t$ with $s_t = [\,S_t - S_{ss};\, r_{t-1} - \bar{r};\, \varepsilon_t\,]$, and set $J_t = \bar{J} + dJ_t$, $M_t^{0^-} = \bar{M}^{0^-} + dM_t^{0^-}$. This is the linear decision rule evaluated at the *exactly transitioned* distribution and is the natural generalization of (23).

3. *Exact scalar clearing.* Solve Y_t from the goods-market condition with saving $A[\mathcal{T}(F_t, m_t^B; \bar{x} + dc_t, J_t, M_t^{0^-}, \iota(Y_t, T_t), R_t)]$. Proposition 2 applies with F_t in place of F_{ss}: as long as the fixed saving maps remain nondecreasing, the scalar residual is strictly increasing. We use the same safeguarded Newton/bisection routine in every quarter.

4. *Advance exactly.* Set $(F_{t+1}, m_{t+1}^B) = \mathcal{T}(F_t, m_t^B; \bar{x} + dc_t, J_t, M_t^{0^-}, \iota(Y_t^*, T_t), R_t)$, obtain debt from bond clearing and r_t from the price closure, and use these values as the next quarter's inputs.

After quarter K the state $s_{K+1} = [\,S_{K+1} - S_{ss};\, r_K - \bar{r};\, 0\,]$ is passed to (21) exactly as before.

The only forward-looking term. Everything in the window pass is backward-looking except expected inflation $\mathbb{E}_t \pi_{t+1}$ in the Phillips-curve/Fisher block. We solve this K-vector by iteration. The distribution is not another fixed-point unknown: for every candidate inflation path, the exact DEGM map advances it once and returns the associated consumption, intermediation cost, and entry flow.

We initialize expected inflation with the linear path, run the window, replace each expected-inflation value with the inflation rate implied by that run, and repeat until the values stop changing. At the window boundary, π_{K+1} comes from the linear continuation at the continuation state. Thus inflation is the only intertemporal fixed-point variable at each window date, even though it depends on the full distribution through aggregate consumption.

What the window does and does not correct. The window corrects the distribution movement in every quarter it covers. It does not re-solve household policies. Policies inside the window come from g_s, so their response to future prices is still first-order. This leaves a separate error from using first-order household expectations. Section 6 measures the size of this remaining error.[14]

6 Quantitative Results

This section connects the transfer-induced shift in the distribution to the aggregate transfer response. We proceed in three steps. First, we measure how transfers move households across zero liquid wealth in the model and the SCF. Second, we hold prices fixed and isolate the resulting nonlinearity in household consumption. Third, we allow prices and income to adjust and compare NDU with the linear and fully nonlinear general-equilibrium solutions.

6.1 The Zero-Crossing Flow

For type θ, a transfer moves the mass

$$F_\theta(\bar{x}_\theta) - F_\theta(\bar{x}_\theta - \Delta)$$

from the high-MPC branch to the low-MPC branch. The exact policy boundary $\bar{x}_\theta$ differs across households and is not observed in survey data. We therefore use zero liquid wealth as a common empirical boundary. The *zero-crossing curve* reports, for each transfer size, the share of households that begins with weakly negative liquid wealth and ends with strictly positive liquid wealth after its consumption response.[15]

We do not assume that zero is every household's exact policy kink. It is a common observable balance-sheet boundary, and in every model variant households at or below zero have a substantially higher average MPC than households above it. The comparison is reported in Table B.7. The crossing curve is therefore a diagnostic of the distributional source of curvature, not a sufficient statistic for aggregate consumption.

Taking this statistic to the model first requires household debt. The Bianchi–Kaplan calibration has a zero borrowing limit and therefore cannot represent the 17.6 percent of SCF households with negative liquid wealth. It instead places households that would like to borrow at a mass point at zero. A transfer moves much of this mass at once, so the model cannot generate the gradual flow from negative to positive wealth seen in the data.

We study this issue with a ladder of four models. *No debt* is the Bianchi–Kaplan calibration. *Debt access* adds a borrowing-rate wedge and a fixed debt-entry cost while retaining the original income process and equal preference-type masses. *Debt + preference recalibration* keeps that income process but recalibrates the discount-factor support, type masses, and debt parameters. *Debt + preference and income-process recalibration* also recalibrates the income process and is our baseline debt model. Appendix B reports the parameters, targets, and fit of all four variants.

[14] One can also update household expectations inside the window. We leave this extension for future work.

[15] Appendix B.5.2 reports the exact construction for each model variant.

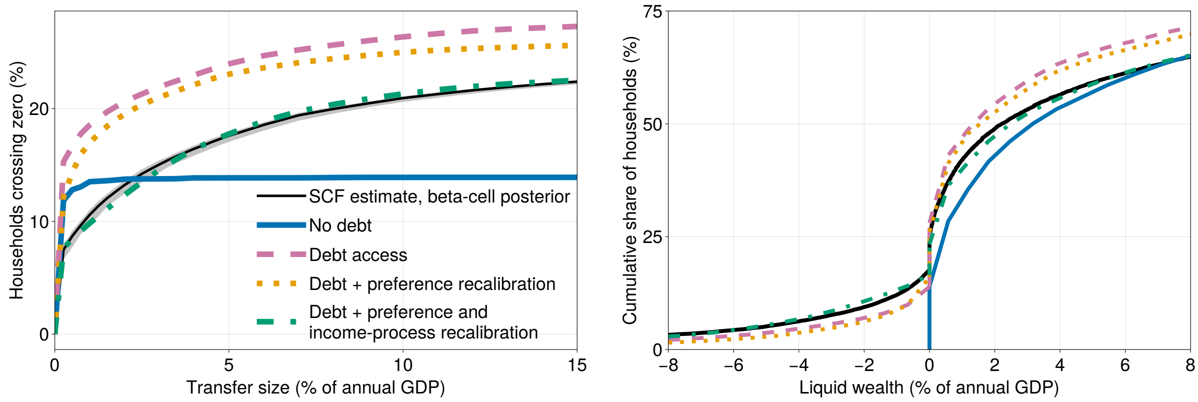

(a) Zero-crossing share

(b) Liquid-wealth CDF around zero

Figure 1: **Zero-crossing shares and liquid-wealth CDFs around zero.**

Notes: Panel (a) shows the share of households that moves from weakly negative to strictly positive liquid wealth; the threshold is exactly zero. The black SCF line is the pointwise mean of four estimates obtained by applying each model's policies and income process to the SCF distribution; the shaded band spans their pointwise minimum and maximum. Panel (b) shows stationary liquid-wealth CDFs over $[-8, 8]$ percent of annual GDP; vertical segments show point masses at zero. Colors and line styles identify the same variants in both panels. Appendices B.5 and B.5.2 give definitions and construction details.

Panel (a) of Figure 1 shows the zero-crossing curves. In the no-debt model, the zero mass crosses abruptly, after which its curve becomes nearly flat. Allowing debt creates a negative tail, and the subsequent recalibrations move the model curve progressively closer to the SCF curve. The SCF curve is constructed by applying the policy functions from the four model variants to the observed liquid-wealth distribution and computing the resulting post-transfer shifts. The implied SCF crossing curve barely changes across variants. It is therefore, to a close approximation, a direct reading of the liquid-wealth distribution around zero rather than a consequence of a particular model policy.

Panel (b) shows that the baseline debt model best matches the liquid-wealth distribution around zero. Among the debt specifications, the RMSE across 14 distributional moments falls from 4.68 percentage points in the debt-access model to 0.96 percentage points in the baseline debt model (Table B.6).

Debt alone is not enough to match the local shape of the distribution. The SCF contains relatively little shallow debt but a dense region of small positive wealth. A borrowing wedge changes desired debt positions but does not create a discrete barrier to taking a small loan. A fixed entry cost does. Households close to the margin remain at zero or just above it, while households that pay the cost choose a nontrivial negative position. The boundary-flow term (13) that we derive analytically illustrates the importance of placing mass on both sides of the threshold in line with the data for generating the correct crossing shares in the model.[16]

[16] Appendix B.3 compares this mechanism with a wedge-only model. We interpret the entry cost as a reduced form for debt-access margins; the data identify the asymmetry around zero, not a particular institutional source of that asymmetry.

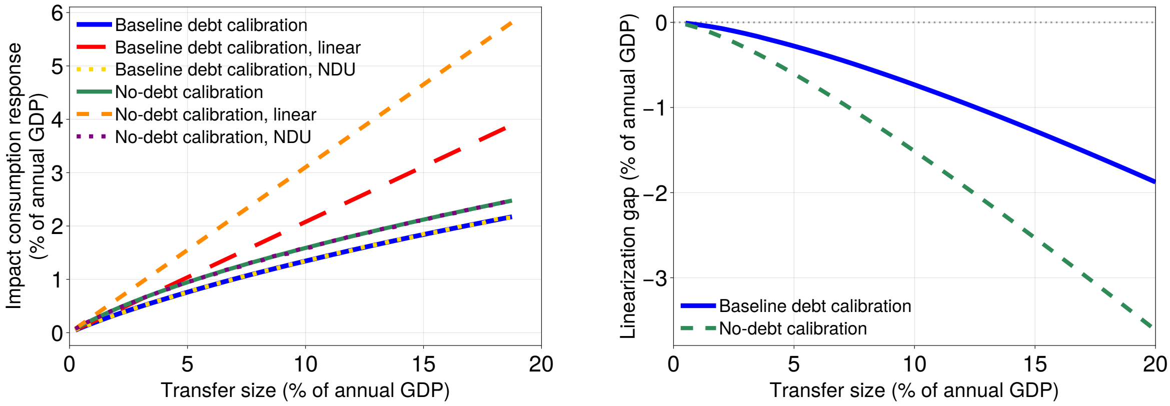

(a) Impact response

(b) Linearization gap

Figure 2: **Partial-equilibrium impact consumption responses.**

Notes: Prices and policies remain at their steady-state values. Panel (a) compares the nonlinear and linear impact responses in the baseline debt and no-debt models. The Nonlinear DEGM Update (NDU) computes the baseline-model response from the shifted distribution using (27) and the resource constraint. Panel (b) reports nonlinear minus linear within each model; negative values mean that the linear approximation overstates the response. Transfers, consumption responses, and gaps are percentages of annual GDP.

6.2 Partial-Equilibrium Responses

Prices remain at their steady-state values, and households receive an unanticipated one-time transfer of size Δ. Because the continuation problem is unchanged, the impact consumption and saving policies are exactly their steady-state policies, c_θ and g_θ. This experiment isolates the effect of moving the initial distribution across policy boundaries.

The impact response can be computed in two equivalent ways:

$$
\begin{aligned}
C(\Delta) &= \int_\Theta \int c_\theta(x+\Delta)\, dF_\theta(x)\, \mu(d\theta) && (26)\\
&= \mathcal{X}(\Delta) - \int_\Theta \int g_\theta(x+\Delta)\, dF_\theta(x)\, \mu(d\theta) - EC(\Delta) \\
&= \mathcal{X}(\Delta) - \int_\Theta \int a'\, d\widetilde{F}^a_{\theta,\Delta}(a')\, \mu(d\theta) - EC(\Delta), && (27)
\end{aligned}
$$

where $\mathcal{X}(\Delta) = \int_\Theta \int (x+\Delta)\, dF_\theta(x)\, \mu(d\theta)$ is aggregate cash-on-hand including the transfer. The first line integrates the consumption policy at shifted cash-on-hand over the stationary distribution. The second line uses the household budget identity, including the entry cost paid by non-waived entrants. The third line changes variables from cash-on-hand to end-of-period assets. The distribution $\widetilde{F}^a_{\theta,\Delta}$ is exactly the shifted distribution produced by (6)–(7). In the debt-entry calibration, $EC(\Delta) = \kappa e(\Delta)$, with the entry flow defined in (50).

The two expressions are identical, but they use different intermediate calculations. Equation (26) interpolates the consumption policy while keeping the initial distribution fixed. Equation (27) interpolates the updated CDF on a fixed asset grid and then recovers consumption from the resource constraint. Their agreement therefore tests the DEGM change of variables. The second representation is also the one needed for NDU: the shifted distribution is an aggregate state that can be passed directly to the linear state-space model.

Panel (a) of Figure 2 first compares the two models under the same nonlinear solution. The

Table 1: First-quarter output response: state-space methods

Transfer size	Nonlinear	Linear	NDU(1)	NDU(6)
0.5% GDP	0.390	0.598	0.383	0.383
1.0% GDP	0.698	1.196	0.685	0.685
2.0% GDP	1.293	2.392	1.273	1.273
5.0% GDP	2.775	5.980	2.724	2.724
10.0% GDP	4.481	11.961	4.523	4.523

Notes: Sticky-wage HANK economy, state-space solution, $N_A = 120$, log-log grid, $T = 150$. Output responses in percent deviation from steady state. "Nonlinear" is the fully nonlinear path solve; "Linear" scales the first-order state-space impulse response. NDU(1) and NDU(6) share the identical period-1 solve, so their impacts coincide by construction; they differ over quarters 2–K.

baseline debt model predicts less consumption at every reported transfer size. The difference grows from 0.04 percent of annual GDP at a 0.5-percent transfer to 0.31 percent at a 20-percent transfer. This is the partial-equilibrium implication of the gradual movement of indebted households into the lower-MPC positive-wealth region.

Panel (a) also compares solution methods. The first-order response is the aggregate MPC times the transfer and therefore misses the declining MPC along the crossing curve. By contrast, the direct integral and the Nonlinear DEGM Update[17] are visually indistinguishable: their largest difference is 0.006 percent of annual GDP over the reported range. Panel (b) isolates the resulting linearization gap. It is negative in both models and larger without debt, meaning that the linear approximation overstates the nonlinear response. In the baseline debt model, the overstatement is 55 percent at a 10-percent transfer and 82 percent at a 20-percent transfer. The DEGM approach of shifting the distribution captures the distributional curvature exactly, without needing to approximate the moved cutoff via expansion. Appendix D.3 reports the same experiment, for both calibrations, in table form and adds the second-order approximation, including the policy-boundary terms derived in Appendix D.1.

The partial-equilibrium exercise establishes two points. The transfer response is less concave in the baseline debt model than in the no-debt model, but the linear approximation still substantially overstates large transfers. At the same time, the DEGM distribution update reproduces the nonlinear response almost exactly. We next embed this update in general equilibrium, where output feeds back into household income and the corrected distribution affects the entire transition path.

6.3 General-Equilibrium Responses

All state-space results use a 120-point log-log asset grid on $[-1, 4000]$, plain CDF state coordinates, and first-order household policies inside the NDU window.[18] The fully nonlinear benchmark solves a transition path of length $T = 150$ using the same DEGM distribution update. The linear, NDU, and nonlinear solutions therefore differ in how they treat the transfer-induced

[17] In partial equilibrium, the Nonlinear DEGM Update refers simply to computing aggregate consumption by integrating over the shifted end-of-period asset distribution, as in Equation (27). No iteration is needed in that case.

[18] Bayer et al. (2026) show that DEGM converges much faster in the distribution than the lottery method.

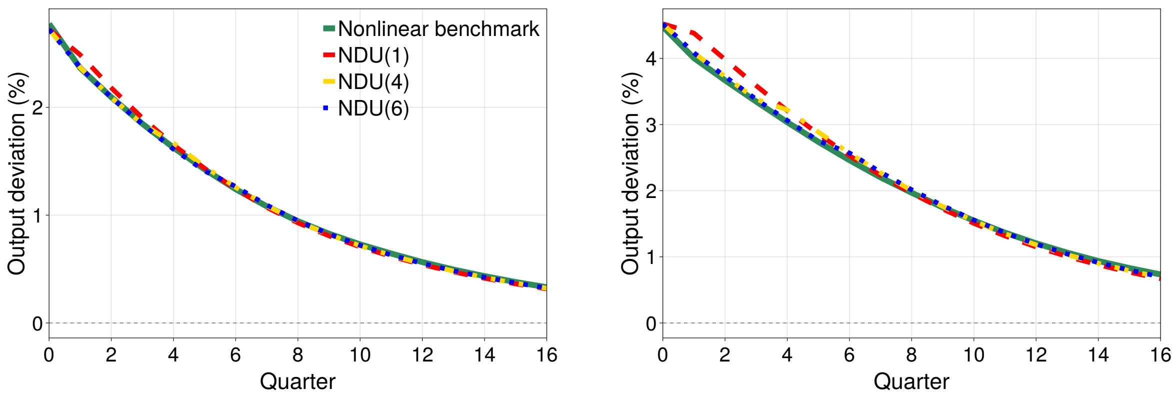

(a) 5% transfer

(b) 10% transfer

Figure 3: **Output impulse responses: nonlinear benchmark and Nonlinear DEGM Update.**

Notes: Panels (a) and (b) show one-time transfers equal to 5 and 10 percent of annual GDP. Green is the fully nonlinear benchmark. NDU(1), in red, updates the first quarter nonlinearly and uses the linear law of motion thereafter. NDU(4) and NDU(6), in orange and blue, extend the nonlinear update through quarters 4 and 6. Output is measured as a percentage deviation from steady state.

movement of the distribution, not in the underlying discretization.

Table 1 begins with the first-quarter output response. The linear solution substantially overstates the effect of large transfers: relative to the nonlinear benchmark, its error is 116 percent at a 5-percent transfer and 167 percent at a 10-percent transfer. NDU(1) replaces the first distribution transition and the corresponding goods-market equation by their exact nonlinear versions. This removes nearly all of the impact error. Across transfer sizes up to 10 percent of annual GDP, the largest remaining difference is 5.1 basis points.

Panels (a) and (b) of Figure 3 show the full transition paths after transfers equal to 5 and 10 percent of annual GDP. In both panels, NDU(1) matches the nonlinear solution in the first quarter, but its linear continuation lies above the benchmark during the next several quarters. The separation is larger for the 10-percent transfer in Panel (b). The reason is that the distribution remains far from steady state after impact. Handing it immediately to the linear law of motion applies steady-state, tangent MPC behavior too early. NDU(4) and NDU(6) continue to move the distribution with the exact update while this adjustment is largest, bringing the output path progressively closer to the nonlinear benchmark in both panels.

Table 2 summarizes the accuracy and cost of this extension. The accuracy measure is the largest absolute output difference from the nonlinear benchmark over the first five quarters. NDU(1) removes the large impact error but leaves a visible continuation error. At a 10-percent transfer, increasing the exact-update window from one to six quarters lowers the peak early-path error from 0.377 to 0.079 percentage points. At smaller transfers, the remaining differences are correspondingly smaller.

Table 2: Early-path accuracy and per-shock cost of the Nonlinear DEGM Update

	Peak early-path error (pp)			Per-shock solve time (s)			
Transfer size	$K = 1$	$K = 4$	$K = 6$	NDU(1)	NDU(4)	NDU(6)	Nonlinear
0.5% GDP	0.007	0.003	0.003	0.35	1.53	3.29	262.05
1.0% GDP	0.041	0.009	0.004	0.32	1.59	3.22	8.03
5.0% GDP	0.123	0.043	0.015	0.33	1.73	3.61	326.42
10.0% GDP	0.377	0.187	0.079	0.32	1.42	3.24	187.76

Notes: Sticky-wage HANK economy, state-space solution, $N_A = 120$, log-log grid, $T = 150$. Peak early-path error is the maximum absolute output deviation from the nonlinear benchmark over quarters 1–5, in percentage points. NDU times are per-transfer window solves, excluding the shared linear state-space solution. Nonlinear times are the corresponding conditional benchmark cost: linear initial path plus Newton solve; the computation time of the sequence-space Jacobian is excluded. Nonlinear times vary with Newton convergence and include cases that stop with discretization residuals of roughly 10^{-6}; NDU times are stable across shock sizes. All times are measured after untimed warm-up evaluations with 8 Julia and 8 OpenBLAS threads.

The computational cost is stable across transfer sizes. NDU(6) takes about three seconds per experiment, while the conditional nonlinear benchmark takes between 8 and 326 seconds. Averaged across the four transfer sizes, the nonlinear transition costs 196 seconds and NDU(6) costs 3.3 seconds. The expected-inflation iteration requires 12–13 window passes for NDU(4) and 20 passes for NDU(6). Even with these passes, the nonlinear part of NDU remains a sequence of distribution updates and scalar goods-market solves. It does not solve the full nonlinear transition system. This distinction becomes more valuable as the household state space grows.

7 Conclusion

Large fiscal transfers are a challenge for linear methods in HANK models (Bianchi and Kaplan, 2026). The source of this nonlinearity is not policy-function curvature per se. Aggregate consumption is an integral over the household distribution, and a transfer moves the high-MPC threshold through that distribution. The key quantity is therefore the fraction of households that cross the threshold. The DEGM representation of the KFE directly tracks this fraction by expressing the nonlinear movement of the distribution as a closed-form update.

A state-space solution makes this update particularly useful because the corrected distribution remains in the system. Our Nonlinear DEGM Update (NDU) replaces the linear distribution transition and goods-market equation by their exact counterparts over a short window. Each window quarter requires one CDF update and one scalar goods-market solve. In the fiscal-transfer experiments, NDU matches the fully nonlinear solution on impact. With a six-quarter window, it also closely follows the nonlinear IRFs over the full horizon, while adding only seconds per experiment after a one-time linearization.

Beyond the solution method, we contribute a measurable target—the zero-crossing curve—and show that household debt changes how transfer effects decline with size. The curve records the share of households that a transfer moves from weakly negative to positive liquid wealth. In the SCF, shallow debt is sparse and negative positions are dispersed away from zero, so the crossing share rises gradually with transfer size. Standard no-debt and borrowing-wedge models miss this shape, whereas a stochastic debt-entry cost reproduces it. The resulting baseline debt model remains in the conventional MPC range for transfers around $1,000 but predicts a smaller response to large transfers because indebted households deleverage gradually into a lower-MPC region. Large-transfer nonlinearities therefore depend not only on average MPCs, but also on the distribution of liquid wealth around zero.

A Model

This appendix presents the economic model underlying the quantitative exercises. The model builds on the sticky-wage HANK economy in Bianchi and Kaplan (2026) and extends it to allow household debt.

The general formulation nests the four model variants used in the empirical comparison. Appendix B defines and calibrates these variants and documents the empirical mapping, while Appendix C provides the numerical implementation. We first describe the aggregate environment, then the income process and the household decision problem.

A.1 Aggregate Environment

Time is discrete and starts at $t = 1$.

Production. A representative competitive firm produces the final good using labor according to the linear technology

$$Y_t = N_t.$$

Under the competitive linear technology, the real wage is normalized to one and firm profits are zero.

Households. Households have CRRA utility over consumption and separable disutility from labor, with

$$u'(c) = c^{-\sigma}, \qquad v'(n) = \chi n^{1/\nu}.$$

They trade a one-period liquid asset whose net supply is government debt. The gross return on nonnegative beginning-of-period assets is $R_t = 1 + r_{t-1}$, while negative positions face the gross return $R_t + \rho_0$.

Let A_t denote aggregate beginning-of-period household assets and $A_t^- \leq 0$ the aggregate of negative individual asset positions. The borrowing wedge absorbs

$$IC_t = -\rho_0 A_t^- \geq 0$$

units of goods. Let e_t denote the mass of households that pay the fixed debt-entry cost and let

$$EC_t = \kappa e_t$$

be the associated resource cost. The income process, the household decision problem, and the aggregation of e_t are defined below.

Wage Phillips curve. Nominal wages are set by a representative labor union subject to quadratic adjustment costs. The union uses the population-weighted mean discount factor

$$\bar{\beta} = \sum_h p_h \beta_h.$$

Optimality condition implies the wage Phillips curve

$$\pi_t = \kappa_w \left[v'(N_t) - (1-\tau)u'(C_t) \right] + \bar{\beta}\, \pi_{t+1}, \tag{28}$$

where τ is the proportional labor-income tax.

The wage Phillips curve depends on aggregate consumption, which cannot be set equal to output net of government expenditure in models with debt since intermediation and debt-entry costs absorb goods: $C_t \neq Y_t - G$ (G is constant government expenditure, see below).

Monetary and fiscal policy. The monetary authority sets the net nominal interest rate i_t according to a Taylor rule. The Taylor rule and Fisher equation determine the nominal and real interest rates:

$$i_t = \bar{r} + \phi_\pi \pi_t, \qquad 1 + r_t = \frac{1 + i_t}{1 + \pi_{t+1}}, \tag{29}$$

where $\bar{r}$ is the steady-state net real interest rate.

The government purchases G units of the final good, taxes labor income at rate τ, makes the lump-sum transfer T_t, and issues one-period debt. The transfer experiment is implemented through

$$T_1 = \bar{T} + \Delta, \qquad T_t = \bar{T} - \phi_B \bar{r} \left(B_{t-1} - \bar{B} \right), \quad t \geq 2, \tag{30}$$

where Δ is the one-time transfer innovation and $\bar{T}$ and $\bar{B}$ denote steady-state transfers and government debt.

Neither intermediation nor entry costs are government revenue. The government budget constraint is therefore

$$B_t = (1 + r_{t-1}) B_{t-1} + G + T_t - \tau Y_t, \tag{31}$$

where B_t is end-of-period government debt.

Aggregation and market clearing. Aggregating the household budget constraints and using the normalization that average labor efficiency equals one gives

$$C_t = R_t A_t + \rho_0 A_t^- + (1-\tau) Y_t + T_t - EC_t - A_{t+1}. \tag{32}$$

Goods and bond-market clearing require

$$Y_t = C_t + G + IC_t + EC_t, \qquad A_{t+1} = B_t. \tag{33}$$

Appendix B.1 reports the common macro calibration and steady-state normalization used in the quantitative exercises.

A.2 Income Process and Household States

Household states. Households differ permanently in their discount factors and face persistent and transitory labor-income risk. Let $b \in \{1, \ldots, N_\beta\}$ index a permanent discount-factor type with discount factor β_b and population mass p_b^β, and let $z_t \in \mathcal{Z}$ denote the persistent income state. We combine these states as

$$h_t = (b, z_t).$$

with a combined transition matrix $\Pi_{h,h'}$ and stationary population mass p_h.

Labor-income process. Following Bianchi and Kaplan (2026), we use the income-risk specification discussed by Kaplan and Violante (2022): log labor efficiency is the sum of an infrequently changing persistent component and an infrequent transitory component. In our notation,

$$\log \ell_t = z_t + \epsilon_t - \log \mathbb{E}[\exp(z_t + \epsilon_t)], \tag{34}$$

so average labor efficiency is one. The expectation is taken under the stationary distributions of the persistent and transitory components.

Let $Q(\phi_z, \sigma_\eta^2)$ denote the transition associated with the AR(1) process

$$z_t = \phi_z z_{t-1} + \eta_t, \qquad \eta_t \sim N(0, \sigma_\eta^2).$$

To capture infrequent innovations, the transition used in the model is

$$P_z = (1 - \lambda_\eta) I + \lambda_\eta Q(\phi_z, \sigma_\eta^2). \tag{35}$$

Thus, λ_η is the probability that the persistent state is refreshed; otherwise it remains unchanged.

A transitory shock occurs with probability λ_ϵ, so that

$$\epsilon_t \sim N(-\sigma_\epsilon^2/2, \sigma_\epsilon^2) \quad \text{with probability } \lambda_\epsilon, \qquad \epsilon_t = 0 \quad \text{with probability } 1 - \lambda_\epsilon. \tag{36}$$

Let k index the transitory realizations ϵ_k, with probabilities ω_k. The corresponding labor-efficiency realization is

$$w_{h,k} = \exp\left(z + \epsilon_k - \log \mathbb{E}[\exp(z_t + \epsilon_t)]\right), \tag{37}$$

where $w_{h,k}$ depends on h only through its persistent-income component z. Pre-tax labor income in state (h, k) is

$$y_{hk,t} = w_{h,k} Y_t. \tag{38}$$

The normalization in (34) implies

$$\sum_h p_h \sum_k \omega_k w_{h,k} = 1, \tag{39}$$

so aggregate labor income equals Y_t. The transitory realization k is not carried across periods and therefore is not part of the persistent household state h. Appendix B reports the parameter values and discretization used for the four model variants.

A.3 Household Decision Problem and Debt Access

Debt access and timing. Households enter period t with beginning-of-period assets a and choose end-of-period assets $a' \geq \underline{a}$, where $\underline{a} \leq 0$. Beginning-of-period borrower status provides a reduced-form representation of persistent credit access without adding credit history as a separate individual state. A household is an incumbent borrower exactly when $a < 0$. Incumbent borrowers may continue borrowing or exit freely.

At the beginning of each period, a non-borrower with $a \geq 0$ enters an i.i.d. waiver-lottery. Let q denote the waiver-lottery outcome, where F denotes the fee-free draw and K the draw under which entering debt is costly:

$$q = F \quad \text{with probability } \lambda, \qquad q = K \quad \text{with probability } 1 - \lambda.$$

With a waiver, the household may enter debt without paying the fixed cost. Without a waiver, choosing $a' < 0$ costs κ units of consumption goods. Entry is the only event that incurs the cost: continuing debt and returning to nonnegative assets are free.

Let $y_{hk,t} = w_{h,k} Y_t$ denote pre-tax labor income in state (h, k). Define disposable nonasset resources as

$$\iota_{hk,t} = (1 - \tau) y_{hk,t} + T_t. \tag{40}$$

When time or state indices are suppressed, we write these objects as y_{hk}, y, or ι_{hk}.

Define the gross saving return $R_t = 1 + r_{t-1}$ and the return on beginning-of-period assets by

$$R_t(a) = R_t + \rho_0 \mathbf{1}\{a < 0\}. \tag{41}$$

Cash-on-hand is then

$$x_{hk,t}(a) = R_t(a) a + \iota_{hk,t}. \tag{42}$$

Household problem. A household in state h, with discount factor β_h, solves

$$V^q_{t,h}(a) = \max_{a' \geq \underline{a}} \left\{ u(c) + \beta_h \mathbb{E}_t \left[V_{t+1,h'}(a') \right] \right\}, \tag{43}$$

$$\text{s.t.} \qquad c + a' = x - \kappa \mathbf{1}\{a \geq 0,\ a' < 0,\ q = K\}. \tag{44}$$

Here $x = x_{hk,t}(a)$ is the present period's cash-on-hand.

The choice a' determines next period's borrower status. If $a' < 0$, the household enters the next period as an incumbent borrower and the new waiver draw is irrelevant; denote the corresponding value by

$$V^B_{t+1,h'}(a').$$

If $a' \geq 0$, the household enters as a non-borrower and draws a new waiver before choosing. Its ex-ante value is therefore

$$V^N_{t+1,h'}(a') = \lambda V^F_{t+1,h'}(a') + (1 - \lambda) V^K_{t+1,h'}(a'). \tag{45}$$

Thus,

$$V_{t+1,h'}(a') = \begin{cases} V^B_{t+1,h'}(a'), & a' < 0, \\ V^N_{t+1,h'}(a'), & a' \geq 0. \end{cases} \tag{46}$$

The expectation in (43) is thus over the next persistent state h' and transitory-income realization k', and the waiver-lottery if $a' \geq 0$.

Policy branches. Incumbent borrowers and waived non-borrowers face the same unrestricted problem; denote its asset policy by $g^F_h(x)$. A non-waived non-borrower compares three branches. For low cash-on-hand, it pays the cost and enters debt; over an intermediate region, it bunches at $a' = 0$; above that region, it saves.

Suppressing time indices, let $x_{0,h}$ denote the unrestricted cash-on-hand threshold associated with $a' = 0$, so that $g^F_h(x_{0,h}) = 0$. The non-waived policy is

$$g^K_h(x) = \begin{cases} g^E_h(x), & x < x^\star_h, \\ 0, & x^\star_h \leq x \leq x_{0,h}, \\ g^S_h(x), & x > x_{0,h}, \end{cases} \tag{47}$$

where g^E_h is the debt-entry branch and g^S_h is the positive-saving branch. The fixed cost shifts the cash-on-hand required to reach a given negative asset choice upward by κ. The debt-entry threshold x^*_h is determined by value matching:

$$V^{K,E}_h\left(\frac{x^*_h - \iota_{hk}}{R}\right) = V^{K,0}_h\left(\frac{x^*_h - \iota_{hk}}{R}\right), \tag{48}$$

where the second superscript in $V^{K,E\,\vee\,0}_h$ denotes whether the household enters into debt, $a' < 0$, or remains at zero, $a' = 0$.

The asset policy can jump at $x^\star_h$: paying a fixed cost is worthwhile only if the household chooses a nontrivial debt position. This discontinuity is the household mechanism that produces relatively little shallow debt and a dense small-positive region in the stationary distribution. We focus on the ordering

$$x^\star_h < x_{0,h},$$

which holds for all calibrated household states.

The no-debt model is nested by setting $\underline{a} = 0$, which makes the debt-entry block inactive.

A.4 Distribution, Aggregation, and Equilibrium

Let F_t denote the beginning-of-period joint distribution of assets and persistent household states, with

$$F_{t,h}(a) = \Pr(a_t \leq a,\ h_t = h).$$

Given F_t and the household policies, the distribution F_{t+1} is induced by households' end-of-period asset choices, the transition matrix Π, and the transitory-income and waiver draws.

Aggregate total and negative assets are

$$A_t = \sum_h \int_{\underline{a}}^{\infty} a \, dF_{t,h}(a), \qquad A_t^- = \sum_h \int_{\underline{a}}^{0} a \, dF_{t,h}(a). \tag{49}$$

Aggregate consumption is obtained by integrating individual consumption over the household distribution and the idiosyncratic draws. The mass of non-waived non-borrowers that enter debt and its resource cost are

$$e_t = (1-\lambda) \sum_h \sum_k \omega_k F_{h,t}^{N,<}\left(\frac{x_{h,t}^* - \iota_{hk,t}}{R_t}\right), \qquad EC_t = \kappa e_t, \tag{50}$$

where we write $F_{h,t}^{s,<}(a) := F_{h,t}^{s}(a^-) = \Pr(a_t < a, s_t = s, h_t = h)$ for the corresponding strict CDF on side $s \in \{B, N\}$. Appendix C describes the CDF representation used to implement the distribution law of motion.

Given initial conditions and the transfer innovation Δ, an equilibrium is a sequence of household value and policy functions, distributions, and aggregate quantities such that, for every t:

(i) household value and policy functions solve the problem in Section A.3, given prices, income, and transfers;

(ii) the distribution F_{t+1} is induced by F_t, the household policies, and the idiosyncratic-state transitions;

(iii) aggregate assets, consumption, and the debt-entry flow are consistent with individual choices and the household distribution; and

(iv) production, wage setting, monetary and fiscal policy, the government budget constraint, and goods- and bond-market clearing satisfy the conditions in Section A.1.

B Calibration and Empirical Mapping

This appendix first defines the model variants and reports their macro and micro calibration and fit. It then explains why a fixed debt-entry cost is needed and documents the MPC contrast around zero liquid wealth. Finally, it maps the Survey of Consumer Finances (SCF) into the models. The model ladder organizes a sequence of specification and recalibration steps, while the SCF mapping constructs the empirical objects used in the crossing-share and paired consumption-response comparisons.

Two qualifications are important for interpreting these comparisons. First, the steps in the model ladder are not controlled one-parameter experiments: later variants recalibrate several parameters jointly. Second, the SCF consumption response is necessarily model-informed, because consumption is not observed for the counterfactual transfer. We therefore compare each model with an SCF response evaluated through that same model's policies.

B.1 Model Variants and Macro Calibration

All four variants share the aggregate environment described in Appendix A. A period is a quarter; quarterly GDP is normalized to one, so annual GDP is four. The annual real interest rate is two percent. Utility is logarithmic, and the proportional labor-income tax and steady-state lump-sum transfer are taken from Bianchi and Kaplan (2026). Table B.1 summarizes the common macro calibration.

Table B.1: Common macro calibration

Object	Value	Description
$\bar{Y}$	1	Quarterly output normalization
Real interest rate	2% annually	Steady-state real rate
σ	1	Log utility
ν	1	Frisch elasticity
τ	0.274	Labor-income tax rate
$\bar{T}$	0.15	Steady-state transfer relative to quarterly GDP
κ_w	0.01	Wage Phillips-curve slope
ϕ_π	1.5	Taylor-rule coefficient
ϕ_B	8	Fiscal-rule coefficient
χ	Chosen so that $\bar{\pi} = 0$	Labor-disutility scale

Notes: The equations in Appendix A.1 use the quarterly real interest rate corresponding to the reported annual rate.

All four models have five permanent discount-factor types and eleven persistent-income states. They differ in debt access, the calibration of preferences, and the income process, as summarized in Table B.2.

Table B.2: The four model variants

Model	Definition
A	No household debt (borrowing limit at zero), with the original Bianchi–Kaplan income process.
B1	Adds debt, a borrowing-rate wedge, and the entry-cost/waiver block; retains the original income process.
B2	Retains B1's structure and income process, but recalibrates the discount-factor support, type masses, and debt parameters.
C	Recalibrates the income process, discount-factor distribution, and debt parameters; this is the baseline debt model.

Model A is obtained from the general formulation in Appendix A by imposing $\underline{a} = 0$, which makes the borrowing wedge and the entry-cost/waiver block inactive. Models B1–C allow $\underline{a} < 0$ and use the same debt-capable model structure, but differ in the calibration of the discount-factor distribution, debt parameters, and, for model C, the income process. The population-weighted discount factor $\bar{\beta}$ in the wage Phillips curve is implied by the discount-factor support and type masses reported below.

B.2 Micro Calibration

The micro calibration comprises the income process, discount-factor heterogeneity, debt parameters, and the distributional targets. Appendices A.2 and A.3 describe the corresponding functional forms and household problem. This section reports the discretization, parameter values, and calibration sequence used for the four model variants.

B.2.1 Income-Process Calibration

The persistent component is represented on eleven Rouwenhorst nodes. The Gaussian component of the transitory shock is integrated with five Gauss–Hermite nodes, so the discretized income process has eleven persistent states and six transitory outcomes. The mean correction in (36), followed by the normalization in (34), keeps average labor efficiency fixed when the risk parameters change.

The benchmark process used by A, B1, and B2 is the Bianchi–Kaplan implementation of the Kaplan–Violante specification. C retains the same functional form and the same eleven-by-six discretization but recalibrates four parameters to improve the joint income–wealth and liquid-wealth fit; the persistent innovation variance remains fixed. The complete values are reported in Table B.3.

Table B.3: Income-process parameters by model variant

Specification	ϕ_z	σ^2_η	λ_η	σ^2_ϵ	λ_ϵ
A, B1, B2	0.9878	0.0439	0.2500	0.6376	0.2500
C	0.98213	0.0439	0.2350	0.69196	0.2725

The A–B1 and B1–B2 comparisons hold the income process fixed. B2–C changes the income process together with preferences and debt parameters; it is therefore a sequential fit comparison rather than a one-parameter experiment. Figure B.2 shows the resulting income–wealth sorting directly.

B.2.2 Discount-Factor Heterogeneity and Debt Parameters

All variants have five permanent discount-factor types, but neither their support nor their masses are generally fixed across calibrations. The calibrated values used in the results are reported in Table B.4.

Table B.4: Discount-factor support and type masses by model variant

Model	Discount-factor support	Type masses
A	(0.9599, 0.9676, 0.9758, 0.9835, 0.9911)	(0.2, 0.2, 0.2, 0.2, 0.2)
B1	(0.9600, 0.9677, 0.9759, 0.9836, 0.9912)	(0.2, 0.2, 0.2, 0.2, 0.2)
B2	(0.9720, 0.9785, 0.9793, 0.9857, 0.9920)	(1, 2, 3, 2, 1)/9
C	(0.9659, 0.9723, 0.9791, 0.9856, 0.9919)	(1, 6, 6, 3, 3)/19

B1 preserves A's beta spacing and equal masses but raises every support point by only 8.1×10^{-5}. The A–B1 comparison is therefore extremely close to a pure debt-access experiment, but not literally one. B2 changes both the shape of the support and the masses while retaining the benchmark income process. C changes the beta distribution again while also recalibrating the income process. These preference changes matter both for average MPCs and for the composition of discount types across wealth–income cells used in the paired SCF calculation.

In A, $\underline{a} = 0$, and the debt wedge, entry cost, and waiver probability are all zero. The debt variants use $\underline{a} = -1$. B1 uses

$$(\rho_0, \kappa, \lambda) = (0.05, 0.03, 0.25),$$

B2 uses

$$(\rho_0, \kappa, \lambda) = (0.038, 0.02, 0.45),$$

and C uses

$$(\rho_0, \kappa, \lambda) = (0.0385, 0.0125, 0.13).$$

Thus, the A–B1 step is dominated by making the debt region feasible, whereas B1–B2 and B2–C should be read as sequential specification and recalibration steps rather than as isolated causal effects of individual parameters.

B.2.3 Calibration Targets and Fit

The empirical target is the marginal liquid-wealth distribution in the Survey of Consumer Finances (SCF), pooling the twelve triennial waves from 1989 through 2022 with equal wave weights. Appendix B.5 describes the sample, liquid-wealth definition, implicate treatment, and assignment of households to model income states. Wealth is expressed as a percent of annual GDP. The distributional criterion contains fourteen moments: thirteen CDF levels and the discrete mass at zero, all in percentage points, as reported in Table B.5.

Table B.5: SCF liquid-wealth distribution targets

a	-12	-9	-6	-3	-1	-0.5	0^-
$F^{\text{SCF}}(a)$	1.84	2.79	4.31	7.66	12.11	14.32	17.60
a	mass at 0	0.5	1	2	5	10	15
Target	6.68	36.20	42.00	48.90	59.10	67.50	72.40

The full calibration objective sums squared deviations from the targeted liquid-wealth moments and adds an aggregate-asset target and a one-sided penalty for violating an upper bound on the income-rank/wealth gradient. No MPC moment enters the selection criterion for these four specifications. The discrete mass at zero wealth is targeted separately from $F(0^-)$: the entry cost is intended to govern precisely how mass is split between debt and the kink at zero.

Raw values of the full calibration objective are not directly comparable across models when the number of reachable liquid-wealth moments differs. For model M, let $\mathcal{M}_M$ be its set of reachable distributional moments, let $n = |\mathcal{M}_M|$, and let m_M and m^{SCF} denote the corresponding model and empirical moment values. We therefore report the distributional fit as

$$\text{RMSE}_M = \left[\frac{1}{n}\sum_{m\in\mathcal{M}_M}\left(m_M - m^{\text{SCF}}\right)^2\right]^{1/2}, \tag{51}$$

in percentage points.

Model A has no negative support at any parameter value, so its seven negative-side moments are mechanically unreachable and its reported RMSE uses the seven nonnegative moments. B1–C use all fourteen moments. B1, B2, and C minimize the full calibration objective over their feasible parameter sets. A is calibrated to the aggregate-asset and average-MPC targets of Bianchi and Kaplan (2026) and is evaluated, but not optimized, on the distributional criterion.

Table B.6: Fit to the marginal liquid-wealth distribution

Model	Moments n	RMSE (pp)	RMSE on common seven moments (pp)
A	7	5.94	5.94
B1	14	4.68	6.22
B2	14	3.47	4.34
C	14	0.96	1.07

Notes: The common seven moments are the moments reachable by the no-debt model. A has 100 nonnegative asset nodes; B1–C have the same 100 nonnegative nodes plus 20 negative nodes.

The common-support column clarifies where the improvement comes from. B1 fits the shared region slightly worse than A; its entire gain comes from replacing A's empty debt region with a negative tail. On the squared-error scale, B1 reduces the negative-side discrepancy from 749.9 to 35.6 while increasing the common-support discrepancy from 246.6 to 270.8. B2 and especially C then improve both regions. Debt is therefore indispensable for matching the empirical shape around zero, while the baseline debt model also requires recalibration within the positive region and of the income process.

B.3 Why a Fixed Entry Cost?

The empirical motivation is a sharp asymmetry around zero. Immediately to the right of zero, the SCF contains a dense group of households with small positive liquid-wealth positions.

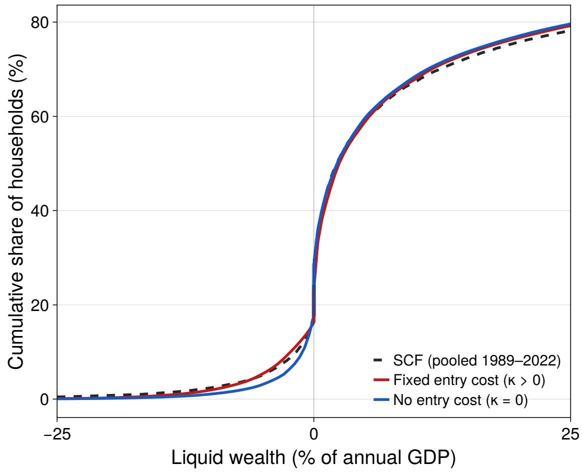

(a) Liquid-wealth CDF

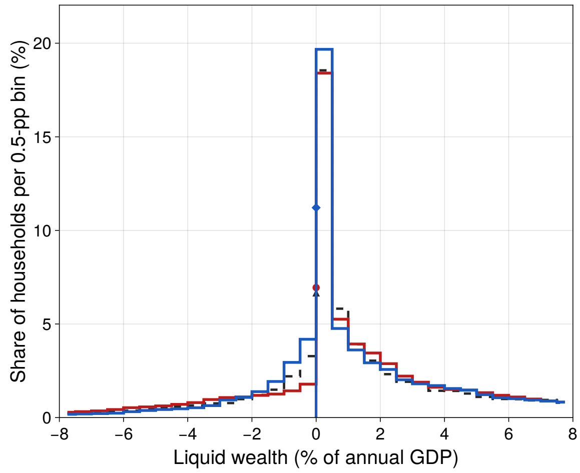

(b) Liquid-wealth density near zero

Figure B.1: **Liquid wealth around zero with and without a fixed entry cost.**

Notes: The entry-cost and wedge-only calibrations are independently optimized under the same objective and evaluated on the same grid. Panel (a) shows the liquid-wealth CDF. Panel (b) shows probability mass in 0.5-percentage-point bins around zero, with point masses at zero displayed separately. The SCF combines a dense small-positive-wealth region with relatively little shallow debt. Grid size 180, with 30 nodes in the negative range.

Immediately to the left, the shallow-debt region is much less densely populated. The distribution therefore does not pass smoothly through zero: many households remain just outside debt, while households that borrow are comparatively less likely to hold an arbitrarily small negative balance. A standard debt model with only a borrowing limit and a spread ρ_0 on negative assets ($\kappa = 0$) does not reproduce this local asymmetry.

The spread ρ_0 is a recurring price on the *stock* of debt. As the desired debt position approaches zero, the associated interest payment also approaches zero. The wedge therefore gives a household with small positive wealth no discrete reason to avoid a small negative position. Changing ρ_0 shifts desired borrowing throughout the negative region, but it cannot create a pronounced difference between the densities immediately on the two sides of zero.

A fixed cost κ, by contrast, is paid when a non-borrower first enters debt and does not vanish with the size of the initial balance. It creates an inaction region: households near the margin prefer zero or a small positive position to paying the cost for a shallow loan. Conditional on paying the cost, borrowing only a negligible amount is unattractive, so entrants choose debt positions farther to the left. The model consequently produces a dense small-positive region, relatively little shallow debt, and a nontrivial deeper debt tail. The waiver probability λ allows some households to enter without paying the cost and controls the strength of this separation.

Panel (a) of Figure B.1 shows that the aggregate borrower share by itself does not distinguish the specifications. The best wedge-only model matches $F(0^-)$ almost exactly: its error is $+0.04$ percentage points, compared with -1.31 for the entry-cost model. Yet its errors at $F(-6)$ and $F(-3)$ are -1.89 and -2.50 percentage points, compared with -0.05 and $+0.79$ with the entry cost. Panel (b) shows the local source of this difference. The wedge-only model concentrates too much probability immediately below and exactly at zero and too little in the deeper negative tail; its zero-wealth mass exceeds the 6.68-percent target by 4.53 points, compared with 0.26

points in the entry-cost model. Thus the wedge-only model gets approximately the right *number* of borrowers but not the observed asymmetry between shallow debt and small positive wealth.

This is a like-for-like calibration comparison. The entry-cost specification, $(\kappa, \lambda, \rho_0) = (0.0125, 0.13, 0.0385)$, and the wedge-only specification, $(\kappa, \rho_0) = (0, 0.056)$, independently minimize the same objective. Although the entry-cost triple equals C's parameters in Appendix B.2.2, this auxiliary robustness specification is not the small-grid model C used in the four-model comparison: it is evaluated on the common 180/30 comparison grid against the independently optimized wedge-only specification.

We interpret (κ, λ) as a reduced form for access margins not explicitly tracked in the model—for example, whether a household already has an available credit line and can draw on it without incurring an origination or search cost. We do not claim that the SCF wealth distribution separately identifies those institutional components. The narrower empirical claim is that an entry hurdle is sufficient to reproduce the local asymmetry at zero, whereas a borrowing wedge is not.

B.4 MPCs Around Zero Liquid Wealth

Section 6.1 uses zero liquid wealth as a common observable balance-sheet boundary. We do not assume that zero is every household's exact policy kink. This subsection instead documents its interpretation as a boundary between groups with different MPCs.

Table B.7 decomposes each model's exact fixed-price response by beginning liquid wealth. The statistic is the stationary-distribution-weighted effective MPC,

$$\frac{c(x + \Delta) - c(x)}{\Delta},$$

conditional on being weakly below or strictly above zero. The discrete mass at zero wealth is included in the weakly non-positive group, and consumption after the transfer is evaluated on the household's beginning-of-period borrower or non-borrower policy branch. Thus, the entry-cost policy branches and the mass at zero are handled explicitly rather than inferred from CDF node masses.

Table B.7: Conditional MPCs on either side of zero liquid wealth

Specification	$\Pr(a \leq 0)$	$\text{MPC}(a \leq 0)$	$\text{MPC}(a > 0)$	Ratio
No debt (A)	0.139	0.469	0.231	2.03
Debt access (B1)	0.280	0.359	0.193	1.86
Debt + preference recalibration (B2)	0.261	0.307	0.169	1.82
Baseline debt model (C)	0.233	0.241	0.172	1.40

Notes: Exact partial-equilibrium decomposition at a transfer of 0.5 percent of annual GDP, with fixed prices and steady-state policies. Conditional effective MPCs are weighted by each model's stationary distribution. The $a \leq 0$ group includes the discrete mass at zero; the no-debt model has no negative-wealth region, so its first group consists only of that mass.

Across all four variants, households at or below zero have a substantially higher average

MPC than households with positive liquid wealth. The comparison is strongest in the no-debt benchmark, where the first group consists of the mass exactly at the borrowing limit, and remains quantitatively meaningful in every debt specification.

This does not require the MPC to jump downward exactly at zero. In the debt models, households just above zero can remain high-MPC; the lower conditional average on the positive side emerges as MPCs decline with the accumulation of liquid buffers. Zero liquid wealth therefore provides a useful common empirical boundary between, on average, higher- and lower-MPC regions of the wealth distribution.

B.5 Mapping the SCF to the Model

This section describes how we map households in the Survey of Consumer Finances (SCF) to the model. The mapping has two conceptually separate parts. First, we construct the empirical wealth and income inputs and use residualized-income rank to assign households to persistent model income states. Second, we put SCF wealth, transfers, and model consumption in common GDP units and average model policies over the model-implied distribution of discount-factor types within wealth–income rank cells.

The resulting paired SCF is used both to construct the zero-crossing curve in the main text and to provide an additional model-informed validation of the finite-transfer consumption response. We first describe the SCF data and income-state assignment, then construct the paired SCF, and finally report the paired SCF consumption response.

B.5.1 SCF Data and Income-State Assignment

We use the Survey of Consumer Finances summary extracts for the twelve triennial waves 1989–2022 (Aladangady et al., 2023), pooling the five implicates with survey weights and giving each wave equal weight. Liquid wealth follows Kaplan and Violante (2014a): transaction accounts and directly held liquid financial assets, net of revolving credit-card debt. This definition produces the negative liquid-wealth observations that distinguish the debt models from the no-debt one.

The mapping has two conceptually separate parts. First, a common-GDP conversion puts SCF wealth, transfers, and model consumption in the same units. Second, residualized-income rank assigns a persistent model state. The next subsection states the conversion explicitly when constructing the paired SCF response.

Within each wave, we run a weighted regression of log income (floored at $1,000) on a cubic in age, indicators for education, marital status, race, and sex of the household head, and the number of children. The rank of the residual is matched to the cumulative stationary probabilities of the model's eleven Rouwenhorst states. Cardinal household income neither rescales the policy function nor enters the zero threshold in the crossing exercise.

Income is gross pre-tax total household income. Because it enters only through its within-wave rank, the tax treatment has little effect on the assignment. Re-ranking households using TAXSIM disposable income, which is available for the 1989–2019 waves, leaves 86 percent in the same state and more than 99 percent within one state.

After this step, each SCF household i is characterized by its observed liquid wealth W_i, its survey weight, and its assigned persistent-income state z_i. Section B.5.2 adds the common-GDP conversion, model-specific policy evaluation, and posterior distribution of discount-factor types used to construct the paired SCF.

B.5.2 Constructing the paired SCF

Because consumption following a hypothetical transfer is not observed in the SCF, we evaluate each household using model-implied consumption policies. For each model calibration, the paired SCF preserves observed wealth levels and the observed wealth–income joint distribution, assigns the persistent-income state z_i using the residualized-income rank described in Section B.5.1, and averages model policies over the posterior distribution of discount-factor types in household i's wealth-rank and income-rank cell.

The construction is model specific: each model, indexed here by M, is paired with an SCF evaluation based on that model's policies, income process, and implied discount-type composition. We use the same paired-SCF construction for both empirical applications in the paper. It determines post-transfer liquid wealth and hence the zero-crossing curve in Figure 1, and it provides the paired SCF consumption-response benchmark reported in Section B.5.3.

Common-output conversion. All dollar quantities are put in common output units. The model has no capital, so its output concept is not measured GDP but output net of fixed capital formation. If annual GDP per household is Y^{ann} dollars and s_I is the share of output devoted to fixed capital formation, one model unit therefore equals $(1 - s_I)Y^{ann}/4$ dollars. We set $s_I = 0.20$, since gross private domestic investment averaged close to one-fifth of US GDP over 1989–2022. Writing $g_Y = (1 - s_I)Y^{ann}/4$, we obtain for SCF wealth W_i and a dollar transfer D the normalized values

$$a_i = \frac{W_i}{g_Y}, \qquad \Delta = \frac{D}{g_Y}. \tag{52}$$

Thus Δ is the dollar-to-model-unit conversion, while the transfer itself enters cash-on-hand one-for-one. The transfer is a share of actual GDP and the response is reported as a percent of actual annual GDP; only the mapping into model units carries $1 - s_I$. The same conversion is used for the crossing share of Figure 1 and for every consumption response reported below, so the paper carries a single dollar-to-model-unit conversion throughout.

Model-specific policy and type assignment. Each model is paired with an SCF estimate that preserves the observed wealth–income joint distribution and observed wealth levels, assigns z_i by residualized-income rank, and averages over the model-implied posterior

$$\pi_i^M(\beta) := \Pr_M(\beta \mid \text{wealth-rank cell of } i, \text{income-rank cell of } i).$$

The income rank is taken on residualized income. In particular, for the zero-crossing exercise, we map each household's residual-income rank to one realized model-income cell (z, k) and evaluate the policy at those selected persistent- and transitory-income states. The paired-SCF consumption response, instead, averages its policy response over the model's transitory-income distribution. The partition uses 40 equal-width wealth-rank bins and income-rank cells cut at the model's persistent-state masses, so every occupied cell has model support.

Cash-on-hand is then constructed as

$$x_{izk} = R(a_i)a_i + \iota_{zk},$$

where $R(a)$ is the model-implied gross quarterly return on observed beginning-of-period assets a_i, and ι_{zk} is model income for the data-implied ranks (z, k).

In debt models, negative SCF wealth uses the incumbent-borrower policy and return; formally, $s_i = B$ if $a_i < 0$ and $s_i = N$ otherwise. For nonborrowers in the debt models, c_M^N is the explicit waiver mixture $\lambda c_M^F + (1-\lambda)c_M^K$.

Zero-crossing construction. Household i is counted as crossing zero at transfer size Δ iff

$$a_i \leq 0 \wedge a_i + \Delta - \left[c_M^{s_i}(x_{izk} + \Delta; \beta, z_i, k_i) - c_M^{s_i}(x_{izk}; \beta, z_i, k_i)\right] > 0. \tag{53}$$

That is, a household crosses zero at transfer size Δ if they start at weakly negative wealth and if the transfer induces them to save a positive amount.

Applying the construction separately to the four model variants gives four model-paired SCF crossing curves. The black SCF line in Figure 1 is their pointwise mean, and the shaded band spans their pointwise minimum and maximum. The four curves differ by at most about 1.5 percentage points. The empirical crossing share is therefore governed primarily by the mass of the observed liquid-wealth distribution near zero; the precise model policy matters relatively little for this binary outcome. Aggregate consumption remains more model dependent because it depends on the amount consumed by every household rather than only on whether a household crosses zero. Section B.5.3 uses the same paired-SCF construction to evaluate this intensive-margin response.

B.5.3 Paired SCF Consumption Response

The paired-SCF construction is defined as

$$C_i(\Delta) = g_Y \sum_\beta \pi_i^M(\beta) \sum_k \omega_k \left[c_M^{s_i}(x_{izk} + \Delta; \beta, z_i, k) - c_M^{s_i}(x_{izk}; \beta, z_i, k)\right]. \tag{54}$$

It provides a model-informed benchmark for aggregate consumption. Because consumption following a hypothetical transfer is not observed in the SCF, this exercise is not an independently observed calibration target. Instead, it provides suggestive evidence on whether each calibration maps the observed SCF wealth–income distribution into a finite-transfer consumption response close to the corresponding model response.

Table B.8: Model response minus its paired SCF response

Model	$\Delta = 0.5\%$	$\Delta = 5\%$	$\Delta = 15\%$
A	0.0039	-0.0109	-0.0774
B1	0.0208	0.0727	0.0869
B2	0.0182	0.0568	0.0676
C	0.0090	0.0252	0.0351

Notes: Entries are the model response minus its paired SCF response, in percentage points of annual GDP. Each row uses that model's policies and discount-type posterior to construct its own paired empirical benchmark; the rows are not four estimates of one model-free empirical curve.

Table B.8 shows close agreement between the model and paired SCF responses as the paired gaps are small throughout. They increase modestly with transfer size for the debt specifications, while model A's gap turns slightly negative at larger transfers. Among the debt specifications, model C delivers the smallest gap at every reported transfer size. Each entry is a net aggregate gap, so discrepancies across different regions of the wealth distribution can offset, which is the case for model A where the missing borrower contribution offsets discrepancies on the nonnegative wealth support. Figure B.2 shows that model A overstates income–wealth sorting more strongly than the debt models.

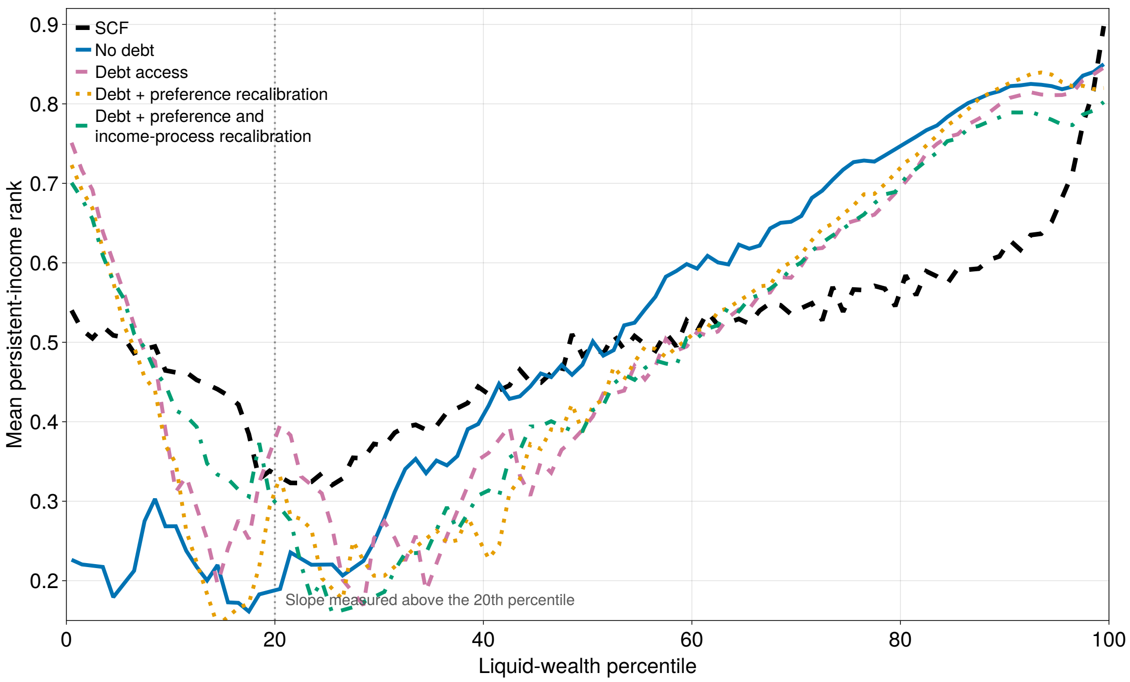

Figure B.2: **Mean persistent-income rank by liquid-wealth percentile.**

Notes: At each liquid-wealth percentile, the figure reports the mean rank of the persistent-income state. Both the SCF and model curves use the same persistent-state rank definition and population weighting; model curves use a five-bin moving average. The dotted vertical line marks the 20th liquid-wealth percentile. The slopes from the 20th to the 100th percentile are 0.113 in the SCF, 0.240 with no debt, 0.201 with debt access, 0.233 after preference recalibration, and 0.195 in the baseline model after preference and income-process recalibration.

Income–wealth sorting. Households' policies are determined not only by their liquid-wealth position but also by their persistent-income states. Hence, divergence in the *joint* income–wealth distribution can explain remaining gaps in Table B.8, even for model C, which matches the *marginal* empirical wealth distribution very well.

Figure B.2 shows the joint sorting of persistent income and liquid wealth observed in the SCF and implied by the different model calibrations. In the data, income follows a "hockey-stick" pattern over wealth: income falls with lower debt but rises with net wealth once wealth is positive (around the 20th percentile). All models with debt reproduce this pattern qualitatively. In the SCF, however, income is relatively flat over most of the wealth distribution, whereas all four models sort persistent income too strongly into wealth. The baseline debt model is closest to the data.

C Numerical Implementation

This appendix describes the numerical implementation of the model in Appendix A. The nonstandard part is the treatment of debt entry. The return kink and fixed entry cost create a discontinuous household policy, a point mass at zero, and a distinct mass strictly below zero. The implementation must therefore recover the endogenous entry threshold, track the two sides of the asset distribution separately, and preserve differentiability of the distribution transition.

We first describe the additional controls used to implement debt entry and the resulting distribution update. We then report the state-space solution and numerical settings and close with differentiability and validation checks.

C.1 Debt Entry and the Distribution Update

Boundary controls and the entry threshold. Appendix A.3 defines the debt-entry policy and the value-matching threshold $x^{\star}_{h,t}$. To impose value matching in the state-space system without storing the full value function, we recover continuation-value differences from the marginal policy.

Let $c^{\mathrm{EGM}}_{t,h}(a')$ denote the unrestricted (endogenous-grid) consumption schedule for household state h, evaluated at *end-of-period* assets a' before the household chooses among the debt-entry, zero, and saving branches. Let $\mathcal{V}_{t,h}(a') := \mathbb{E}_t V_{t+1,h'}(a')$ denote the expected continuation value conditional on choosing a'. We define the boundary premium

$$J_{t,h} := \beta_h \left[\mathcal{V}_{t,h}(0) - \mathcal{V}_{t,h}(0^-)\right]. \tag{55}$$

The premium measures the difference between entering the next period as a non-borrower and as an incumbent borrower at the boundary, i.e. the value of arriving at essentially zero assets with different *debt access status*. It generally holds that $J_{t,h} \leq 0$. Optimality in the interior (under CRRA preferences) requires that

$$\beta_h \mathcal{V}'_{t,h}(a') = c^{\mathrm{EGM}}_{t,h}(a')^{-\sigma}. \tag{56}$$

Discounted continuation-value differences can then be reconstructed from the marginal policy by the fundamental theorem of calculus (where the integral is oriented):

$$\mathcal{D}_{t,h}(\underline{\gamma}, \bar{\gamma}; c^{\mathrm{EGM}}_{t,h}, J_{t,h}) := \int_{\underline{\gamma}}^{\bar{\gamma}} c^{\mathrm{EGM}}_{t,h}(a')^{-\sigma}\, da' + \Big[\mathbf{1}\{\bar{\gamma} \geq 0\} - \mathbf{1}\{\underline{\gamma} \geq 0\}\Big] J_{t,h}. \tag{57}$$

The integral recovers continuation-value differences within either side of zero, while $J_{t,h}$ accounts for crossing the borrower-status boundary.

We define superscript B as the *incumbent*-borrower branch and superscript K as the non-waived non-borrower branch. The objects c^B, c^K and a^B, a^K denote the consumption and asset choices at the *common limiting cash-on-hand* defined by initial wealth approaching the status boundary from either side:

$$a \nearrow 0^- \quad \text{or} \quad a \searrow 0^+, \qquad x \to \iota_{hk}. \tag{58}$$

Taken together, we find that the boundary premium satisfies the following recursion:

$$J_{t,h} = \beta_h(1-\lambda)\sum_{h'}\Pi_{hh'}\sum_k \omega_k\Big[u\Big(c^K_{h'k,t+1}\Big) - u\Big(c^B_{h'k,t+1}\Big) \\ + \mathcal{D}_{t+1,h'}\left(a^B_{h'k,t+1}, a^K_{h'k,t+1}; c^{\text{EGM}}_{t+1,h'}, J_{t+1,h'}\right)\Big]. \tag{59}$$

The two households arrive at the limiting cash-on-hand associated with assets 0^- and 0, but their borrower and non-borrower statuses generate different conditional policies. The waived branch (with probability weight λ) contributes zero to the difference because for those households the optimization problem is identical to that of an incumbent borrower.

A second boundary control records the borrower-side marginal-utility endpoint

$$M^{0^-}_{t,h} := \lim_{a\nearrow 0} c^{\text{EGM}}_{t,h}(a)^{-\sigma}.$$

Using the disposable nonasset resources $\iota_{hk,t}$ defined in (40), the left-limit Euler equation determines $M^{0^-}_{h,t}$:

$$M^{0^-}_{h,t} = \beta_h(R_{t+1}+\rho_0)\sum_{h'}\Pi_{hh'}\sum_k \omega_k\left[c^B_{h',t+1}\left(\iota_{h'k,t+1}\right)^{-\sigma}\right]. \tag{60}$$

This is distinct from the marginal value at the non-borrower node $a=0$, which uses the saving return (no borrowing wedge) and the waiver mixture. While wealth levels infinitesimally close to 0 have measure zero, we must track this term due to the *finite grid* we employ in practice: we represent the marginal utilities to the left of zero as a cubic spline, with the end node given by $(0^-, M^{0^-}_{t,h})$.[19] This endpoint enters the system as an extra control for each state h.

Conditional on $(c_t, J_t, M^{0^-}_t)$, we can compute the first-order deviation of the debt-entry threshold x^*_h from its steady-state value in the following way. We define and rewrite the value-matching residual as

$$\mathcal{G}_{t,h}(x^*_{h,t}) := V^{K,E}_{t,h}\left(\frac{x^*_{h,t}-\iota_{hk}}{R}\right) - V^{K,0}_{t,h}\left(\frac{x^*_{h,t}-\iota_{hk}}{R}\right) \tag{61}$$

$$= u(x^*_{h,t}-\kappa-a^E_{t,h}) - u(x^*_{h,t}-a^0_h) + \beta_h\left(\mathcal{V}_{t,h}(a^E_{t,h}) - \mathcal{V}_{t,h}(a^0_h)\right) \tag{62}$$

$$= u(x^*_{h,t}-\kappa-a^E_{t,h}) - u(x^*_{h,t}) - \mathcal{D}_{t,h}(a^E_{t,h}, 0; c^{EGM}_{t,h}, J_{t,h}) \tag{63}$$

$$= u(x^*_{h,t}-\kappa-a^E_{t,h}) - u(x^*_{h,t}) - \int_{a^E_{t,h}}^0 c^{\text{EGM}}_{t,h}(a')^{-\sigma}da' - J_{t,h}, \tag{64}$$

where $a^E_{t,h} < 0$ and $a^0_h = 0$ are optimal savings choices at the threshold. The approach is to linearize the value-matching condition

$$\mathcal{G}_{t,h}\left(x^\star_{h,t}\right) \stackrel{!}{=} 0$$

around steady-state values $\bar{x}^*_h, \bar{J}_h, \bar{M}^{0^-}_h, \bar{c}^{\text{EGM}}_h$.

For the derivative calculation, represent the savings choices by cash-on-hand x and consump-

[19] The cubic spline allows for numerical integration that can be differentiated by automatic differentiation.

tion choices:

$$a^E_{t,h}(x) = x - \kappa - c^E_{t,h}(x), \qquad a^0_h(x) = 0, \quad c^0_h(x) = x.$$

Holding the policy schedule fixed,

$$\frac{\partial \mathcal{G}_{t,h}}{\partial J_{t,h}} = -1, \qquad \frac{\partial \mathcal{G}_{t,h}}{\partial M^{0-}_{t,h}} = -\int_{a^E_h(x^\star_h)}^{0} \frac{\partial c^{\mathrm{EGM}}_{t,h}(a')^{-\sigma}}{\partial M^{0-}_{t,h}}\, da'. \tag{65}$$

The derivative with respect to the threshold argument x^*_h is

$$\begin{aligned}\frac{\partial \mathcal{G}_{t,h}}{\partial x^*_{h,t}} &= u'\left(c^E_{t,h}\right) c^{E\,\prime}_{t,h}(x^*_{h,t}) - u'\left(x^*_{h,t}\right) + c^{\mathrm{EGM}}_{t,h}\left(a^E_{t,h}\right)^{-\sigma} (1 - c^{E\,\prime}_{t,h}(x^*_{h,t})) \\ &= u'(c^E_{t,h}) - u'(x^*_{h,t}),\end{aligned} \tag{66}$$

where the second line follows from $c^{\mathrm{EGM}}_{t,h}\left(a^E_{t,h}\right)^{-\sigma} = u'(c^E_{t,h})$.

Finally, let $\mathcal{P}_h$ denote the state of the policy schedule (the consumption interpolants, stemming from EGM, and their implied marginal-utility interpolants), which are the typical control variables of heterogeneous agent models. The derivative with respect to $\mathcal{P}_h$, holding the evaluation argument fixed, is

$$\frac{\partial \mathcal{G}_{t,h}}{\partial \mathcal{P}_{t,h}} = \frac{\partial c^E_{t,h}}{\partial \mathcal{P}_{t,h}} \underbrace{\left(u'(c^E_{t,h}) - c^{\mathrm{EGM}}_{t,h}(a^E_{t,h})^{-\sigma}\right)}_{=0} - \int_{a^E_{t,h}}^{0} \frac{\partial c^{\mathrm{EGM}}_{t,h}(a')^{-\sigma}}{\partial \mathcal{P}_{t,h}} da'. \tag{67}$$

The complete total differential of the threshold condition is then

$$0 = \left.\frac{\partial \mathcal{G}_{t,h}}{\partial x^*_{h,t}}\right|_{SS} dx^*_{h,t} - dJ_{t,h} + \left.\frac{\partial \mathcal{G}_{t,h}}{\partial M^{0-}_{t,h}}\right|_{SS} dM^{0-}_{t,h} + \left.\frac{\partial \mathcal{G}_{t,h}}{\partial \mathcal{P}_{t,h}}\right|_{SS} d\mathcal{P}_{t,h} \tag{68}$$

By the implicit-function theorem, the first-order deviation of the threshold $x^*_{t,h}$ around $\bar{x}^*_h$ is determined by rearranging Equation (68) for $dx^*_{h,t}$, which is possible as long as $a^E_h(\bar{x}^*_h) \neq -\kappa$. The deviation of $x^*_{t,h}$ is thus a static variable, determined by the forward-looking controls J and M^{0-} that are added in the state space for the models with debt-entry fixed costs.

Computing the steady-state debt-entry threshold. There generally is no closed-form expression for $x^\star_h$, because the policy and marginal-utility schedules are interpolated numerical objects. We rely on these state-space objects also when computing the threshold in steady state, so that the linearized system remains internally consistent.[20] We first test whether $\mathcal{G}_h(10^{-8}) \leq 0$, in which case we set $x^\star_h = 0$ (entry never dominates). Otherwise our algorithm brackets the unique root of $\mathcal{G}_h(x) = 0$ with initial upper endpoint $x_{h,\mathrm{enter}}(a^-_{\max})$, the shifted cash-on-hand value at the negative asset node closest to zero; it repeatedly expands that endpoint until $\mathcal{G}_h(x) < 0$, and searches for the root using bisection. This scalar solve is embedded in a joint household fixed point: Starting from a candidate policy and boundary objects, the algorithm updates the EGM consumption schedule, the boundary-value recursion for J_h, the borrower-side

[20] Alternatively, one could employ a separate value-function iteration to compute the value-matching residual that determines the steady-state thresholds. This will give the same solution in the limit of infinite grid sizes, but yields errors in the practical application since the numerical representations are not perfectly aligned.

endpoint Euler condition for $M_h^{0^-}$, and then solves $\mathcal{G}_h(x_h^\star; J_h, M_h^{0^-}) = 0$ for every h; it iterates these objects jointly to convergence.

Distribution at zero. For each combined discount-factor and persistent-income state h, let

$$F_{t,h}(a) = \Pr(a_t \le a,\ h_t = h)$$

denote the joint beginning-of-period asset CDF. The full CDF value at zero combines the mass below zero and the discrete mass at zero. We therefore define the borrower mass separately as

$$m_{t,h}^B = F_{t,h}(0^-), \qquad F_{t,h}(0) = m_{t,h}^B + \Pr(a_t = 0,\ h_t = h). \tag{69}$$

The numerical state contains the free CDF coordinates and one separate borrower mass $m_{t,h}^B$ for every h. The terminal CDF value fixes the total population mass p_h of each column, but no normalization fixes its borrower mass.

For readability, suppress time subscripts in the following definitions. CDF evaluation is separated by side. On the borrower side, a monotone interpolant $\mathcal{I}_h^B$ is constructed through the negative asset nodes and the endpoint $(0, m_h^B)$:

$$F_h^B(a) = \begin{cases} 0, & a < a_{\min}, \\ \mathcal{I}_h^B(a), & a_{\min} \le a < 0, \\ m_h^B, & a \ge 0. \end{cases} \tag{70}$$

On the non-borrower side, $\mathcal{I}_h^N$ is constructed from the full CDF on the nonnegative nodes, and the borrower mass is subtracted after evaluation:

$$F_h^N(a) = \begin{cases} 0, & a < 0, \\ \mathcal{I}_h^N(a) - m_h^B, & 0 \le a < a_{\max}, \\ p_h - m_h^B, & a \ge a_{\max}. \end{cases} \tag{71}$$

In particular,

$$F_h^N(0) = F_h(0) - m_h^B$$

returns the mass at zero exactly. Imposing the probability bounds through the endpoint data avoids additional minimum or maximum corrections after interpolation.

Paid entry flow. The mass of non-waived non-borrowers that enter debt is

$$e_t = (1-\lambda) \sum_h \sum_k \omega_k F_{t,h}^{N,<}\left(\frac{x_{h,t}^\star - \iota_{hk,t}}{R_t}\right), \qquad EC_t = \kappa e_t. \tag{72}$$

The argument of $F_{t,h}^N$ is the beginning-asset threshold at which current cash-on-hand equals the entry threshold.

Exact distribution transition. The transition evaluates the borrower and non-borrower parts of the beginning-of-period distribution separately. Let

$$R_{B,t} := R_t + \rho_0, \qquad R_{S,t} := R_t$$

denote the gross returns applied to negative and nonnegative beginning assets. For every savings choice a_j, let $x_{j,h}$ be the unrestricted endogenous-grid cash-on-hand preimage at which the household chooses a_j. The waived policy uses

$$x^F_{j,h} = x_{j,h}.$$

For the non-waived fraction, a negative savings choice uses the entry branch, with its cash-on-hand threshold capped at $x^\star_h$. A nonnegative savings choice uses the policy that combines remaining at zero and saving. The two CDF evaluations are averaged over the transitory shock and waiver draw and then mixed by the persistent-state transition matrix Π. More explicitly, before this transition, the savings CDF column at savings choice a_j is $\tilde{F}[j,h]$; primes denote next-period quantities:

$$\begin{aligned}\tilde{F}[j,h] = \sum_k \omega_k \Bigg[F^B_h\left(\frac{x_{j,h} - \iota_{hk}}{R_B}\right) + \lambda F^N_h\left(\frac{x^F_{j,h} - \iota_{hk}}{R_S}\right) \\ + (1-\lambda)\mathfrak{F}^K_{hjk}\Bigg], \qquad F' = \tilde{F}\Pi.\end{aligned} \tag{73}$$

where

$$\mathfrak{F}^K_{hjk} = \begin{cases} F^N_h\left(\dfrac{x_{j,h} + \kappa - \iota_{hk}}{R_S}\right), & a_j < 0,\ x_{j,h} + \kappa < x^\star_h, \\ F^{N,<}_h\left(\dfrac{x^\star_h - \iota_{hk}}{R_S}\right), & a_j < 0,\ x_{j,h} + \kappa \geq x^\star_h, \\ F^N_h\left(\dfrac{x_{j,h} - \iota_{hk}}{R_S}\right), & a_j \geq 0. \end{cases} \tag{74}$$

The second case in (74) is the entry cap: a non-waived household enters only *strictly* below the value-matching threshold, even if a deeper negative savings choice would otherwise imply a larger beginning-asset threshold. The left limit preserves this strict inequality.

The CDF transition at $a_j = 0$ does not separately identify next period's borrower mass because it also includes the discrete mass at zero. We therefore evaluate the borrower-mass law directly at the output threshold $a' = 0$:

$$\begin{aligned}\widetilde{m}^{B\prime}_h = \sum_k \omega_k \Bigg[F^{B,<}_h\left(\frac{x_{0,h} - \iota_{hk}}{R_B}\right) + \lambda F^{N,<}_h\left(\frac{x_{0,h} - \iota_{hk}}{R_S}\right) \\ + (1-\lambda)F^{N,<}_h\left(\frac{x^\star_h - \iota_{hk}}{R_S}\right)\Bigg], \qquad m^{B\prime} = \widetilde{m}^{B\prime}\Pi.\end{aligned} \tag{75}$$

The three terms correspond to incumbent borrowers below the limiting zero threshold, waived non-borrowers below the unrestricted threshold, and non-waived non-borrowers below the entry threshold.

Denoting the complete distribution update by $\mathcal{T}$, its derivative has the coupled form

$$D\mathcal{T} = \begin{bmatrix} \partial F'/\partial F & \partial F'/\partial m^B \\ \partial m^{B\prime}/\partial F & \partial m^{B\prime}/\partial m^B \end{bmatrix}. \tag{76}$$

The off-diagonal blocks are economically required: changing the borrower mass affects both side-specific probabilities, while changing the remaining CDF coordinates affects next period's borrower mass.

Negative assets are aggregated from the borrower-side interpolant:

$$A^- = -\sum_h \int_{a_{\min}}^0 F_h^B(a)\, da. \tag{77}$$

Aggregate total assets are recovered analogously from the full distribution. The same transition therefore determines the next distribution, aggregate saving, negative assets, and the paid-entry flow used in the aggregate equilibrium conditions.

C.2 State-Space Solution and Numerical Settings

State-space representation. The combined idiosyncratic state h has

$$N_H = 5 \times 11 = 55$$

values. Let dF_t stack deviations of the beginning-of-period joint CDF columns from their stationary values at every interior asset-grid node. The distribution is retained at full grid resolution; only the terminal CDF coordinate of each column is omitted, because it is fixed by the population mass p_h.

In the debt models, the distributional state also contains the deviations dm_t^B of the borrower masses defined in (69). The remaining aggregate states are the lagged real rate

$$r_t^{\text{lag}} = r_{t-1} - \bar{r}$$

and a one-time shock state ε_t that introduces the transfer innovation into the fiscal rule. Thus, suppressing transposes, the state vector is

$$s_t = \left(dF_t,\, dm_t^B,\, r_t^{\text{lag}},\, \varepsilon_t\right). \tag{78}$$

The controls contain a compressed representation dc_t of the consumption-policy deviations and the aggregate variables

$$\left(B_t,\, A_t,\, A_t^-,\, Y_t,\, C_t,\, \pi_t,\, i_t,\, r_t,\, T_t\right).$$

With a fixed debt-entry cost, the controls additionally contain the boundary premium J_t and the borrower-side marginal-policy endpoint $M_t^{0^-}$ for every household state h. The entry threshold $x_t^\star$ is recovered from these controls using (68) and is therefore not included as a separate control. In the no-debt model, dm_t^B, J_t, $M_t^{0^-}$, and the negative asset grid are omitted.

Stacking the household optimality conditions, the distribution transition in Section C.1, and the aggregate equilibrium conditions in Appendix A gives a system

$$\mathbb{E}_t \mathcal{F}(s_t, y_t, s_{t+1}, y_{t+1}) = 0. \tag{79}$$

Automatic differentiation provides the Jacobians of this system at the stationary equilibrium. Its first-order expansion is solved using a generalized Schur decomposition, yielding the decision rule and state transition

$$y_t = g_s s_t, \qquad s_{t+1} = h_s s_t. \tag{80}$$

Policy compression. The distribution is not dimensionally reduced, but the consumption-policy control is compressed using a discrete cosine transform. The vector dc_t contains the retained DCT coefficients of the policy deviations for all household states. We retain a small share of coefficients at the baseline tolerance. In the entry-cost model, the $2N_H$ controls J_t and M_t^{0-} are added separately because they determine the behavior of the policy at the debt-entry boundary.

Asset grids. The nonnegative asset grid contains

$$N_A^+ = 100$$

points on $[0, 4000]$, including zero. Its nodes are uniformly spaced under the transformation

$$q(a) = \log(1 + \log(1 + a)),$$

which concentrates points in the region where most households are located.

The debt models add

$$N_A^- = 20$$

strictly negative nodes on $[\underline{a}, 0)$, with $\underline{a} = -1$. Their absolute values are uniformly spaced under the same transformation and then reflected around zero. Thus,

$$N_A = N_A^+ + N_A^- = 120$$

in the debt models, and zero appears only once, as part of the nonnegative grid.

Interpolation and differentiation. The CDF interpolant is monotone cubic and uses a Fritsch–Carlson construction with a smooth adjustment where neighboring secant slopes change sign. In the debt models, the negative and nonnegative parts of the CDF are interpolated separately as described in Section C.1, and the borrower mass has its own law of motion. This construction preserves differentiability when automatic differentiation forms the system Jacobians.

NDU and the nonlinear benchmark. NDU uses the same first-order decision rule g_s and state transition h_s. Within its exact-update window, household-policy responses remain first-order, while the distribution transition and associated aggregate-savings condition are evaluated using the nonlinear operator $\mathcal{T}$ defined in Section C.1.

The fully nonlinear benchmark uses the same asset grid and DEGM distribution transition over a path of length $T = 150$. It is solved with a safeguarded Newton method initialized from the linearized transition path. At the finite grid used in the reported results, the remaining nonlinear-system residuals are approximately 10^{-6}.

C.3 Differentiability and Validation

Differentiability. The CDF interpolator is monotone cubic except in the terminal upper-tail cell, which is linear. The borrower-side interpolant remains cubic through its endpoint at zero. Constructing the endpoints separately for the borrower and non-borrower sides imposes the relevant probability bounds without non-differentiable minimum or maximum corrections. The linear terminal cell prevents the last free upper-tail CDF coordinate from generating unstable feedback. CDF integration for aggregate assets uses the same interpolation rule within each cell.

Validation. We check the implementation in four ways. First, automatic-differentiation Jacobian actions agree with two-sided finite differences of the nonlinear distribution transition. Second, each type-income column retains its calibrated terminal mass, total probability remains one, and the borrower- and non-borrower-side probability bounds hold without projection. Third, initialization and warm starts preserve the calibrated, potentially unequal discount-type masses, as required by the block-diagonal transition matrix. Fourth, the generalized eigensystem has the required stable-root count at the tested grid sizes, and the no-debt implementation reproduces the root count of the corresponding model without debt.

The same nonlinear distribution operator is used to construct the state-space Jacobians, during the exact-update window of NDU, and in the fully nonlinear benchmark. These checks therefore apply to the distribution transition used across all three solution methods.

D Policy Boundaries under Debt Entry Costs

This appendix provides the additional boundary analysis required by the entry-cost model. It first generalizes Proposition 1 to discontinuous policies. It then derives the policy-only approximations used in the quantitative decomposition and reports the corresponding partial-equilibrium evidence.

D.1 Crossing Terms for Discontinuous Policies

With a debt entry cost $\kappa > 0$, the entry threshold $x^\star_\theta$ solves $V_{\text{enter}}(x) = V_{\text{stay}}(x)$. This is indifference in *value*, not in consumption, so the consumption levels on the two branches need not agree. In fact, consumption must jump at the threshold. Without an increase in consumption, entering debt would give weakly lower current utility, a worse continuation asset position, and the additional cost κ, making indifference impossible. The entry-cost model therefore violates the continuity assumption in Proposition 1.

Let the household draw an i.i.d. fee waiver with probability λ before choosing saving, so that the type-θ policy is the mixture

$$c_\theta = \lambda\, c_\theta^w + (1-\lambda)\, c_\theta^n \tag{81}$$

of the fee-free policy c_θ^w and the costly policy c_θ^n. Both are piecewise smooth, but only c_θ^n has a cutoff at $x_\theta^\star$. The mixture therefore inherits $(1-\lambda)$ times the jump in the costly policy, so the waiver probability scales the boundary term. The function c_θ is the average policy for type θ, not the policy of a particular household.

Proposition 3 (Crossing terms for discontinuous policies)**.** *Let $C_\theta(\Delta)$ denote the conditional type-θ contribution to* (9)*, with c_θ the mixture* (81)*. For μ-almost every θ, suppose c_θ is piecewise C^2 with $K_\theta < \infty$ cutoffs $\bar{x}_{\theta,1} < \cdots < \bar{x}_{\theta,K_\theta}$, each fixed in terms of cash-on-hand. A lump-sum transfer then shifts every corresponding cutoff in the initial distribution one-for-one, $q_{\theta,k}(\Delta) = \bar{x}_{\theta,k} - \Delta$. Suppose c_θ is C^2 up to each cutoff from either side, with finite one-sided limits*

$$J_{\theta,k} := c_\theta(\bar{x}_{\theta,k}^+) - c_\theta(\bar{x}_{\theta,k}^-), \qquad [m]_{\theta,k} := m_\theta^+(\bar{x}_{\theta,k}^+) - m_\theta^-(\bar{x}_{\theta,k}^-).$$

Suppose the continuous part of F_θ has a differentiable density f_θ at every moved cutoff, that any mass points are isolated and none coincides with a cutoff, and that the set of θ for which some $q_{\theta,k}(\Delta)$ coincides with a mass point has μ-measure zero. Suppose finally that the one-sided fields $m_\theta^\pm$ and $m_{\theta,x}^\pm$ are dominated, so that each branch may be differentiated under the integral. Then, for regular θ,

$$C_\theta'(\Delta) = \int m_\theta(x+\Delta)\, dF_\theta(x) + \sum_{k=1}^{K_\theta} J_{\theta,k}\, f_\theta(q_{\theta,k}(\Delta))\,, \tag{82}$$

and

$$\begin{aligned} C_\theta''(\Delta) = &\int m_{\theta,x}(x+\Delta)\, dF_\theta(x) \\ &+ \sum_{k=1}^{K_\theta} \left[[m]_{\theta,k}\, f_\theta(q_{\theta,k}(\Delta)) - J_{\theta,k}\, f_\theta'(q_{\theta,k}(\Delta)) \right], \end{aligned} \tag{83}$$

both integrals taken branch-wise, that is, with the integrand evaluated on the branch containing $x+\Delta$. Aggregate derivatives follow by integrating over $\mu(d\theta)$.

Proof. Fix a regular θ and write $q_k = q_{\theta,k}(\Delta)$, $q_0 = -\infty$, $q_{K_\theta+1} = +\infty$. Split the type-θ contribution at the moving cutoffs,

$$C_\theta(\Delta) = \sum_{k=0}^{K_\theta} \int_{q_k}^{q_{k+1}} c_\theta^{(k)}(x+\Delta)\, dF_\theta(x), \tag{84}$$

with $c_\theta^{(k)}$ the branch on $(\bar{x}_{\theta,k}, \bar{x}_{\theta,k+1})$. No mass point lies exactly at a cutoff, so every mass point is in the interior of a branch, where the integrand is smooth in Δ.

Differentiate (84) by Leibniz' rule. Because every cutoff is fixed in terms of cash-on-hand, $q_k'(\Delta) = -1$, and the two moving-limit terms that meet at cutoff k combine into

$$\left[c_\theta^{(k)}(\bar{x}_{\theta,k}^+) - c_\theta^{(k-1)}(\bar{x}_{\theta,k}^-) \right] f_\theta(q_k) = J_{\theta,k}\, f_\theta(q_k),$$

while the interior terms give the branch-wise integral of $m_\theta(x+\Delta)$. This is (82). The boundary terms cancel only where consumption is continuous. Where it jumps, the mass swept through the moved cutoff carries the level gap with it, and the crossing flow appears already in the first derivative.

Differentiate once more. The branch-wise integral of m_θ yields the integral of $m_{\theta,x}$ plus, at each cutoff, the same combination applied to the slopes, $[m]_{\theta,k} f_\theta(q_k)$. Each first-derivative boundary term contributes $J_{\theta,k} f'_\theta(q_k)\, q'_k(\Delta) = -J_{\theta,k} f'_\theta(q_k)$, which is where differentiability of the density is used. Summing gives (83). Integrating over θ gives the aggregate derivatives; the mass-point exceptions do not contribute because the set of affected types has μ-measure zero. $\square$

Setting every $J_{\theta,k} = 0$ and $K_\theta = 1$ returns Proposition 1. In the model at hand the two kinds of cutoff separate cleanly. At the hard borrowing constraint and at the bunching edges consumption is continuous and only its slope jumps, so $J_{\theta,k} = 0$ and the crossing flow is the second-order term (13). At the entry threshold consumption itself jumps, and because the jump is downward the crossing flow subtracts from the aggregate response at first order.

Remark 2. *The order at which a policy boundary first matters is set by the smoothness of the policy there. If the policy is C^{r-1} but not C^r at a cutoff, the crossing term first appears in the aggregate derivative of order $r+1$. A kink is $r=1$ and shows up in the curvature of the aggregate response; a jump is $r=0$ and shows up in its slope. A fixed cost of entering debt therefore renders the threshold-crossing of households first-order for aggregate consumption.*

D.2 Policy-Only Approximations

Proposition 3 separates aggregate derivatives into within-branch policy terms and boundary terms. We now derive the within-branch terms. The policy-only first- and second-order approximations follow from implicit differentiation of the Euler equation $u'(c) = \beta R\,\mathbb{E}[u'(c')]$ on the slack branch. Here y denotes the current idiosyncratic state; primes on c', m', and m'_x denote next-period quantities, while primes on u denote derivatives:

$$m(x,y) = \frac{\Lambda}{u''(c)+\Lambda}, \qquad \Lambda \equiv \beta R^2\,\mathbb{E}\left[u''(c')\,m' \mid y\right], \tag{85}$$

$$m_x(x,y) = \frac{[R(1-m)]^2\Psi - u'''(c)\,m^2}{u''(c)+\Lambda}, \qquad \Psi \equiv \beta R\,\mathbb{E}\left[u'''(c')(m')^2 + u''(c')\,m'_x \mid y\right], \tag{86}$$

with $m=1$ and $m_x = 0$ on the constrained branch, where consumption is affine and these derivatives are exact.

The expectations $\mathbb{E}[m]$ and $\mathbb{E}[m_x]$, taken over the fixed steady-state distribution, are not by themselves the derivatives of aggregate consumption. The policy is differentiable at all but finitely many cash-on-hand values for each type, so m and m_x are the correct pointwise derivatives almost everywhere. At the cutoffs, however, differentiation and integration do not commute. The boundary sums in Proposition 3 provide the missing terms. Table D.1 combines those sums with $\mathbb{E}[m]$ and $\mathbb{E}[m_x]$.

D.3 Partial-Equilibrium Evidence

Table D.1 complements the partial-equilibrium figure of Section 6.2 with Taylor approximations of the aggregate response. Prices, continuation values, and household policies are held fixed at the steady state. FOPS uses $C'(0)\Delta$, while SOPS adds $\frac{1}{2}C''(0)\Delta^2$. Each coefficient combines the household-level moment of Appendix D.2 with the policy-boundary sums in Proposition 3.

Table D.1: Partial-equilibrium impact consumption response

Transfer size	Nonlinear	FOPS	SOPS	NDU
0.5% GDP	0.0953	0.1118	0.1095	0.1006
1.0% GDP	0.1808	0.2236	0.2143	0.1778
2.0% GDP	0.3429	0.4472	0.4100	0.3437
5.0% GDP	0.7601	1.1181	0.8856	0.7556
10.0% GDP	1.3437	2.2362	1.3063	1.3461
20.0% GDP	2.2780	4.4724	0.7529	2.2755

Notes: Responses in percent of annual GDP, for the baseline debt model on the same solve as Figure 2 ($N_A = 120$, $N_A^- = 20$, log-log grid, $\kappa = 0.0125$, $\lambda = 0.13$); the Nonlinear and NDU (Nonlinear DEGM Update) columns are exactly the blue and gold curves of that figure, read at these transfer sizes. FOPS and SOPS are the aggregate Taylor coefficients $C'(0)\Delta$ and $C'(0)\Delta + \frac{1}{2}C''(0)\Delta^2$ of Proposition 3. Each combines a household-level moment, obtained by implicit differentiation of the Euler equation (Appendix D.2) using the same steady-state DEGM policy and distribution as the other columns, with the policy-boundary sums over the genuine cutoffs of each type: the borrowing constraint, the bunching edges, and, since $\kappa > 0$, the debt-entry threshold. FOPS and the linear approximation plotted in Figure 2 are two computations of the same derivative, $C'(0)$, one analytic and one a finite-difference tangent of the nonlinear benchmark; they *differ slightly* because interpolation and differentiation do not commute on a discretized policy. NDU pushes the steady-state distribution through the steady-state policy using the shifted distribution $\tilde{F}_\Delta$.

At first order, the household-level average MPC is ($\mathbb{E}[m] = 0.2297$). The entry-threshold jump contributes ($\sum_k J_k f(\bar{x}_k) = -0.0061$), giving ($C'(0) = 0.2236$). The extensive margin therefore lowers the first-order aggregate response, but only by about three percent.

At second order, the boundary terms dominate. Household-level curvature contributes ($\mathbb{E}[m_x] = -0.1063$), while the boundary sums contribute (-0.3586), giving ($C''(0) = -0.4649$). Thus, more than two thirds of the aggregate second derivative comes from households moving across policy branches rather than from curvature within a branch. Most of the boundary contribution is the slope-jump term ($\sum_k [m]_k f(\bar{x}_k) = -0.2882$) at the borrowing constraint and the bunching edges. The density-slope term ($-\sum_k J_k f'(\bar{x}_k) = -0.0704$) is also material.

Because most of the aggregate concavity comes from boundary crossing, the second-order correction becomes more useful as the transfer grows. It removes 14 percent of the first-order error at a 0.5-percent transfer and 22 percent at a 1-percent transfer, compared with 65 percent at 5 percent and 96 percent at 10 percent. At a 20-percent transfer, however, the second-order approximation overshoots and predicts less consumption than the nonlinear solution, by 1.53 percentage points. NDU avoids this limitation because it evaluates the moved cutoff in (7) directly rather than expanding around zero. Table D.2 shows the same pattern for the model calibration without debt. Since the nonlinearity in that model is more extreme, the local Taylor expansion performs worse, while our method continues to track the nonlinear solution closely.

Table D.2: Partial-equilibrium impact consumption response: no-debt calibration

Transfer size	Nonlinear	FOPS	SOPS	NDU
0.5% GDP	0.1321	0.1592	0.1532	0.1379
1.0% GDP	0.2471	0.3185	0.2945	0.2453
2.0% GDP	0.4503	0.6369	0.5409	0.4440
5.0% GDP	0.9462	1.5923	0.9921	0.9502
10.0% GDP	1.5886	3.1846	0.7839	1.5828
20.0% GDP	2.5889	6.3692	-3.2336	2.5859

Notes: Responses in percent of annual GDP, for the no-debt calibration on the same ($N_A = 100$), log-log-grid steady-state solve as the no-debt curve in Figure 2. The Nonlinear and NDU columns are the direct pointwise benchmark and one-shot DEGM CDF update, respectively. FOPS and SOPS are $C'(0)\Delta$ and $C'(0)\Delta + \frac{1}{2}C''(0)\Delta^2$, computed from implicit policy derivatives and the no-debt policy-boundary terms on the same steady-state DEGM objects. Here $C'(0) = 0.3184$ and $C''(0) = -1.200356$; the nonlinear finite-difference tangent is 0.310201.

E The Sequence-Space Savings Correction

In sequence space (Auclert et al., 2021), the distributional correction enters as a one-period savings correction. This achieves an (imperfect) nonlinear correction for the impact response, but it does not carry the corrected distribution forward, and thus fails to track the nonlinear solution over the horizon (this limitation motivates the state-space design in the main text). As in the state-space correction (NDU), the correction never re-solves the household problem, and the only nonlinear object it evaluates is the exact DEGM distribution transition of the impact period.

Equilibrium system. The unknown is the output path $Y = (Y_1, \ldots, Y_T)$, with $T = 150$. Let $\bar{Y}$ denote the constant path at steady-state output and e_1 the first coordinate vector. Given Y and Δ, the macro block computes prices and fiscal variables in three steps: a backward wage-Phillips recursion $\pi_t = \kappa_w \widehat{g}_t + \bar{\beta}\pi_{t+1}$, with $\widehat{g}_t = v'(Y_t) - (1-\tau)u'(C_t^{pc})$ the wage-Phillips gap of Section 4, terminal condition $\pi_{T+1} = 0$, and the intermediation resource cost held at its steady-state level inside $C_t^{pc} = Y_t - G - \overline{IC} - \overline{EC}$; the Taylor rule and Fisher equation $i_t = \bar{r} + \phi_\pi \pi_t$, $1 + r_t = (1+i_t)/(1+\pi_{t+1})$; and the fiscal block, in which the period-1 lump sum is $\bar{T} + \Delta$, later lump sums follow the debt rule $\bar{T} - \phi_B \bar{r}(B_{t-1} - \bar{B})$, and B_t accumulates at the lagged real rate. The household block then maps (r, Y, T) into aggregate goods demand $C(Y, \Delta)$, which in the debt economy is household consumption plus the intermediation and entry resource costs, since that is the object entering goods market clearing. Market clearing is the residual $\mathcal{R}(Y, \Delta) = C(Y, \Delta) + G - Y$. With steady-state sequence-space Jacobians

$$J_Y = \left.\frac{\partial C(Y, 0)}{\partial Y}\right|_{Y=\bar{Y}}, \qquad J_\Delta = \left.\frac{\partial C(\bar{Y}, \Delta)}{\partial \Delta}\right|_{\Delta=0},$$

the linear solution solves $(J_Y - I)(Y - \bar{Y}) + J_\Delta \Delta = 0$. Both Jacobians are built once by finite differences at the steady-state path and are never refreshed: column t of $J_Y - I$ is $[\mathcal{R}(\bar{Y} + h_J e_t, 0) - \mathcal{R}(\bar{Y}, 0)]/h_J$ with $h_J = 10^{-5}$, and J_Δ is the one-sided quotient $[\mathcal{R}(\bar{Y}, \varepsilon) - \mathcal{R}(\bar{Y}, 0)]/\varepsilon$ at a transfer ε equal to 0.01 percent of annual GDP. The matrix $J_Y - I$ is factorized once and reused for every transfer size.

The nonlinear remainder. The nonlinear correction enters through first-period saving. Starting from the steady-state distribution and policy, we apply one DEGM forward step with cash-on-hand shifted by income Y_1 and the transfer Δ. Concretely, let $\iota(Y_1, \bar{T} + \Delta)$ be the period-1 net-income matrix over persistent and transitory income states, let (F_{ss}, m_{ss}^B) be the per-type stationary DEGM distribution and its borrower mass, and let $\bar{x}$ and $\bar{x}^\star$ be the steady-state endogenous-grid cash-on-hand nodes and entry thresholds. End-of-period saving is

$$\tilde{A}_1(Y_1, \Delta) = \sum_i p_i^\beta \, A\Big[\mathcal{T}\Big(F_{ss,i}, m_{ss,i}^B;\, \bar{x}_i, \bar{x}_i^\star, \iota(Y_1, \bar{T} + \Delta), \bar{R}\Big)\Big], \tag{87}$$

where i indexes discount-factor types, p_i^β is the calibrated type mass, $\mathcal{T}$ is the same nonlinear side-aware DEGM operator used by NDU, and $A[\cdot]$ is the asset aggregate of a CDF and its borrower mass. Two features of (87) define the exercise. First, the policy arguments are the steady-state ones—$\bar{x}$ rather than $\bar{x} + dc_1$, and $\bar{x}^\star$ rather than the value-matched threshold at the shifted controls—so no policy or threshold response enters. Second, the return is the predetermined steady-state gross return $\bar{R}$, and only current resources move with (Y_1, Δ).

Let $\tilde{A}_1^{lin}$ be the first-order expansion of (87) at $(\bar{Y}, 0)$, constructed from the same finite-difference steps as the Jacobians: a central difference in Y_1 with step h_J and a one-sided difference in Δ with step ε,

$$\tilde{A}_1^{lin}(Y_1, \Delta) = \tilde{A}_1(\bar{Y}, 0) + \frac{\partial \tilde{A}_1}{\partial Y_1}(Y_1 - \bar{Y}) + \frac{\partial \tilde{A}_1}{\partial \Delta}\Delta,$$

and define

$$\gamma_A(Y_1, \Delta) = \tilde{A}_1(Y_1, \Delta) - \tilde{A}_1^{lin}(Y_1, \Delta) \tag{88}$$

the saving gap. It is the entire nonlinear content of the method: a single scalar per candidate (Y_1, Δ). It enters consumption with the opposite sign, so the corrected residual is

$$(J_Y - I)(Y - \bar{Y}) + J_\Delta \Delta - e_1 \gamma_A(Y_1, \Delta) = 0. \tag{89}$$

Because the correction γ_A depends on the output path only through Y_1, writing $v(\Delta) = \bar{Y} - (J_Y - I)^{-1} J_\Delta \Delta$ and $w = (J_Y - I)^{-1} e_1$ collapses the system to the scalar fixed point

$$Y_1 = v_1(\Delta) + w_1 \gamma_A(Y_1, \Delta), \tag{90}$$

which is solved by safeguarded Newton/bisection without re-solving the household problem. Because the tangent $\tilde{A}_1^{lin}$ in (88) is built with the same finite-difference steps as J_Δ, any small-step noise from the atom-crossing kinks in the DEGM response is common to γ_A and to the linear part of (89) and cancels in the corrected impact row.

Algorithm. The complete procedure for one transfer size Δ is:

1. Solve the steady state; store the per-type stationary pair $(F_{ss,i}, m^B_{ss,i})$ obtained by iterating $\mathcal{T}$ at steady-state income and policy to a fixed point (tolerance 10^{-10} on the CDF and on m^B), together with $\bar{x}_i$, $\bar{x}^\star_i$, and p^β_i.

2. Build and factorize $J_Y - I$ column by column, and form J_Δ, as described above. Compute $w_1 = \left[(J_Y - I)^{-1}e_1\right]_1$.

3. Evaluate the three tangent quantities of $\tilde{A}_1$: the level $\tilde{A}_1(\bar{Y}, 0)$, the central difference in Y_1, and the one-sided difference in Δ. These are four evaluations of (87) in total and are reused for every transfer size.

4. Compute the linear path $v(\Delta) = \bar{Y} - (J_Y - I)^{-1}J_\Delta\Delta$ and take $v_1(\Delta)$ as the starting value.

5. Solve (90) for Y_1. Write $\phi(Y_1) = Y_1 - v_1(\Delta) - w_1\gamma_A(Y_1, \Delta)$. Bracket the root on $[\bar{Y} - 0.5, \bar{Y} + 0.5]$ and verify a sign change. Then iterate: evaluate ϕ, update the bracket by the sign of ϕ, take the Newton step using a central difference of ϕ with step $100h_J$, accept it if it stays inside the bracket and bisect otherwise, and stop when the step falls below 10^{-9}. Each ϕ evaluation costs one evaluation of (87), i.e. one DEGM transition per type; no household problem and no Jacobian is touched.

6. With the converged γ_A, solve the full path once, $Y = \bar{Y} - (J_Y - I)^{-1}\left[J_\Delta\Delta - e_1\gamma_A\right]$, and pass Y through the macro block to recover inflation, the nominal and real rates, transfers, and debt.

Steps 1–3 are shared across transfer sizes, so the marginal cost of an additional size is the handful of DEGM transitions consumed by step 5 plus one back-substitution.

Results. Figures E.1 and E.2 compare this correction with NDU(6) and the nonlinear benchmark at transfers of 1, 5, 10, and 15 percent of annual GDP. Figure E.1 removes the entry cost, while Figure E.2 uses the baseline debt-entry cost. This comparison isolates whether a one-period correction is enough to recover the transition path once the impact nonlinearity has been (partly) corrected.

Without the entry cost, Figure E.1 shows that the sequence-space correction matches impact but overshoots the nonlinear response after the first period at every transfer size. With the baseline entry cost, Figure E.2 shows that it remains close at transfers of 1 and 5 percent of annual GDP but departs visibly at 10 and 15 percent, also on impact (the sequence-space correction deviates by 1.2pp for the impact-response to the 15 percent transfer). NDU(6) remains close to the nonlinear benchmark in both figures. The reason is direct. Sequence space substitutes the distribution out of the equilibrium system. The one-period saving correction therefore captures the nonlinear first-period update evaluated at the steady-state distribution, but the shifted distribution is not available as an input in later periods.

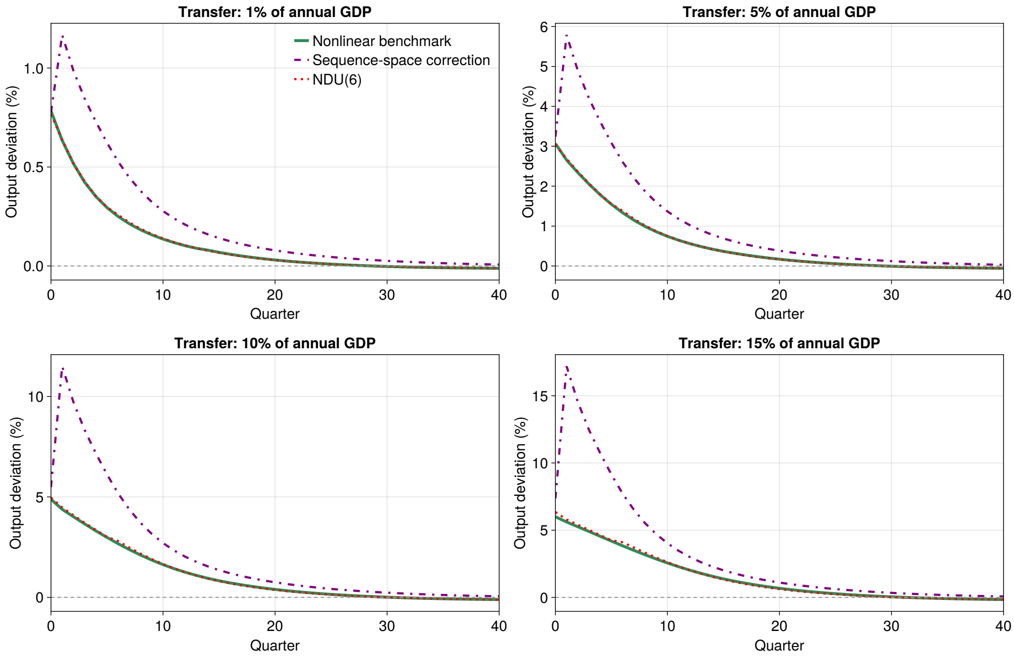

Figure E.1: **Output responses without a fixed debt-entry cost ($\kappa = 0$).**

Notes: Panels, read from upper left to lower right, show one-time transfers equal to 1, 5, 10, and 15 percent of annual GDP. Green is the nonlinear benchmark, purple dashed is the sequence-space correction in (90), and red dotted is NDU(6). The model calibration is that from Appendix B.3. Output is measured as a percentage deviation from steady state.

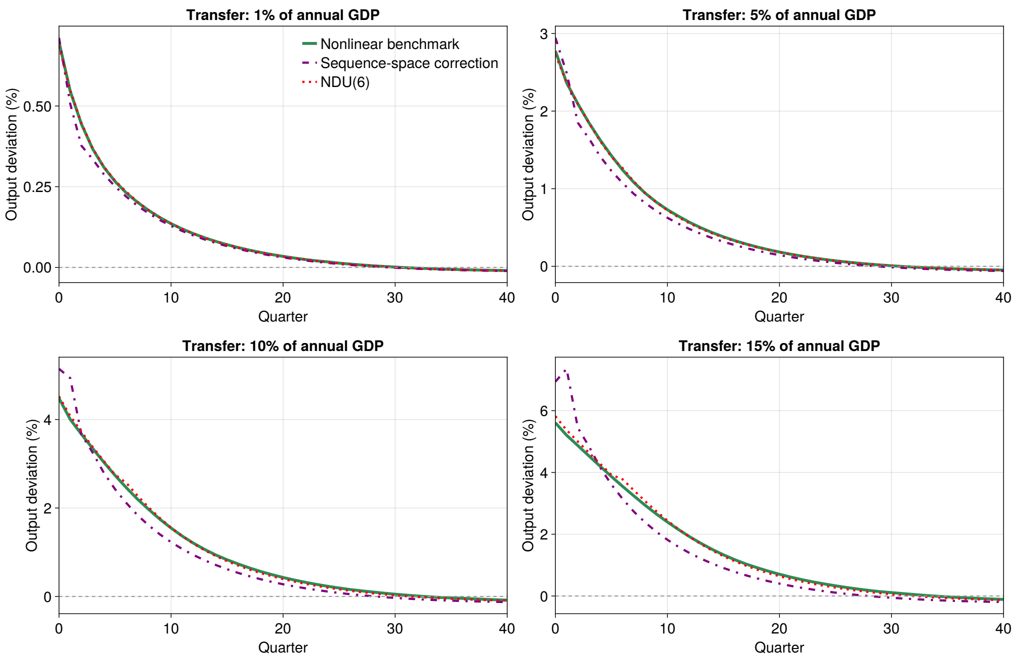

Figure E.2: **Output responses with a fixed debt-entry cost ($\kappa > 0$).**

Notes: Panels, read from upper left to lower right, show one-time transfers equal to 1, 5, 10, and 15 percent of annual GDP. Green is the nonlinear benchmark, purple dashed is the sequence-space correction in (90), and red dotted is NDU(6). The calibration is the baseline debt calibration. Output is measured as a percentage deviation from steady state.

References

Agarwal, S., Liu, C., and Souleles, N. S. (2007). The Reaction of Consumer Spending and Debt to Tax Rebates—Evidence from Consumer Credit Data. *Journal of Political Economy*, 115(6):986–1019.

Aladangady, A. et al. (2023). Changes in U.S. Family Finances from 2019 to 2022: Evidence from the Survey of Consumer Finances. Technical report, Board of Governors of the Federal Reserve System.

Andreasen, M. M. and Kronborg, A. F. (2022). The Extended Perturbation Method: With Applications to the New Keynesian Model and the Zero Lower Bound. *Quantitative Economics*, 13(3):1171–1202.

Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2021). Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica*, 89(5):2375–2408.

Auclert, A., Rigato, R., Rognlie, M., and Straub, L. (2026). Higher-Order Perturbation in Sequence Space: The Certainty Correspondence. Working Paper.

Auclert, A., Rognlie, M., and Straub, L. (2025). Fiscal and Monetary Policy with Heterogeneous Agents. *Annual Review of Economics*, 17:539–562.

Bayer, C. and Luetticke, R. (2020). Solving Discrete Time Heterogeneous Agent Models with Aggregate Risk and Many Idiosyncratic States by Perturbation. *Quantitative Economics*, 11(4):1253–1288.

Bayer, C., Luetticke, R., Pham-Dao, L., and Tjaden, V. (2019). Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk. *Econometrica*, 87(1):255–290.

Bayer, C., Luetticke, R., Weiss, M., and Winkelmann, Y. (2026). An Endogenous Gridpoint Method for Distributional Dynamics. *Journal of Monetary Economics*, 158:103895.

Bianchi, J. and Kaplan, G. (2026). How Small Is Small? Non-Linearities in Heterogeneous Agent Models. NBER Working Paper 35311, National Bureau of Economic Research.

Boutros, M. and Mijakovic, A. (2024). The Macroeconomic Implications of Coholding. Staff Working Paper 2024-16, Bank of Canada.

Carroll, C. D. (2006). The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems. *Economics Letters*, 91(3):312–320.

Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., and van Rooij, M. (2019). Asymmetric Consumption Effects of Transitory Income Shocks. *The Economic Journal*, 129(622):2322–2341.

den Haan, W. J., Kobielarz, M. L., and Rendahl, P. (2016). Exact Present Solution with Consistent Future Approximation: A Gridless Algorithm to Solve Stochastic Dynamic Models. Working Paper.

Evans, R. W. and Phillips, K. L. (2017). Solving and Simulating Unbalanced Growth Models Using Linearization about the Current State. *Economics Letters*, 151:35–38.

Fagereng, A., Holm, M. B., and Natvik, G. J. (2021). MPC Heterogeneity and Household Balance Sheets. *American Economic Journal: Macroeconomics*, 13(4):1–54.

Fulford, S. L. (2015). How Important Is Variability in Consumer Credit Limits? *Journal of Monetary Economics*, 72:42–63.

Fuster, A., Kaplan, G., and Zafar, B. (2021). What Would You Do with $500? Spending Responses to Gains, Losses, News, and Loans. *The Review of Economic Studies*, 88(4):1760–1795.

Guerreiro, J., Hazell, J., Känzig, D. R., and Manuel, E. (2026). The Macroeconomic Effect of Stimulus Checks: Evidence from Postwar Veterans' Payments. Working Paper.

Herkenhoff, K. F. (2019). The Impact of Consumer Credit Access on Unemployment. *The Review of Economic Studies*, 86(6):2605–2642.

Jappelli, T., Savoia, E., and Sciacchetano, A. (2026). Intertemporal MPC and Shock Size. *European Economic Review*, 186:105303.

Kaplan, G. and Violante, G. L. (2014a). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica*, 82(4):1199–1239.

Kaplan, G. and Violante, G. L. (2014b). A Tale of Two Stimulus Payments: 2001 versus 2008. *American Economic Review*, 104(5):116–121.

Kaplan, G. and Violante, G. L. (2022). The Marginal Propensity to Consume in Heterogeneous Agent Models. *Annual Review of Economics*, 14:747–775.

Koşar, G., Melcangi, D., Pilossoph, L., and Wiczer, D. G. (2025). Stimulus through Insurance: The Marginal Propensity to Repay Debt. NBER Working Paper 34399, National Bureau of Economic Research.

Mennuni, A., Rubio-Ramírez, J. F., and Stepanchuk, S. (2025). Dynamic Perturbation. *The Review of Economic Studies*, 92(2):1157–1192.

Reiter, M. (2009). Solving Heterogeneous-Agent Models by Projection and Perturbation. *Journal of Economic Dynamics and Control*, 33(3):649–665.

Young, E. R. (2010). Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell-Smith Algorithm and Non-Stochastic Simulations. *Journal of Economic Dynamics and Control*, 34(1):36–41.