ECONtribute
MARKETS & PUBLIC POLICY

ECONtribute
Discussion Paper No. 450

Local Methods for Large Transfers

Christian Bayer
Luigi-Maria Briglia
Ralph Luetticke
Maximilian Weiss
Yannik Winkelmann

November 2026

www.econtribute.de

UNIVERSITÄT BONN
UNIVERSITÄT ZU KÖLN

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/2 – 390838866 – is gratefully acknowledged.

Local Methods for Large Transfers*

Christian Bayer† Luigi-Maria Briglia‡

Ralph Luetticke§ Maximilian Weiss¶

Yannik Winkelmann‖

November 11, 2026

Abstract

Large fiscal transfers can move households across high- and low-MPC regions of the asset distribution, making aggregate responses nonlinear. We use the distributional endogenous-grid method to capture this adjustment and develop the Nonlinear DEGM Update (NDU), which applies this update over a short window while retaining first-order household policies. We discipline this margin with a zero-crossing curve that combines the empirical liquid-wealth distribution with model-implied consumption policies. In the data, the distribution is asymmetric around zero, a pattern standard models miss. A debt-access cost reproduces this pattern. For a transfer of 10 percent of annual GDP, output rises by 4.9 percent. A linear solution predicts 12.3 percent, whereas NDU matches the nonlinear solution at near-linear computational cost while retaining the distribution as a state.

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–EXC2126/2–390838866. Luigi-Maria Briglia, Ralph Luetticke, Maximilian Weiss, and Yannik Winkelmann gratefully acknowledge funding by the European Union (ERC, AIR-MAC, 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 constrained households, where MPCs fall sharply. We make three contributions. First, we provide an efficient method for handling this nonlinearity in a macroeconomic model. 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.

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 the resulting method the Nonlinear DEGM Update (NDU). It is nearly identical to the fully nonlinear solution and almost as fast as the linear solution.

Keeping the distribution as an explicit state is key. In sequence space, a one-period savings correction fixes the impact response, but by construction it does not track the later dynamics of the corrected distribution; Appendix D 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 4.85 percent under the fully nonlinear solution, 12.25 percent under the linearized solution, and 4.92 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, 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) 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 stochastic fixed cost of obtaining access to a nonzero debt limit delivers the empirically observed asymmetry around zero liquid wealth.

Matching the crossing curve matters for the aggregate nonlinearity of transfers. Because a large and dispersed share of households has negative liquid wealth in the data, the gap between nonlinear and linearized solutions remains sizable. It is nevertheless 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 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 exactly captures this 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. This literature establishes that both shock size and the timing of household spending matter.[1] 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.

[1] 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).

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.

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 policy approximations along the equilibrium path. Andreasen and Kronborg (2022) instead combine a nonlinear extended-path calculation with perturbation corrections.[2] NDU takes a more selective 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 KFE 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 documents the calibration and empirical mapping. Appendix B extends the policy-boundary analysis to debt-entry costs, Appendix C gives the complete state-space implementation, and Appendix D reports the sequence-space alternative.

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 the nominal return on the asset. Appendix A 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 $F^a_\theta(a) = \Pr(A \leq a \mid \theta)$ denote the share of households in persistent state $\theta \in \Theta$ whose assets are weakly below a, and let $\mu(d\theta)$ be the population measure over these states. We use the left-limit notation $F^a_\theta(a^-) = \Pr(A < a \mid \theta)$ whenever the underlying choice inequality is strict.

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

Cash-on-hand is obtained from assets through the monotone map

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

where $\iota_{\theta k}$ is current income conditional on transitory-income draw k, which occurs with probability w_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

$$F_\theta^x(q) = \sum_k w_k F_\theta^a\left(\phi_{\theta k}^{-1}(q)\right), \qquad \phi_{\theta k}^{-1}(q) = \begin{cases} \dfrac{q - \iota_{\theta k}}{R}, & q - \iota_{\theta k} \geq 0, \\ \dfrac{q - \iota_{\theta k}}{R + \rho_0}, & q - \iota_{\theta k} < 0. \end{cases} \tag{2}$$

For a common gross return R and abstracting from transitory-income draws, Equation (2) reduces to $F_\theta^x(q) = F_\theta^a((q - \iota_\theta)/R)$. For readability, we write $F_\theta := F_\theta^x$ and use f_θ for its density throughout this section. The state-space representation later returns to the asset CDF F_θ^a, which is the distributional state of the model.[3]

The notation in this and the next section treats θ as the index of the current policy that a household faces. In the debt-entry model, this index is augmented by beginning-of-period borrower status and, for non-borrowers, the waiver draw. The measure μ then includes their probabilities. A cash transfer shifts x while holding beginning status fixed, and the formulas are applied to each policy branch before averaging over the waiver. This is only a notational augmentation: borrower status remains encoded by the sign of beginning assets, and the waiver is iid rather than an additional persistent state. The continuous-policy argument below uses $c = x - g$. For a non-waived entrant the corresponding identity is $c = x - \kappa - g$, so the aggregate formulas subtract the entry resource cost EC. Appendix C writes these branches separately.

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

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

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 $c_\theta^j(x) = x - g_\theta^j(x)$, and the MPC on branch $j \in \{-, +\}$ is

$$m_\theta^j(x) = 1 - g_{\theta,x}^j(x). \tag{4}$$

[3]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.

The economically relevant case has $m_\theta^-(\bar{x}_\theta^-) > m_\theta^+(\bar{x}_\theta^+)$: 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

$$x_\Delta = x + \Delta. \tag{5}$$

The distribution after the policy choice is the distribution of end-of-period assets,

$$\widetilde{F}^a_{\theta,\Delta}(a') = \Pr\left(g_\theta(x + \Delta) \leq a'\right). \tag{6}$$

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

$$\widetilde{F}^a_{\theta,\Delta}(a') = F_\theta\left((g_\theta^j)^{-1}(a') - \Delta\right), \qquad j = \begin{cases} - & a' \leq \bar{a}_\theta, \\ + & a' > \bar{a}_\theta. \end{cases} \tag{7}$$

At the policy boundary itself,

$$\widetilde{F}^a_{\theta,\Delta}(\bar{a}_\theta) = F_\theta\left(\bar{x}_\theta - \Delta\right). \tag{8}$$

This formula also covers a hard borrowing limit. When $g_\theta^-(x) = \underline{a}$, the lower branch is flat rather than invertible, and (8) becomes the mass at the constraint:

$$\boxed{\widetilde{F}^a_{\theta,\Delta}(\underline{a}) = F_\theta\left(\bar{x}_\theta - \Delta\right).} \tag{9}$$

Equations (7)–(9) 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

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

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 branches. Aggregate first-period consumption is[4]

$$C(\Delta) = \int_{\Theta} \int c_\theta(x + \Delta)\, dF_\theta(x)\, \mu(d\theta). \tag{11}$$

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.*[5]

Proposition 1 (Policy-boundary crossing terms)**.** *Let $C_\theta(\Delta)$ denote the conditional type-θ contribution to* (11)*. 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,*

$$C'_\theta(\Delta) = \left[\int_{-\infty}^{\bar{x}_\theta - \Delta} m_\theta^-(x + \Delta) dF_\theta(x) + \int_{\bar{x}_\theta - \Delta}^{\infty} m_\theta^+(x + \Delta) dF_\theta(x) \right], \tag{12}$$

and

$$\begin{aligned} C''_\theta(\Delta) = & \left[m_\theta^+(\bar{x}_\theta^+) - m_\theta^-(\bar{x}_\theta^-) \right] f_\theta(\bar{x}_\theta - \Delta) \\ & + \int_{-\infty}^{\bar{x}_\theta - \Delta} m_{\theta,x}^-(x + \Delta) dF_\theta(x) \\ & + \int_{\bar{x}_\theta - \Delta}^{\infty} m_{\theta,x}^+(x + \Delta) dF_\theta(x). \end{aligned} \tag{13}$$

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:

$$C_\theta(\Delta) = \int_{-\infty}^{q_\theta(\Delta)} c_\theta^-(x + \Delta)\, dF_\theta(x) + \int_{q_\theta(\Delta)}^{\infty} c_\theta^+(x + \Delta)\, dF_\theta(x). \tag{14}$$

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

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

$$C(\Delta) = \mathcal{X}(\Delta) - \int_{\Theta} \int a'\, d\widetilde{F}^a_{\theta,\Delta}(a')\, \mu(d\theta) - EC(\Delta), \tag{10}$$

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.

[5] 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 B.1 states and proves the version for discontinuous policies in our stochastic entry-cost model.

MPC on each side, which gives (12). Differentiating once more with Leibniz' rule gives (13): 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. □

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

$$\left(m_\theta^+(\bar{x}_\theta^+) - m_\theta^-(\bar{x}_\theta^-)\right) f_\theta(\bar{x}_\theta - \Delta). \tag{15}$$

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 (15) 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

$$g_\theta(x) = \begin{cases} \bar{a}_\theta + s_\theta^-(x - \bar{x}_\theta), & x \le \bar{x}_\theta, \\ \bar{a}_\theta + s_\theta^+(x - \bar{x}_\theta), & x > \bar{x}_\theta, \end{cases} \qquad 0 \le s_\theta^- < s_\theta^+ \le 1. \tag{16}$$

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

$$C_\theta(\Delta) - C_\theta(0) = m_\theta^+ \Delta + (m_\theta^- - m_\theta^+) \int_0^\Delta F_\theta(\bar{x}_\theta - u)\, du. \tag{17}$$

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

$$\begin{aligned} C_\theta(\Delta) - C_\theta(0) = &\left[m_\theta^+ + (m_\theta^- - m_\theta^+) F_\theta(\bar{x}_\theta)\right] \Delta \\ &- \frac{1}{2}(m_\theta^- - m_\theta^+) f_\theta(\bar{x}_\theta)\Delta^2 + \frac{1}{6}(m_\theta^- - m_\theta^+) f_\theta'(\bar{x}_\theta)\Delta^3 + O(\Delta^4). \end{aligned} \tag{18}$$

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 $s_\theta^- = 0$, or $m_\theta^- = 1$. With heterogeneous types, the aggregate expansion integrates (18) 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. The debt entry margin creates a discrete mass at zero and requires us to track the distribution strictly below zero separately. We keep the notation here at the aggregate-block level; Appendix C states the household problem, policy branches, distribution transition, resource costs, and equilibrium conditions needed to reconstruct the implementation. Appendix A gives the calibration. In the state-space system, the linear laws of motion for these quantities can be replaced by exact updates in quarters when a large transfer moves the distribution substantially.

States and controls. Let $h = (\text{discount-factor type}, z)$ 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 x_t stacks

- the distribution coordinates X_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 $m^B_{t,h} = F_{t,h}(0^-)$.[6]
- 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 $(B_t, A_t, A^-_t, Y_t, C_t, EC_t, \pi_t, i_t, r_t, T_t)$. 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, EC_t is the debt-entry resource cost, 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

[6] 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.

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 $a' = x - c(x)$.

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 w_k is the probability of draw k. Aggregate labor income is $(1-\tau)\omega Y$, with $\omega = \sum_h p_h \sum_k w_k w_{h,k} = 1$ 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,

$$a' \le a_j \quad \Longleftrightarrow \quad a \le a^{\text{pre}}_{j,h,k}, \qquad a^{\text{pre}}_{j,h,k} = \begin{cases} \dfrac{x_{j,h} - \iota_{h,k}}{R_t}, & x_{j,h} - \iota_{h,k} \ge 0, \\ \dfrac{x_{j,h} - \iota_{h,k}}{R_t + \rho_0}, & x_{j,h} - \iota_{h,k} < 0. \end{cases} \tag{19}$$

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 $a^{\text{pre}}_{j,h,k}$,

$$\tilde{F}_{j,h} = \sum_k w_k \, \hat{F}_{t,h}\left(a^{\text{pre}}_{j,h,k}\right), \qquad F_{t+1,j,h'} = \sum_h \Pi_{hh'} \, \tilde{F}_{j,h} \;\equiv\; \mathcal{T}(F_t;\, x,\, \iota,\, R_t)_{j,h'}\,, \tag{20}$$

where $\hat{F}_{t,h}$ is a monotone-cubic interpolant of column h and w_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 operator 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 $\{x_t, y_t\}$ satisfies

$$\mathbb{E}_t \, \mathcal{F}(x_t,\, y_t,\, x_{t+1},\, y_{t+1}) = 0, \tag{21}$$

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

$$\mathcal{F} = \begin{pmatrix} (F_{t+1}, m^B_{t+1}) - \mathcal{T}\Big(F_t, m^B_t;\, \bar{x} + dc_t, J_t, M^{0^-}_t, Y_t, T_t, R_t\Big) & \text{(distribution: KFE rows)} \\ A_t - A[F_t, m^B_t], \qquad A^-_t - A^-[F_t, m^B_t] & \text{(distribution aggregates)} \\ dc_t - \mathcal{E}(dc_{t+1},\, Y_{t+1},\, T_{t+1},\, r_t) & \text{(household policy: Euler rows)} \\ J_t - \mathcal{J}\Big(dc_{t+1}, J_{t+1}, M^{0^-}_{t+1}, Y_{t+1}, T_{t+1}, r_t\Big) & \text{(boundary-value rows)} \\ M^{0^-}_t - \mathcal{M}\Big(dc_{t+1}, J_{t+1}, M^{0^-}_{t+1}, Y_{t+1}, T_{t+1}, r_t\Big) & \text{(left-limit Euler rows)} \\ EC_t - \kappa\, e\Big(F_t, m^B_t; dc_t, J_t, M^{0^-}_t, Y_t, T_t, R_t\Big) & \text{(entry-resource row)} \\ C_t - \Big[R_t A_t + \rho_0 A^-_t + (1-\tau)\omega Y_t + T_t - EC_t - A_{t+1}\Big] & \text{(household budget)} \\ Y_t - C_t - G + \rho_0 A^-_t - EC_t, \qquad A_{t+1} - B_t & \text{(goods and bond clearing)} \\ \text{Phillips, Taylor, Fisher, fiscal rule} & \text{(macro rows)} \end{pmatrix}, \tag{22}$$

where $\mathcal{E}$ is the household block's backward Euler operator, while $\mathcal{J}$ and $\mathcal{M}$ are the boundary recursions stated 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 $A^-_t \le 0$, intermediation costs are $-\rho_0 A^-_t \ge 0$; both intermediation and entry costs absorb goods rather than finance the government. Combining the household budget with goods and bond clearing gives the aggregate-savings row on which NDU operates. It also implies the government budget identity in (53), so that identity is a Walras-law check rather than an additional independent row of (22). 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. This is why a debt-entry extension cannot be represented by changing the borrowing constraint alone.

Linearization. The linear solution replaces every row of (22) by its first-order expansion around the steady state and solves the resulting linear system, giving the decision rule and law of motion

$$y_t = g_x x_t, \qquad x_{t+1} = h_x x_t. \tag{23}$$

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

$$dY_t \;=\; \bar{R}\, DA[F_{ss}, m^B_{ss}], dX_t + \bar{A}\, dr^{lag}_t + (1-\tau)\omega\, dY_t + dT_t \;-\; d\,A\big[\mathcal{T}\big], \tag{24}$$

Here $DA[F_{ss}, m^B_{ss}]dX_t$ 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 $A[\mathcal{T}(F, m^B;\, \bar{x} + dc, J, M^{0^-}, \iota, R)]$ with respect to each argument, evaluated at the steady state and applied to $(dX_t, dc_t, dJ_t, dM^{0^-}_t, dY_t, dr^{lag}_t, dT_t)$. The distribution deviation dX_t includes dm^B_t, and $\bar{x}$ 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 (23) from $t = 1$ with $x_1 = \Delta\, e_\varepsilon$, 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 (22) and replaces two parts by their exact nonlinear counterparts:

- the *distribution rows*: instead of the linearized law of motion (the first block of (22) 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 (24), 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_x and h_x 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_x. 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_x and shapes every subsequent period. Operationally, the period-1 solve proceeds in three steps.

Step 1: First-order policy. The period-1 state is $x_1 = \Delta\, e_\varepsilon$: 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,

$$(dc_1, dJ_1, dM_1^{0^-}) = \Delta \cdot (g_x^{(c)}, g_x^{(J)}, g_x^{(M)}) e_\varepsilon. \tag{25}$$

Because g_x solves the full forward-looking system, these controls contain the linear path's response to future prices. Their levels are $J_1 = \bar{J} + dJ_1$ and $M_1^{0^-} = \bar{M}^{0^-} + dM_1^{0^-}$. Today's income enters separately, through $\iota(Y_1, \bar{T} + \Delta)$, inside the distribution update.

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 $\bar{x} + dc_1$,

$$A_2(Y_1) = A\Big[\mathcal{T}\Big(F_{ss}, m_{ss}^B;\, \bar{x} + dc_1, J_1, M_1^{0^-}, \iota(Y_1, \bar{T} + \Delta), \bar{R}\Big)\Big], \tag{26}$$

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

$$\boxed{Y_1 = C_1(Y_1) + G - \rho_0 \bar{A}^- + EC_1(Y_1).} \tag{27}$$

Equation (27) is the aggregate-savings row of (22), restored to its exact form: comparing with the linearized row (24), 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 (27) 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 is 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

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

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

In the continuous-policy benchmark, if the saving policy is also one-Lipschitz, then $0 \leq A_2'(Y_1) \leq (1-\tau)\omega$ wherever the derivative exists. The original Keynesian-cross map then has slope between zero and $(1-\tau)\omega$ and is a contraction. The debt-entry policy instead has an upward saving jump at 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. The baseline implementation therefore solves $H(Y_1) = 0$ by safeguarded Newton/bisection rather than relying on Picard iteration.

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).[7]

[7] The 2×2 system is

$$r_1 = \bar{r} + \phi_\pi \kappa_w \widehat{g} + (\phi_\pi \bar{\beta} - \bar{R})\pi_2, \quad \pi_2 = g_x^{(\pi)} x_2(r_1),$$

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 $x_2(r_1)$ depends affinely on r_1 through its lagged-rate coordinate.

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

$$x_2 = [\, X_2 - X_{ss};\;\; r_1 - \bar{r};\;\; 0\,],$$

where X_2 contains both F_2 and m_2^B. This state is passed to the linear continuation (23) 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.

5.1 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 $X_t = (F_t, m_t^B)$ (with $X_1 = X_{ss}$); 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 $(dc_t, dJ_t, dM_t^{0^-}) = (g_x^{(c)}, g_x^{(J)}, g_x^{(M)}) x_t$ with $x_t = [\, X_t - X_{ss};\, r_{t-1} - \bar{r};\, \varepsilon_t\,]$, 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 (25).

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 $x_{K+1} = [\, X_{K+1} - X_{ss};\, r_K - \bar{r};\, 0\,]$ is passed to (23) 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_x, 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.[8]

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

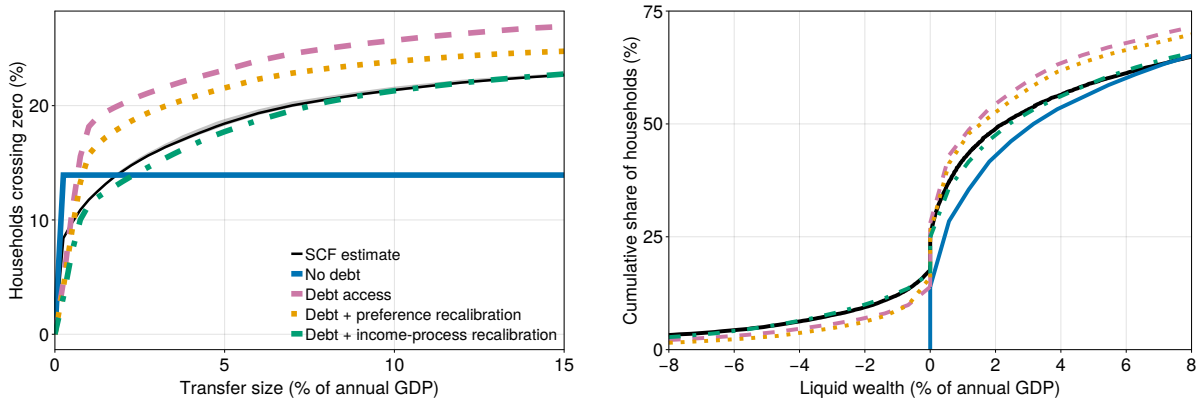
For type θ, a transfer moves the mass

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

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.[9]

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.

[8] One can also update household expectations inside the window. We leave this extension for future work.
[9] Appendix A.7 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 A.1, A.7, and A.6 give definitions and construction details.

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 + income-process recalibration* also recalibrates the income process and is our baseline debt model. Appendix A reports the parameters, targets, and fit of all four variants.

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 1.04 percentage points in the baseline debt model (Table A.5).

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 entry cost therefore places mass on both sides of zero in a way that is consistent with the crossing term in (15).[10]

[10]Appendix A.5 compares this mechanism with a wedge-only model. We interpret the entry cost as a reduced

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. As specified in Section 2, θ here includes the household's beginning borrower status and current waiver branch: the transfer shifts cash-on-hand but does not reclassify an incumbent borrower as a non-borrower before the choice. This experiment isolates the effect of moving the initial distribution across policy boundaries.

The impact response can be computed in two equivalent ways:

$$C(\Delta) = \int_\Theta \int c_\theta(x+\Delta)\, dF_\theta(x)\, \mu(d\theta) \tag{28}$$

$$= \mathcal{X}(\Delta) - \int_\Theta \int g_\theta(x+\Delta)\, dF_\theta(x)\, \mu(d\theta) - EC(\Delta)$$

$$= \mathcal{X}(\Delta) - \int_\Theta \int a'\, d\widetilde{F}^a_{\theta,\Delta}(a')\, \mu(d\theta) - EC(\Delta), \tag{29}$$

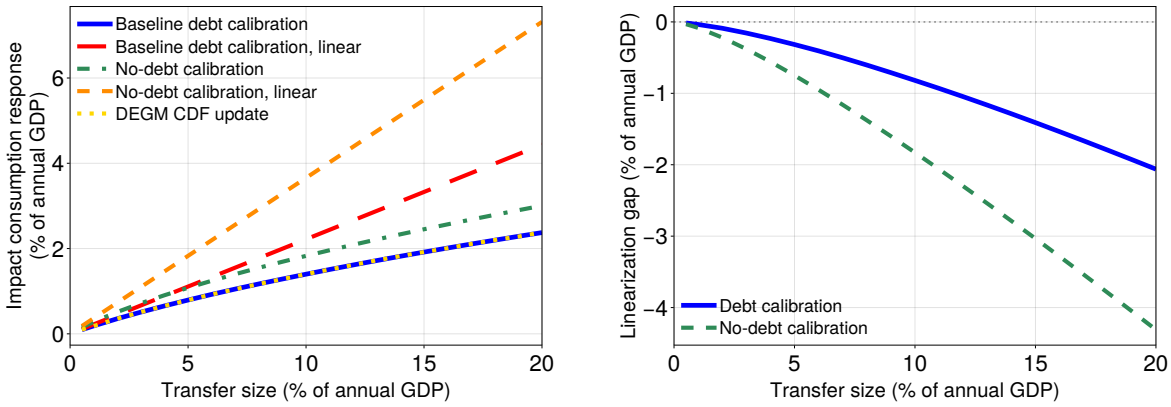
where $\mathcal{X}(\Delta) = \int_\Theta \int (x+\Delta)\, dF_\theta(x)\, \mu(d\theta)$ is aggregate cash-on-hand including the transfer. The first line integrates the consumption policy at shifted cash-on-hand over the stationary distribution. The second line uses the household budget identity, including the entry cost paid by non-waived entrants. The third line changes variables from cash-on-hand to end-of-period assets. With a continuous single policy, $\widetilde{F}^a_{\theta,\Delta}$ is the shifted distribution in (7)–(9). The debt-entry calibration applies the same change of variables branch by branch using the side-aware operator in (55)–(56); in that case $EC(\Delta) = \kappa e(\Delta)$, with the entry flow defined in (50).

The two expressions are identical, but they use different intermediate calculations. Equation (28) interpolates the consumption policy while keeping the initial distribution fixed. Equation (29) 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 baseline debt model predicts less consumption at every reported transfer size. The difference grows from 0.05 percent of annual GDP at a 0.5-percent transfer to 0.63 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 holds the baseline debt model fixed and 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 DEGM CDF update 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 59 percent at a 10-percent transfer

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 DEGM CDF update computes the baseline-model response from the shifted distribution using (29) 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.

and 87 percent at a 20-percent transfer. The nonlinear CDF update therefore captures the distributional curvature without approximating the moved cutoff. Appendix B.3 reports the same experiment in table form and adds the second-order approximation, including the policy-boundary terms derived in Appendix B.1.

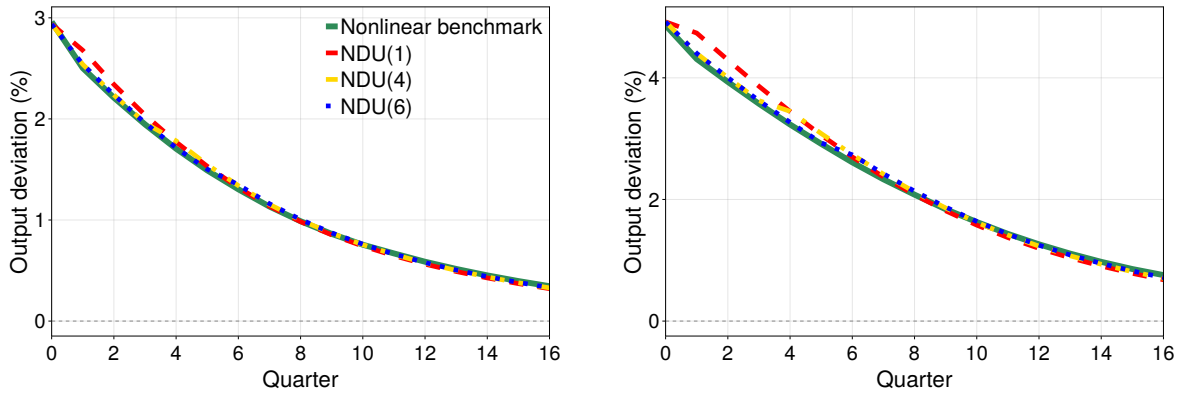
The partial-equilibrium exercise establishes two points. The transfer response is less concave in the baseline debt model than in the no-debt model, but the linear approximation still substantially overstates large transfers. At the same time, the DEGM distribution update reproduces the nonlinear response 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.[11] 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 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 107 percent at a 5-percent transfer and 153 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 6.8 basis points.

[11] 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.

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

Transfer size	Nonlinear	Linear	NDU(1)	NDU(6)
0.5% GDP	0.402	0.613	0.398	0.398
1.0% GDP	0.730	1.225	0.718	0.718
2.0% GDP	1.388	2.451	1.366	1.366
5.0% GDP	2.956	6.127	2.938	2.938
10.0% GDP	4.853	12.254	4.921	4.921

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.

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 post-impact 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 post-impact error from 0.438 to 0.104 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 post-impact 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	0.006	0.001	0.001	1.66	2.79	4.46	8.58
1.0% GDP	0.043	0.007	0.007	1.65	2.99	4.63	7.68
5.0% GDP	0.181	0.083	0.041	1.66	3.04	5.11	107.76
10.0% GDP	0.438	0.230	0.104	1.65	2.86	4.49	141.60

Notes: Sticky-wage HANK economy, state-space solution, $N_A = 120$, log-log grid, $T = 150$. Peak post-impact 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 five seconds per experiment, while the conditional nonlinear benchmark takes between 8 and 142 seconds. Averaged across the four transfer sizes, the nonlinear transition costs 66 seconds and NDU(6) costs 4.7 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 Four Model Variants, Calibration, and Empirical Fit

This appendix documents the four specifications used in the empirical comparison. It first defines the model ladder and its calibration, then maps the SCF into the models and connects distributional fit to the crossing-share and aggregate-consumption evidence. Two qualifications are important throughout. First, the steps 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.

A.1 Model Ladder and Common Environment

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. The tax-transfer schedule is $y - T(y) = (1-\tau)y + \bar{T}$, with $\tau = 0.274$ and a lump-sum transfer equal to 15 percent of GDP, as in Bianchi and Kaplan (2026). All four models have five 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 A.1.

Table A.1: 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.

A.2 Income Process across the Model Variants

Bianchi and Kaplan follow 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,

$$\log \ell_t = z_t + \epsilon_t - \log \mathbb{E}[\exp(z_t + \epsilon_t)], \tag{30}$$

so average efficiency is one. The persistent component is represented on eleven Rouwenhorst nodes. Let $Q(\phi_z, \sigma_\eta^2)$ denote the Rouwenhorst transition calibrated to the AR(1)

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

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

$$P_z = (1 - \lambda_\eta)I + \lambda_\eta Q(\phi_z, \sigma_\eta^2). \tag{31}$$

Thus λ_η is the probability that the persistent state is refreshed; otherwise it remains unchanged. Conditional on a transitory shock,

$$\epsilon_t \sim N(-\sigma_\epsilon^2/2, \sigma_\epsilon^2) \quad \text{with probability } \lambda_\epsilon, \qquad \epsilon_t = 0 \quad \text{with probability } 1 - \lambda_\epsilon. \tag{32}$$

The Gaussian component is integrated with five Gauss–Hermite nodes, so the discretized income process has eleven persistent states and six transitory outcomes. The mean correction in (32), followed by the normalization in (30), 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 this 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 A.2.

Table A.2: Income-process parameters by model variant

Specification	ϕ_z	σ_η^2	λ_η	σ_ϵ^2	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 comparison therefore holds the income process and uniform discount-type masses fixed. Its beta grid has the same spacing as A's but a very small common level shift, reported below. B1–B2 holds the income process fixed but changes preferences and debt parameters. B2–C changes the income process together with preferences and debt parameters; it is a sequential fit comparison, not a one-parameter experiment. Figure A.4 shows the resulting income–wealth sorting directly.

A.3 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 A.3.

Table A.3: 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	(0.9661, 0.9725, 0.9793, 0.9858, 0.9921)	(1, 2, 2, 1, 1)/7

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. B1 uses $\underline{a} = -1$, $\rho_0 = 0.05$, $\kappa = 0.03$, and waiver probability $\lambda = 0.25$. B2 uses $(\rho_0, \kappa, \lambda) = (0.038, 0.02, 0.45)$, and C uses $(\rho_0, \kappa, \lambda) = (0.042, 0.02, 0.25)$. 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, not as isolated causal effects of one parameter.

A.4 Calibration Target and Fit

The empirical target is the marginal distribution of SCF liquid wealth, pooling the twelve triennial waves from 1989 through 2022 with equal wave weights. Wealth is expressed as a percent of annual GDP. The criterion contains thirteen CDF levels and the discrete mass at zero, all in percentage points (Table A.4):

Table A.4: 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 objective sums squared deviations from the targeted moments and adds an aggregate-asset target and a one-sided 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 objective values cannot be compared when the number of reachable moments differs. For model M, let $\mathcal{M}_M$ be its set of reachable moments, let $n = |\mathcal{M}_M|$, and let m_M and m^{SCF} denote the corresponding model and empirical moment values. We therefore report

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

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 this criterion 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 A.5: 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	1.04	1.34

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.

A.5 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. 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 A.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 -0.94 for the entry-cost model. Yet its errors at $F(-6)$ and $F(-3)$ are -1.89 and -2.50 percentage points, compared with -0.23 and $+0.27$ with the entry

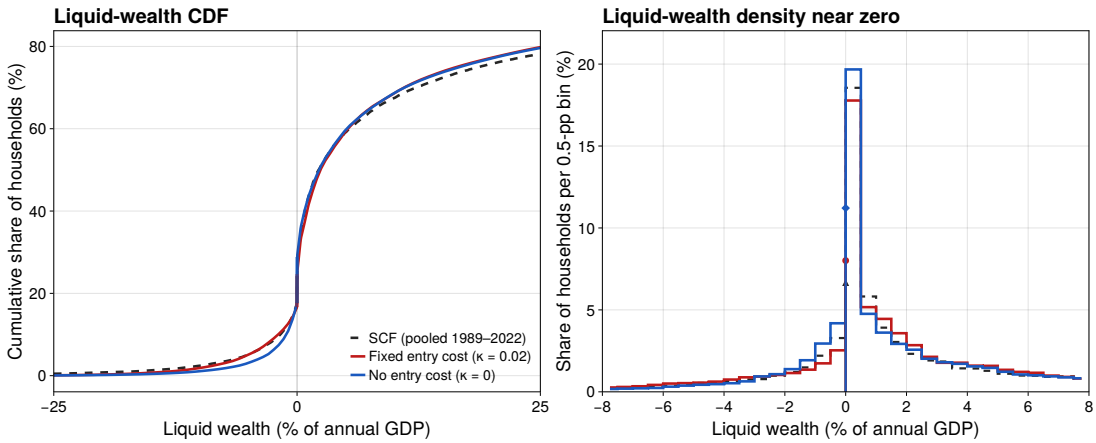


Figure A.1: **Liquid wealth around zero with and without a fixed entry cost.**

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. The fixed entry cost reproduces this asymmetry, whereas the wedge-only calibration places too much probability just below and at zero.

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 1.33 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, $(\kappa, \lambda, \rho_0) = (0.02, 0.25, 0.042)$, and the wedge-only specification, $(\kappa, \rho_0) = (0, 0.056)$, independently minimize the same objective on the same 180-point grid with 30 negative nodes. Although the entry-cost triple equals C's debt parameters in Table A.1, 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. Because both use the same fourteen moments, their raw criteria are comparable: the CDF-anchor contributions are 14.96 and 22.60, and including the zero-wealth mass widens the full comparison to 16.72 versus 57.28. Across the wedge-only parameter grid, zero-wealth masses range from 11.2 to 12.1 percent; every specification that matches the borrower share to within one percentage point violates the calibration's cap on the income–wealth gradient. Variation in the borrowing spread therefore does not resolve the wedge-only model's joint failure.

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.

A.6 Mapping the SCF into Model Objects

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.

For any model, household i's counterfactual consumption response averages the model policy over the posterior distribution of discount types in its wealth-rank and income-rank cell. Beginning wealth plus the transfer, net of this consumption response, gives the household's end-of-period asset position. A household counts as crossing when it begins weakly below zero and ends strictly above zero. We apply the same definition to the model distribution.

In the debt specifications, negative SCF wealth is evaluated on the negative policy support. Only in the no-debt specification do we use the boundary policy for below-zero observations, because negative assets lie outside that model's feasible set. The resulting SCF crossing curves are nearly invariant to the chosen policy system. Aggregate consumption remains model-dependent because it depends on the amount consumed by every household, as explained in Appendix A.7.

A.7 Crossing Shares and Finite-Transfer Consumption

The zero-crossing curve is an extensive-margin diagnostic: it records whether a household moves from weakly negative to positive liquid wealth after its consumption response. The SCF curves obtained from the four policy systems differ by at most about 0.6 percentage points. Hence the empirical crossing share is governed almost entirely by the mass of the wealth distribution near zero; the precise consumption policy matters little for this binary event. Model A's large zero discrete mass produces a low, nearly flat curve after the mass of households at zero-wealth crosses, whereas debt spreads mass below zero and lets the crossing share continue to rise with transfer size.

The threshold also has a direct MPC interpretation. Table A.6 decomposes each model's exact fixed-price response by beginning liquid wealth. The statistic is the stationary-distribution-

weighted effective MPC, $[c(x+\Delta)-c(x)]/\Delta$, conditional on being weakly below or strictly above zero. The zero-wealth discrete mass 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 A.6: 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	0.139	0.469	0.231	2.03
Debt access	0.280	0.359	0.193	1.86
Debt + preference recalibration	0.261	0.307	0.169	1.82
Debt + income-process recalibration	0.249	0.245	0.175	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 is the mass exactly at their 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. The zero-crossing curve therefore measures entry into the side of the wealth distribution along which this decline occurs.

This observation does not make the crossing share a sufficient statistic for aggregate consumption. The latter integrates the amount consumed by every household, rather than a zero-one crossing indicator. Holding a distribution fixed, fewer crossings after the consumption correction naturally correspond to more consumption; there is therefore no contradiction between A's relatively flat crossing curve and its large aggregate response. Across models, however, distributions and policies both change. In particular, B2 has the larger small-transfer response, while C overtakes it at a transfer of roughly 1.5 percent of annual GDP because C's effective MPC declines more slowly. The crossing curve identifies the distributional source of curvature, but it cannot by itself rank finite-transfer consumption responses.

Constructing the paired SCF response. The crossing curve alone does not determine consumption, so we next explain how we construct the paired SCF response used in Section 6.1. All dollar quantities are first put in common GDP units. Let annual GDP per household be Y^{ann} dollars; this is distinct from government spending G. One model unit equals $Y^{\text{ann}}/4$ dollars. For SCF wealth W_i and a dollar transfer D, write $a_i = W_i/(Y^{\text{ann}}/4)$ and $d = D/(Y^{\text{ann}}/4)$. For policy evaluation, let $\widetilde{a}_{M,i} = a_i$ in the debt models and $\widetilde{a}_{A,i} = \max\{a_i, 0\}$ in the no-debt model. Let $s_{M,i} = B$ when $M \in \{B1, B2, C\}$ and $a_i < 0$, and let $s_{M,i} = N$ otherwise. Thus the no-debt

model evaluates negative SCF observations at its zero borrowing limit while the debt models preserve the observed negative position. Define

$$\mathcal{C}^{\mathrm{SCF}}_{M,s}(a,d;\beta,z) = \sum_k w_k \sum_{q\in\mathcal{Q}_s} \pi_M(q \mid s)\, c^q_{M,s}\Big(R_{M,s}a + \iota^M_{zk} + d;\beta,z\Big),$$
$$\Delta C_i = \frac{Y^{\mathrm{ann}}}{4}\left[\mathcal{C}^{\mathrm{SCF}}_{M,s_{M,i}}(\widetilde{a}_{M,i},d;\beta,z_i) - \mathcal{C}^{\mathrm{SCF}}_{M,s_{M,i}}(\widetilde{a}_{M,i},0;\beta,z_i)\right]. \tag{33}$$

Here $M \in \{A, B1, B2, C\}$ indexes the model variant, β is the household's discount factor, z_i is its assigned persistent-income state, and ι^M_{zk} is steady-state disposable non-asset income in persistent state z and transitory state k. In the debt models, $R_{M,B} = \bar{R}_M + \rho_{0,M}$ and $R_{M,N} = \bar{R}_M$. For incumbent borrowers, $\mathcal{Q}_B = \{B\}$ and $\pi_M(B \mid B) = 1$. For non-borrowers, $\mathcal{Q}_N = \{F, K\}$, $\pi_M(F \mid N) = \lambda_M$, and $\pi_M(K \mid N) = 1 - \lambda_M$. In the no-debt model the inner sum collapses to its single policy. The transfer d is added directly to cash-on-hand, and the beginning status $s_{M,i}$ is held fixed inside both policy evaluations. Thus a large transfer can move end-of-period assets across zero without mechanically switching the beginning-of-period return or entry regime. The factor $Y^{\mathrm{ann}}/4$ is the unit conversion; it neither rescales policies by household income nor imposes homogeneity. Income rank assigns z_i, but cardinal household income does not enter the zero-crossing statistic reported in Figure 1.

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

$$\Pr_M(\beta \mid \text{wealth-rank cell}, \text{income-rank cell}).$$

The main benchmark preserves the empirical wealth–income copula and uses the model-implied discount-factor distribution conditional on the household's wealth-rank and income-rank cell. All occupied cells in the 40×40 rank partition have model support. Table A.7 reports the resulting model-minus-SCF gaps at three transfer sizes.

Table A.7: Model response minus its paired SCF response

Model	$\Delta = 0.5\%$	$\Delta = 5\%$	$\Delta = 15\%$
A	0.0373	0.1152	0.0464
B1	0.0221	0.0764	0.0938
B2	0.0176	0.0526	0.0618
C	0.0080	0.0204	0.0326

Notes: 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.

Model C is closest to its paired SCF response throughout the curve, followed by B2 over small and medium transfers; at the largest reported transfer B1 is the worst-fitting specification. A is worst over small and medium transfers, so its smaller endpoint gap is not evidence of a good curve fit. The common-GDP conversion in (33) is essential for this comparison: assets,

cash-on-hand, transfers, and consumption must all be evaluated in the same aggregate units, and no homogeneity restriction permits household-specific rescaling by income levels.

A bridge from the SCF evaluation to the model integral. Figure A.2 changes one part of the empirical evaluation at a time. It begins with the SCF joint distribution in common GDP units. The first step replaces the empirical wealth–income rank copula with the model copula while preserving the represented marginal distributions. The second replaces the residualized-rank assignment of the persistent income state by the model's exact mixture. The third installs the model-implied discount-type composition by wealth–income cell, including the remaining beta–income dependence within each cell. The fourth replaces empirical wealth levels by the model wealth-rank levels at which its policy is evaluated. The top row shows the response curves generated by the same sequence; the bottom row reports the contribution of each step for three transfer sizes.

The bridge separates a mismatch in the wealth–income copula from its effect on the consumption integral. The mismatch is visibly large for the earlier variants and smallest for C, but its effect also depends on the policy system and transfer size. At a 15-percent transfer, replacing the SCF copula with the model copula changes the response by -0.106 percentage points of annual GDP in A, -0.032 in B1, -0.025 in B2, and only -0.006 in C. The effects are smaller at small and medium transfers, though not uniformly zero.

The remaining steps are more modest. Reassigning persistent states has little effect throughout. At a 15-percent transfer, the cell-specific discount-type posterior used in the main paired benchmark reproduces the exact continuous-model sorting effect to within 0.002 percentage points, while the remaining beta–income dependence within a rank cell contributes at most 0.002 points. Replacing empirical wealth levels by model wealth levels provides another observable step. After all four steps, the remaining difference from the exact model response is about 0.04–0.05 percentage points at a 15-percent transfer in every variant. This remainder reflects continuous wealth variation and policy composition within the finite rank cells, together with the difference between the rank-cell calculation and the exact model integral.

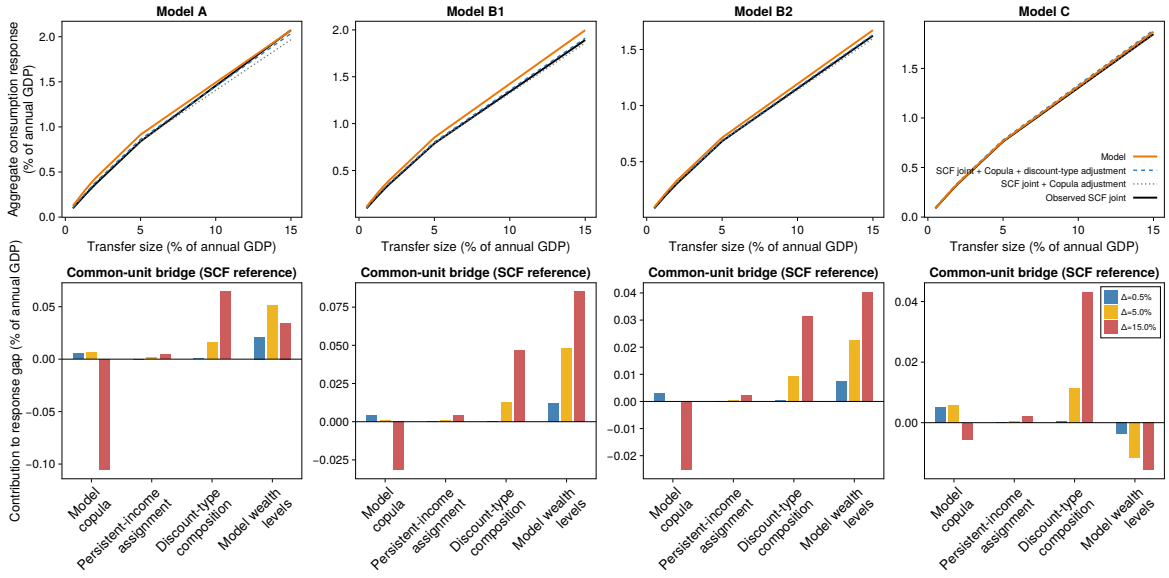


Figure A.2: **Bridge from the SCF distribution to model counterfactuals.**

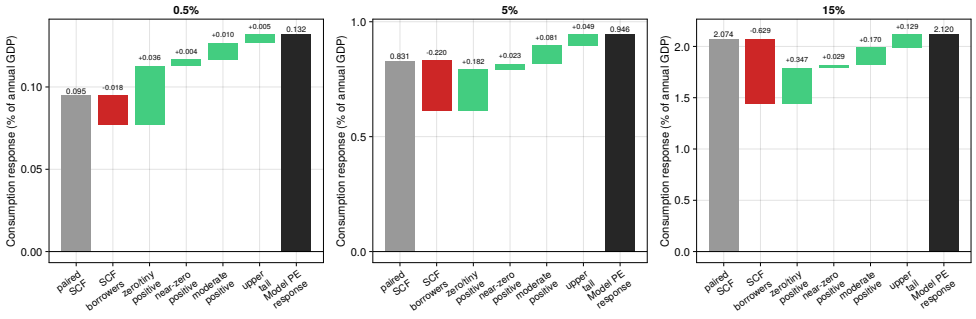
Notes: The top row traces the response from the observed SCF joint distribution to a common-unit rank-cell counterfactual. The bottom row decomposes this change into the model wealth–income copula, persistent-income assignment, discount-type composition, and model wealth levels at transfers of 0.5, 5, and 15 percent of annual GDP. The final orange curve is the rank-cell counterfactual, not the exact model integral. Labels A, B1, B2, and C correspond to the specifications in Table A.1.

These results sharpen the interpretation of the aggregate ranking. C matches the empirical income–wealth sorting much more closely than the earlier variants (Figure A.4), and its copula transplant has correspondingly little effect on consumption. This implication does not extend to every specification: A combines a large sorting discrepancy with a substantial copula contribution at large transfers. In the calibration's directly comparable projection—the slope of mean persistent-income rank on liquid-wealth percentile above p20—the SCF value is 0.113, compared with 0.240 in A, 0.201 in B1, 0.233 in B2, and 0.176 in C. C therefore delivers the closest fit, while the intermediate calibration steps need not improve this joint-distribution moment monotonically.

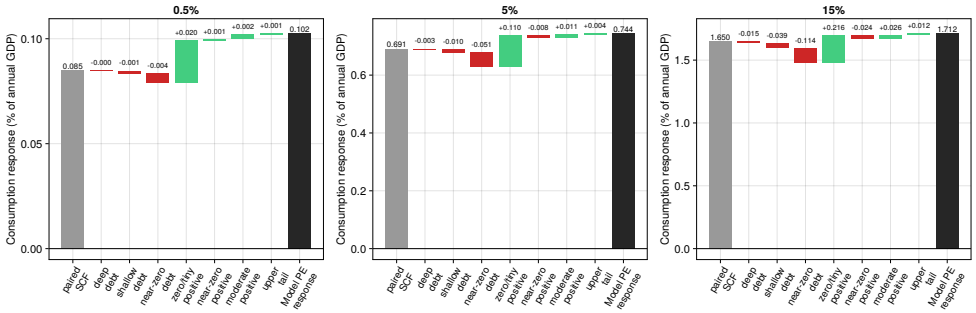
A wealth-bin decomposition explains why A can look close at an isolated transfer size, using the same common-GDP, cell-specific-posterior benchmark as the headline paired curve. The SCF endpoint is the twelve-wave mean of the exact common-GDP paired-SCF bin contributions evaluated with the model-implied posterior $\Pr_M(\beta \mid \text{wealth-rank cell}, \text{income-rank cell})$; the model endpoint is the exact model partial-equilibrium response. Each bar is therefore an exact additive contribution to the displayed model-minus-paired-SCF gap, subject only to the across-wave averaging.

Figures A.3 and A.4 show why the aggregate ranking cannot be read from a single distributional moment. In Figure A.3, each bar is a contribution to an integral, not the response of a representative household in that bin; its sign indicates whether that bin raises or lowers the displayed gap, and the bars sum exactly to the displayed endpoints. Panel (a) shows that the absence of negative wealth in A and the gaps on common support work in opposite directions. Panels (b) and (c) show that B2 and C remove most of the mechanically missing debt-region contribution. C's residual is smaller and spread more evenly across the common support, mir-

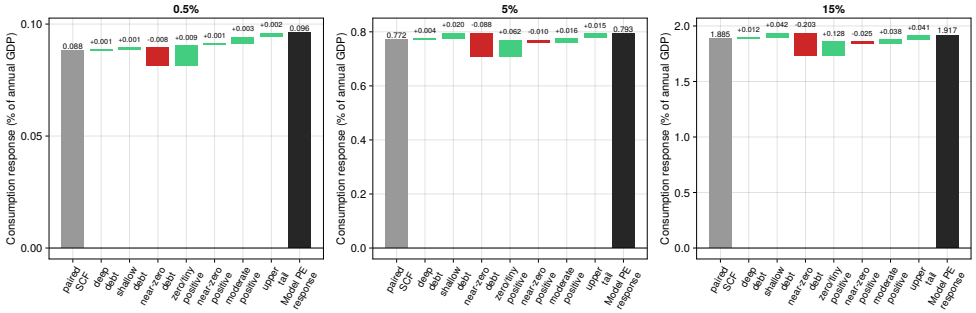
roring its substantially better fit to the marginal wealth distribution (RMSE 1.04 versus 3.47 for B2; Table A.5) rather than any single offsetting bin. Because the waterfalls use the same cell-specific discount-type posteriors as the main paired curve, they provide an exact additive decomposition of the displayed model-minus-paired-SCF gap at each transfer size. Figure A.4 then shows the corresponding income–wealth sorting across the full wealth distribution.



(a) Model A: no debt



(b) Model B2: recalibrated debt model



(c) Model C: baseline debt model

Figure A.3: **Initial-wealth contributions to the paired model-minus-SCF gap.**

Notes: Bars decompose each model's partial-equilibrium consumption-response gap relative to its paired SCF benchmark by initial liquid-wealth bin. The first and last bars are the paired SCF and exact model responses. Positive contributions widen the model-minus-SCF gap; negative contributions narrow it. Values are percentages of annual GDP. The bars sum to the endpoint gap, subject only to across-wave averaging. Equation (33) gives the common-unit mapping.

The income-rank curves illustrate the joint-distribution problem. The SCF relationship is relatively flat over most of the wealth distribution, whereas all four models sort persistent income too strongly into wealth above p20. The no-debt curve is the steepest and the baseline debt model is closest to the data. The sequential transplant shows why visual fit and aggregate importance must nevertheless be reported separately. For C, the remaining copula discrepancy has little effect on the consumption curve. For A, the discrepancy is large and its effect becomes quantitatively important at large transfers; B1 and B2 lie between these cases. Income–wealth sorting is therefore one genuine dimension improved by the baseline debt model, but it is not a sufficient account of the aggregate ranking. Wealth support and levels and the composition of responses within wealth–income cells also enter the consumption integral.

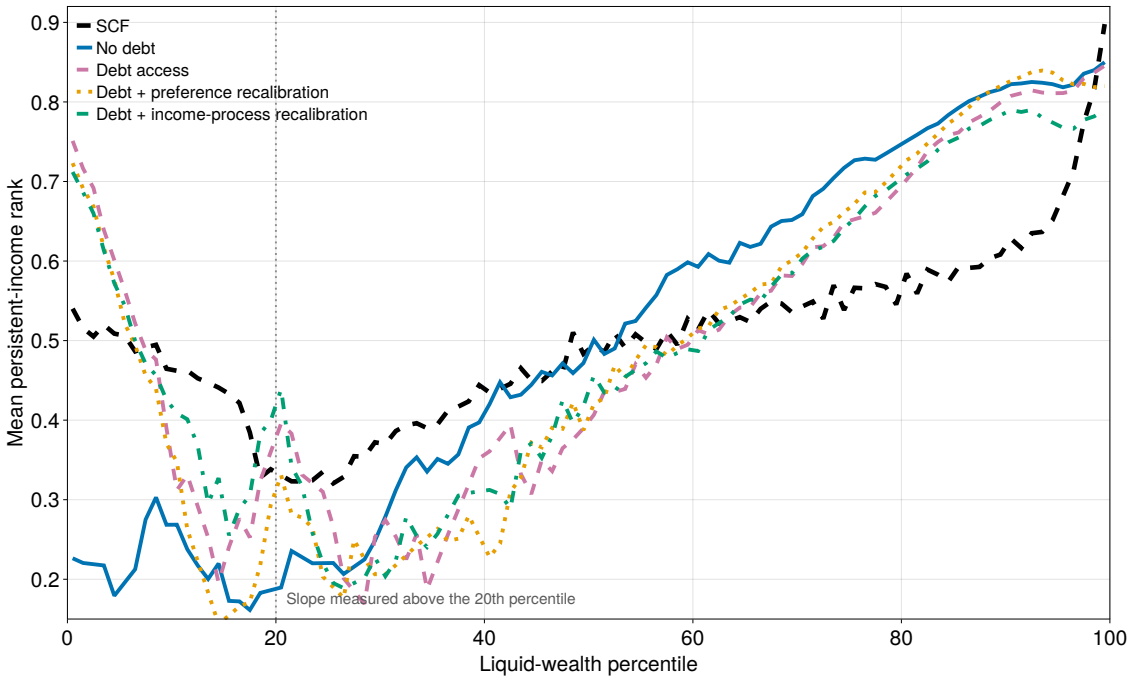


Figure A.4: **Mean persistent-income rank by liquid-wealth percentile.**

Notes: 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 line marks the 20th liquid-wealth percentile, above which the slope is measured. The slopes are 0.113 in the SCF, 0.240 with no debt, 0.201 with debt access, 0.233 after preference recalibration, and 0.176 after income-process recalibration. Figure A.2 measures the effect of replacing the full rank copula.

B 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.

B.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

$$c_\theta = \lambda\, c^w_\theta + (1-\lambda)\, c^n_\theta \tag{34}$$

of the fee-free policy c^w_θ and the costly policy c^n_θ. Both are piecewise smooth, but only c^n_θ has a cutoff at $x^\star_\theta$. 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* (11)*, with c_θ the mixture* (34)*. 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 $m^\pm_\theta$ and $m^\pm_{\theta,x}$ are dominated, so that each branch may be differentiated under the integral. Then, for regular θ,

$$C'_\theta(\Delta) = \int m_\theta(x+\Delta)\, dF_\theta(x) + \sum_{k=1}^{K_\theta} J_{\theta,k}\, f_\theta(q_{\theta,k}(\Delta))\,, \tag{35}$$

and

$$C''_\theta(\Delta) = \int m_{\theta,x}(x+\Delta)\, dF_\theta(x) + \sum_{k=1}^{K_\theta} \left[[m]_{\theta,k}\, f_\theta(q_{\theta,k}(\Delta)) - J_{\theta,k}\, f'_\theta(q_{\theta,k}(\Delta)) \right], \tag{36}$$

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,

$$C_\theta(\Delta) = \sum_{k=0}^{K_\theta} \int_{q_k}^{q_{k+1}} c_\theta^{(k)}(x+\Delta)\, dF_\theta(x), \tag{37}$$

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 (37) 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 (35). 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 (36). Integrating over θ gives the aggregate derivatives; the mass-point exceptions do not contribute because the set of affected types has μ-measure zero. □

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 (15). 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.*

B.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:

$$m(x,y) = \frac{\Lambda}{u''(c) + \Lambda}, \qquad \Lambda \equiv \beta R^2\,\mathbb{E}\left[u''(c')\,m' \mid y\right], \tag{38}$$

$$m_x(x,y) = \frac{[R(1-m)]^2\Psi - u'''(c)\,m^2}{u''(c) + \Lambda}, \qquad \Psi \equiv \beta R\,\mathbb{E}\left[u'''(c')(m')^2 + u''(c')\,m'_x \mid y\right], \tag{39}$$

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 B.1 combines those sums with $\mathbb{E}[m]$ and $\mathbb{E}[m_x]$.

B.3 Partial-Equilibrium Evidence

Table B.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 B.2 with the policy-boundary sums in Proposition 3.

Table B.1: Partial-equilibrium impact consumption response

Transfer size	Nonlinear	FOPS	SOPS	DEGM CDF update
0.5% GDP	0.0961	0.1179	0.1163	0.1006
1.0% GDP	0.1860	0.2358	0.2292	0.1815
2.0% GDP	0.3550	0.4716	0.4452	0.3561
5.0% GDP	0.7925	1.1791	1.0141	0.7897
10.0% GDP	1.3988	2.3582	1.6982	1.4001
20.0% GDP	2.3745	4.7164	2.0764	2.3800

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, $\kappa = 0.02$, $\lambda = 0.25$); the Nonlinear and DEGM CDF 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 B.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. Under the borrowing extension the gross return in the Euler recursions is side-dependent and the continuation expectation carries the entry-fee lottery, exactly as in the household problem; the entry-threshold jump enters the boundary sums weighted by $1 - \lambda$, since the fee-free policy is continuous there. 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. The DEGM CDF update 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 $\mathbb{E}[m] = 0.2417$. The entry-threshold jump contributes $\sum_k J_k f(\bar{x}_k) = -0.0059$, giving $C'(0) = 0.2358$. The extensive margin therefore lowers the first-order aggregate response, but only by two and a half percent.

At second order, the boundary terms dominate. Household-level curvature contributes $\mathbb{E}[m_x] = -0.1031$, while the boundary sums contribute -0.2270, giving $C''(0) = -0.3300$. Thus more than two thirds of the aggregate second derivative comes from households moving across policy branches rather than from curvature within a branch. Almost all of the boundary contribution is the slope-jump term $\sum_k [m]_k f(\bar{x}_k) = -0.2146$ at the borrowing constraint and the bunching edges. The density-slope term $-\sum_k J_k f'(\bar{x}_k) = -0.0124$ accounts for less than one tenth.

Because most of the aggregate concavity comes from boundary crossing, the second-order correction becomes more useful as the transfer grows. It removes 8 percent of the first-order error at a 0.5-percent transfer and 13 percent at a 1-percent transfer, compared with 43 percent at 5 percent and 69 percent at 10 percent. At a 20-percent transfer, however, the second-order approximation overshoots and predicts less consumption than the nonlinear solution. The DEGM CDF update avoids this limitation because it evaluates the moved cutoff in (9) directly rather than expanding around zero.

C Implementing Debt Entry Costs in State Space

This appendix describes how the fixed debt-entry cost is carried through the state-space solution. The central issue is not the household problem alone. The return kink and entry cost create both a point mass at zero and a distinct mass strictly below zero. The nonlinear distribution update must track both masses exactly while remaining differentiable. We first describe the household policy and the additional boundary controls, then derive resource costs, the distribution state at zero, and the exact transition. We close with the aggregate block, numerical implementation, and validation.

C.1 Timing and Household Policy

Beginning-of-period borrower status requires no additional individual state: a household is an incumbent borrower exactly when $a < 0$. Incumbent borrowers may continue borrowing or exit freely. A non-borrower with $a \geq 0$ draws an iid waiver before choosing. With probability λ the household uses the unrestricted policy; otherwise 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. Utility is CRRA with $u'(c) = c^{-\sigma}$; the maintained log-utility calibration sets $\sigma = 1$.

Let $q \in \{F, K\}$ denote the waiver draw: $q = F$ with probability λ, and $q = K$ otherwise. Let $y_{hk,t} = w_{h,k} Y_t$ denote pre-tax labor income in state (h, k), and define disposable non-asset income by $\iota_{hk,t} = (1 - \tau) y_{hk,t} + T_t$. When time or state indices are suppressed, we write these objects as y_{hk} and ι_{hk}. Define the gross saving return $R_t = 1 + r_{t-1}$ and the return on beginning assets by

$$R_t(a) = R_t + \rho_0 \mathbf{1}\{a < 0\}, \qquad x = R_t(a) a + \iota_{hk,t}. \tag{40}$$

A household of permanent type h, with discount factor β_h, solves

$$V^q_{t,h}(x; a) = \max_{a' \geq \underline{a}} \left\{ u(c) + \beta_h \mathbb{E}_t [V_{t+1,h'}(x'; a')] \right\}, \tag{41}$$

$$\text{s.t.} \quad c + a' = x - \kappa \mathbf{1}\{a \geq 0,\ a' < 0,\ q = K\}. \tag{42}$$

Here x' is next period's cash-on-hand, and $V_{t+1,h'}(x'; a')$ is the next-period value before the new waiver is drawn; the expectation integrates over h', the transitory draw, and that waiver. For an incumbent borrower, $a < 0$, the draw is irrelevant: continuing to borrow or exiting is free. For a non-borrower, ex-ante value is $V^N = \lambda V^F + (1 - \lambda) V^K$. The waiver is iid, so no additional household state is required.

Incumbent borrowers and waived non-borrowers use the unrestricted endogenous-grid policy. 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. If

$x_{0,h}$ is the unrestricted cash-on-hand node associated with $a' = 0$, the non-waived policy is

$$g_h^K(x) = \begin{cases} g_h^E(x), & x < x_h^\star, \\ 0, & x_h^\star \le x \le x_{0,h}, \\ g_h^S(x), & x > x_{0,h}. \end{cases} \tag{43}$$

The entry branch uses endogenous-grid cash-on-hand nodes shifted upward by κ. Its threshold is determined by value matching,

$$V_h^E(x_h^\star) = V_h^0(x_h^\star), \tag{44}$$

where V^0 is the value of remaining at zero. The entry policy can jump at $x_h^\star$: 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 maintain and verify $x_h^\star < x_{0,h}$ in every household column. Thus a household that does not enter debt chooses zero assets at the entry threshold, and the linearization can keep this ordering of policy branches fixed.

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 $c_t^{\mathrm{EGM}}(s,h)$ denote the unrestricted endogenous-grid consumption schedule for household type h, evaluated at end-of-period assets s before the household chooses among the branches in (43). Let $\mathcal{V}_t(s,h)$ be expected continuation value conditional on choosing s in state h. We define the boundary premium

$$J_{h,t} = \beta_h \big[\mathcal{V}_t(0,h) - \mathcal{V}_t(0^-,h)\big] \tag{45}$$

We then reconstruct discounted continuation-value differences from the marginal policy. The premium $J_{h,t}$ is distinct from the consumption jump $J_{\theta,k}$ in Proposition 3:

$$\mathcal{D}_{h,t}(a,b;c_t,J_{h,t}) = \int_a^b c_t^{\mathrm{EGM}}(s,h)^{-\sigma}\,ds + \big[\mathbf{1}\{b \ge 0\} - \mathbf{1}\{a \ge 0\}\big] J_{h,t}. \tag{46}$$

In the recursion below, superscript B denotes the incumbent-borrower branch and superscript K the non-waived non-borrower branch; c^B, c^K and a^B, a^K are their consumption and asset choices at the common limiting cash-on-hand. The current boundary premium obeys the forward-looking recursion

$$\begin{aligned} J_{h,t} = \beta_h(1-\lambda) \sum_{h'} \Pi_{hh'} \sum_k w_k \big[& u(c^K_{h'k,t+1}) - u(c^B_{h'k,t+1}) \\ & + \mathcal{D}_{h',t+1}(a^B_{h'k,t+1}, a^K_{h'k,t+1}; c_{t+1}, J_{h',t+1})\big]. \end{aligned} \tag{47}$$

Here the two households arrive at the common limiting cash-on-hand associated with assets 0^- and 0, but their borrower and non-borrower statuses generate the two conditional policies. The waived branch contributes zero because its problem is identical to that of an incumbent borrower.

A second boundary control records the borrower-side marginal-policy endpoint, $M_{h,t}^{0^-} =$

$\lim_{a\uparrow 0} c_{h,t}^{\mathrm{EGM}}(a)^{-\sigma}$. The left-limit Euler equation determines it:

$$M_{h,t}^{0^-} = \beta_h(R_{t+1}+\rho_0)\sum_{h'}\Pi_{hh'}\sum_k w_k \times \left[c^B_{h',t+1}(\iota_{h'k,t+1})\right]^{-\sigma}. \tag{48}$$

Here $R_{t+1} = 1 + r_t$ is the return on the asset choice made at date t. This is not the marginal value at the non-borrower node $a = 0$, which uses the saving return and the waiver mixture. Conditional on $(c_t, J_t, M_t^{0^-})$, the threshold is static. Let $\mathcal{G}_h(x^*; c, J, M)$ denote the value difference $V_h^E(x^*) - V_h^0(x^*)$. Linearizing the value-matching equation $\mathcal{G}_h(x^*_{h,t}; c_t, J_{h,t}, M^{0^-}_{h,t}) = 0$ gives

$$dx^*_{h,t} = -\frac{D_c\mathcal{G}_h\, dc_t + \mathcal{G}_{J,h}\, dJ_{h,t} + \mathcal{G}_{M,h}\, dM^{0^-}_{h,t}}{\mathcal{G}_{x^*,h}}. \tag{49}$$

Thus J and M^{0^-} are forward-looking controls, whereas x^* is a per-column derived quantity rather than an additional QZ control.

C.2 Resource Costs and Market Clearing

Debt entry affects aggregate equilibrium through both the household budget and the goods used to pay the fixed cost. Let A_t be aggregate beginning assets and $A_t^- \le 0$ their negative part. Let $F^N_{h,t}(a)$ denote the joint CDF column for beginning assets among non-borrowers in state h, defined explicitly in (54) and the display below it. We write $F^{s,<}_{h,t}(a) = F^s_{h,t}(a^-) = \Pr(a_t < a, s_t = s, h_t = h)$ for the corresponding strict CDF on side $s \in \{B, N\}$. The mass of non-waived non-borrowers that enter debt and its resource cost are

$$e_t = (1-\lambda)\sum_h\sum_k w_k F^{N,<}_{h,t}\left(\frac{x^*_{h,t} - \iota_{hk,t}}{R_t}\right), \qquad EC_t = \kappa e_t. \tag{50}$$

Aggregating (42) gives

$$C_t = R_t A_t + \rho_0 A_t^- + (1-\tau)Y_t + T_t - EC_t - A_{t+1}. \tag{51}$$

The borrowing wedge absorbs $IC_t = -\rho_0 A_t^- \ge 0$ units of goods. Entry costs are also resource costs. Goods and bond-market clearing are therefore

$$Y_t = C_t + G + IC_t + EC_t, \qquad A_{t+1} = B_t. \tag{52}$$

Neither cost is government revenue, so the government budget remains

$$B_t = (1 + r_{t-1})B_{t-1} + G + T_t - \tau Y_t. \tag{53}$$

The government budget in (53) follows algebraically from the household budget, goods clearing, and bond clearing; it is an accounting identity and a numerical check, not an additional independent equilibrium condition. The household budget and market-clearing equations, together with the fiscal rule, the wage Phillips curve written in terms of C_t, the Taylor rule, and the

Fisher equation, close the aggregate block. They also explain why aggregate consumption must be an explicit control: using the debt-free substitution $C_t = Y_t - G$ would omit both resource costs.

C.3 The Distribution State at Zero

The aggregate equations depend separately on the mass below zero and the mass exactly at zero. A single stored CDF value at zero cannot distinguish them, so the state includes the borrower mass explicitly. For each combined discount-factor and persistent-income state h, let $F_h(a) = \Pr(a_t \leq a, h_t = h)$ be its joint CDF column and define

$$m_h^B = F_h(0^-), \qquad F_h(0) = m_h^B + \Pr(a = 0, h). \tag{54}$$

The stored CDF value at zero therefore cannot identify the borrower mass because it includes the economically important discrete mass at zero wealth. The numerical state stacks the free CDF values and one separate m_h^B for every h. This is distinct from the economic statement that borrower status is already encoded by the sign of assets. The terminal CDF value fixes the total mass of each type-income column, but no normalization fixes its borrower mass; $\sum_h m_h^B$ is endogenous.

CDF evaluation is separated by side rather than corrected after interpolation. On the borrower side, a monotone interpolant $\mathcal{I}_h^B$ is constructed through the negative nodes and the endpoint $(0, m_h^B)$:

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

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 once after evaluation:

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

Thus $F_h^N(0) = F_h(0) - m_h^B$ returns the mass at zero exactly. Because the bounds are imposed through endpoint data, neither evaluation requires a minimum or maximum correction afterward. Such corrections would introduce artificial branch switches into automatic differentiation.

C.4 Exact Transition and Linearization

With the two sides of the distribution defined, the transition can evaluate each policy branch without mixing the mass at zero with the borrower mass. For every output asset node, it finds the corresponding beginning-asset threshold separately for borrowers and non-borrowers. Borrower thresholds use $R_B = R_t + \rho_0$, while non-borrower thresholds use $R_S = R_t$.

The waived fraction uses the unrestricted threshold. For the non-waived fraction, a negative

output node uses the entry branch, with its cash-on-hand threshold capped at $x_h^\star$. A nonnegative output node 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 node a_j is $S[j,h]$; primes below denote next-period quantities:

$$\begin{aligned} S[j,h] = \sum_k w_k \Bigg[F_h^B\left(\frac{x_{j,h} - \iota_{hk}}{R_B}\right) + \lambda F_h^N\left(\frac{x_{j,h}^F - \iota_{hk}}{R_S}\right) \\ + (1-\lambda)\mathfrak{F}_{hjk}^K \Bigg], \qquad F' = S\Pi. \end{aligned} \tag{55}$$

where $x_{j,h}^F = x_{j,h}$ and

$$\mathfrak{F}_{hjk}^K = \begin{cases} F_h^N\left(\dfrac{x_{j,h} + \kappa - \iota_{hk}}{R_S}\right), & a_j < 0,\ x_{j,h} + \kappa < x_h^\star, \\ F_h^{N,<}\left(\dfrac{x_h^\star - \iota_{hk}}{R_S}\right), & a_j < 0,\ x_{j,h} + \kappa \geq x_h^\star, \\ F_h^N\left(\dfrac{x_{j,h} - \iota_{hk}}{R_S}\right), & a_j \geq 0. \end{cases} \tag{56}$$

The second case in (56) is the entry cap: a non-waived household enters only *strictly* below the value-matching threshold, even if a deeper negative output node would otherwise imply a larger beginning-asset threshold. The left limit preserves this strict inequality when the beginning distribution has a point mass.

The CDF transition alone does not recover next period's borrower mass because the value at zero also contains the discrete mass at zero. We therefore evaluate the borrower-mass law directly at the output threshold $a' = 0$. It adds incumbent borrowers below the limiting zero threshold, waived non-borrowers below the unrestricted threshold, and non-waived non-borrowers below the entry threshold. Before applying Π,

$$\begin{aligned} \widetilde{m}_h^{B\prime} = \sum_k w_k \Bigg[F_h^{B,<}\left(\frac{x_{0,h} - \iota_{hk}}{R_B}\right) + \lambda F_h^{N,<}\left(\frac{x_{0,h} - \iota_{hk}}{R_S}\right) \\ + (1-\lambda) F_h^{N,<}\left(\frac{x_h^\star - \iota_{hk}}{R_S}\right) \Bigg], \qquad m^{B\prime} = \widetilde{m}^{B\prime}\Pi. \end{aligned} \tag{57}$$

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

$$D\mathcal{T} = \begin{bmatrix} \partial F'/\partial F & \partial F'/\partial m^B \\ \partial m^{B\prime}/\partial F & \partial m^{B\prime}/\partial m^B \end{bmatrix}. \tag{58}$$

The off-diagonal blocks reflect the economics rather than numerical error: changing the left limit changes both side probabilities, and changing the stored CDF values changes next period's borrower mass. We remove only artificial derivatives caused by clipping after interpolation.

Negative assets are aggregated consistently from the borrower-side interpolant,

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

With wedge ρ_0, intermediation costs are $-\rho_0 A^- > 0$; entry costs add κ times the paid-entry flow. Aggregate consumption is therefore an explicit control rather than being inferred from the debt-free identity $Y - G$. This matters for the wage Phillips curve even when resource-cost terms cancel between the household budget and goods-market-clearing equations.

C.5 Macro Block and Numerical Implementation

We now state the aggregate block and the numerical choices used in all reported solutions. The macro block follows the sticky-wage HANK economy in Bianchi and Kaplan (2026). Hours equal output, $N_t = Y_t$. Consumption utility is CRRA, $u'(C) = C^{-\sigma}$, with $\sigma = 1$ in the log-utility calibration. Labor disutility satisfies $v'(N) = \chi N^{1/\nu}$, with $\nu = 1$ and χ chosen so that steady-state inflation is zero. The coefficient $\bar{\beta} = \sum_h p_h \beta_h$ is the population-weighted mean discount factor. The wage Phillips curve is

$$\pi_t = \kappa_w \left(v'(Y_t) - (1-\tau)u'(C_t)\right) + \bar{\beta}\,\pi_{t+1}, \qquad \kappa_w = 0.01.$$

Writing the Phillips curve in terms of aggregate consumption is important in the debt models: G denotes government spending, and goods used for intermediation and entry imply $C_t \neq Y_t - G$. With log utility, $u'(C_t) = 1/C_t$. This makes the distribution relevant for inflation as well as for goods-market clearing. Consumption enters the wage Phillips curve through the marginal rate of substitution between consumption and labor, while intermediation and debt-entry costs make aggregate consumption depend on the negative-wealth distribution and the entry flow.

This dependence does not require adding the distribution to the nonlinear fixed-point unknowns. Conditional on the inherited distribution and a candidate path for the aggregate variables, the DEGM transition directly returns the next distribution, aggregate saving, negative assets, and entry costs. The Phillips residual can therefore use the consumption aggregate generated by that transition. The outer iteration remains over inflation—one scalar per date in the nonlinear window—rather than over the high-dimensional distribution. Within each date, goods clearing has the strictly increasing one-dimensional residual described in Proposition 2. This separation between an exact distribution update conditional on aggregate variables and a low-dimensional iteration over aggregate prices is a key advantage of the DEGM representation.

The Taylor rule and Fisher equation give nominal and real rates,

$$i_t = \bar{r} + \phi_\pi \pi_t, \qquad \phi_\pi = 1.5, \qquad 1 + r_t = \frac{1 + i_t}{1 + \pi_{t+1}}.$$

The government budget identity with the fiscal rule gives the transfer and debt paths,

$$T_1 = \bar{T} + \Delta, \qquad T_t = \bar{T} - \phi_B \bar{r}(B_{t-1} - \bar{B}), \quad t \geq 2, \qquad \phi_B = 8,$$

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

State-space implementation. The asset grid has $N_A^+ = 100$ points on $[0, 4000]$, uniform in $\log(1 + \log(1 + a))$ ("log-log"), which concentrates nodes where most households are located. The debt models add $N_A^- = 20$ points on $[\underline{a}, 0)$, uniform in the mirrored transformation $\log(1 +$

$\log(1 + |a|)$), with $\underline{a} = -1.0$. Thus $N_A = N_A^+ + N_A^-$.

The combined idiosyncratic state h has $N_H = 5 \times 11 = 55$ values. The state vector contains the interior CDF coordinates, the borrower shares in the debt models, the lagged rate, and the shock states. We compress the policy control using a discrete cosine transform (DCT), retaining a small share of coefficients at the baseline tolerance. In the model with debt-entry costs, $2N_H$ additional controls represent the entry problem.

The CDF interpolant is monotone cubic and uses a Fritsch–Carlson construction with a smooth adjustment where neighboring secant slopes change sign. This preserves differentiability when automatic differentiation constructs the system Jacobians. In the debt models, separate interpolants are used for the negative and nonnegative parts of the wealth distribution because there is a mass point at zero; the borrower share has its own law of motion. The nonlinear benchmark uses the same DEGM distribution update and a safeguarded Newton method initialized from the linearized model path. Its residuals at finite N_A are around 10^{-6}.

C.6 Differentiability and Validation

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

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

NDU uses this same nonlinear operator, which treats negative and nonnegative assets separately, during its exact-update window. In every window quarter it constructs x^* from the first-order policy, J, and M^{0^-} controls through (49); it then updates both F and m^B and includes the paid-entry flow in resources. The resulting distribution—including the discrete mass at zero and the borrower mass—is passed to the linear continuation. No additional approximation is introduced for debt entry beyond the finite exact-update window itself.

D The Sequence-Space Savings Correction

In sequence space (Auclert et al., 2021), the distributional correction enters as a one-period savings correction. This gives a simple impact calculation, but it does not carry the corrected distribution forward. That limitation motivates the state-space design in the main text.

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, and the household block maps them into aggregate consumption $C(Y, \Delta)$. Goods-market clearing is the residual $\mathcal{R}(Y, \Delta) = C(Y, \Delta) - Y$. 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 $(J_Y - I)(Y - \bar{Y}) + J_\Delta \Delta = 0$.

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. Let $\tilde{A}_1(Y_1, \Delta)$ be the resulting end-of-period saving, $\tilde{A}_1^{lin}$ its first-order expansion, and

$$\gamma_A(Y_1, \Delta) = \tilde{A}_1(Y_1, \Delta) - \tilde{A}_1^{lin}(Y_1, \Delta) \tag{59}$$

the saving gap. It enters consumption with the opposite sign, so the corrected residual is

$$(J_Y - I)(Y - \bar{Y}) + J_\Delta \Delta - e_1 \gamma_A(Y_1, \Delta) = 0. \tag{60}$$

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

$$Y_1 = v_1(\Delta) + w_1 \gamma_A(Y_1, \Delta), \tag{61}$$

which is solved by safeguarded Newton/bisection without re-solving the household problem.

Figures D.1 and D.2 compare this correction with NDU(6) and the nonlinear benchmark at transfers of 1, 5, 10, and 15 percent of annual GDP. Figure D.1 removes the entry cost, while Figure D.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 corrected.

Without the entry cost, Figure D.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 D.2 shows that it remains close at transfers of 1 and 5 percent of annual GDP but departs visibly at 10 and 15 percent. 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.

The state-space method retains that distribution as a state and can continue to update it exactly. The difference is therefore not the DEGM formula itself: both methods use the same first-period distribution update. The difference is what happens afterward. In sequence space, the update enters as a one-period correction to market clearing. In state space, it becomes the law of motion for an explicit distributional state, allowing exact updates for several quarters.

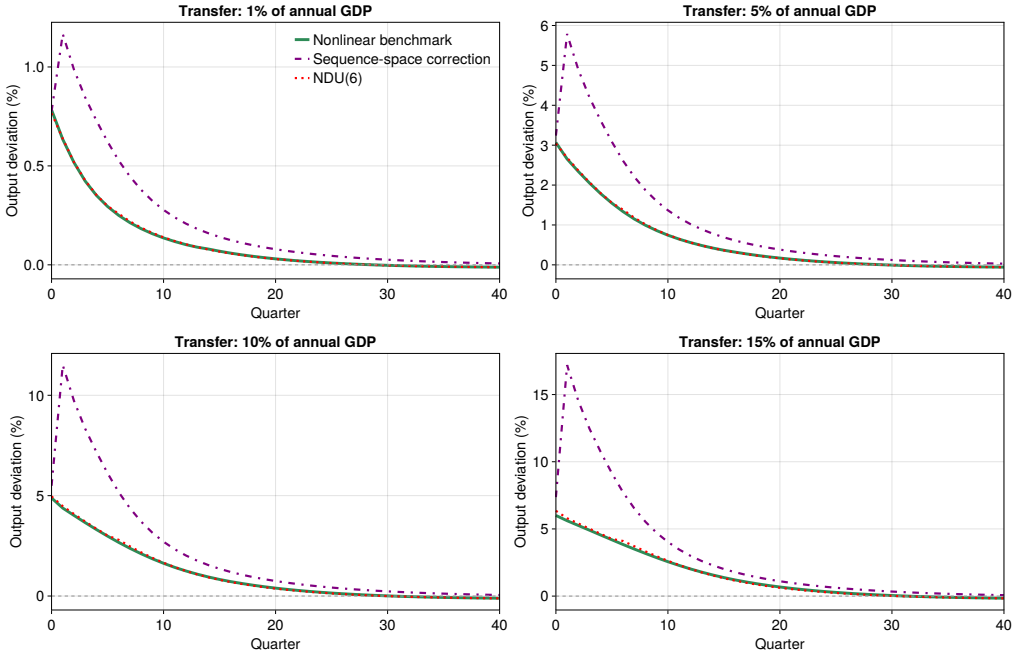


Figure D.1: **Output responses without a fixed debt-entry cost ($\kappa = 0$).**

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 (61), and red dotted is NDU(6). Output is measured as a percentage deviation from steady state.

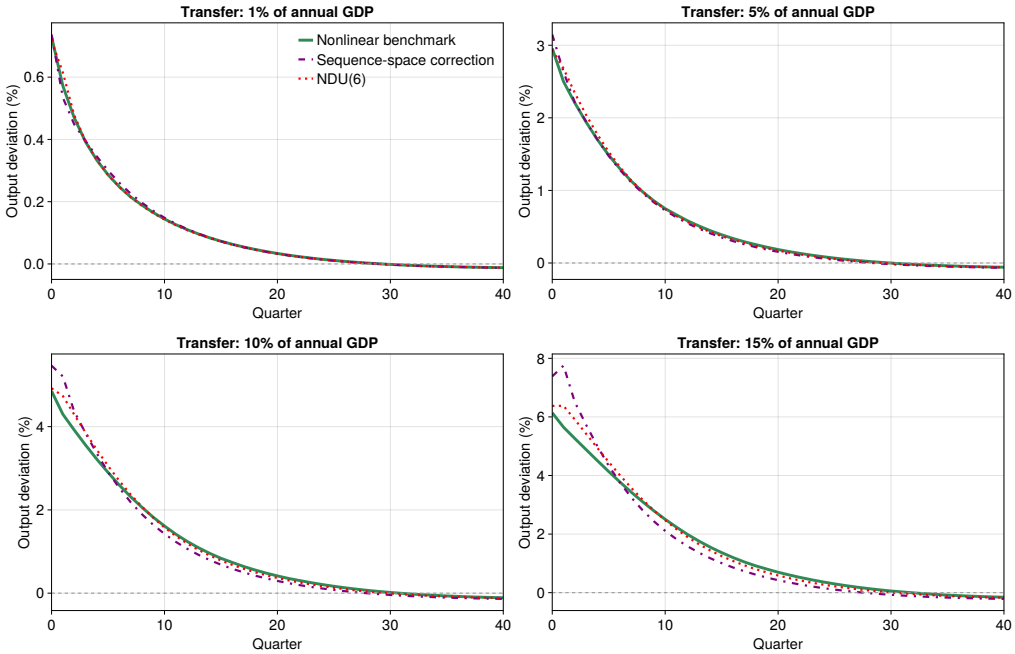


Figure D.2: **Output responses with a fixed debt-entry cost ($\kappa > 0$).**

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 (61), and red dotted is NDU(6). 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.

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.

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.

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.