Local Methods for Large Transfers*

Christian Bayer† Luigi-Maria Briglia‡

Ralph Luetticke§ Maximilian Weiss¶

Yannik Winkelmann‖

July 16, 2026

Abstract

Bianchi and Kaplan (2026) show that linear methods can greatly overstate the effects of large fiscal transfers. The central object is the shift in the asset distribution after a transfer shock, particularly the fraction of households that cross the borrowing-constraint threshold. This is precisely what the distributional endogenous-grid method (DEGM) developed in Bayer et al. (2026) captures. We show that DEGM naturally extends to capture the nonlinear shift in the asset distribution with respect to transfer size. Our Nonlinear DEGM Update (NDU) iterates the distribution exactly for the first period or a short window of periods while maintaining first-order household policies. With a 10 percent transfer, impact output increases by 7 percent with the nonlinear solution, 20 percent with the linear solution, and 7 percent with our novel method at near-linear computational cost.

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

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

†Universität Bonn, CEPR, and IZA, *email:* christian.bayer@uni-bonn.de

‡Universität Tübingen, *email:* luigi.briglia@uni-tuebingen.de

§Universität Tübingen and CEPR, *email:* ralph.luetticke@uni-tuebingen.de

¶Universität Tübingen, *email:* maximilian.weiss@uni-tuebingen.de

‖Universität Tübingen, *email:* yannik.winkelmann@uni-tuebingen.de

1 Introduction

Bianchi and Kaplan (2026) provide an important warning for the heterogeneous-agent literature. In models calibrated to match high and dispersed marginal propensities to consume (MPCs), the aggregate response to a funded fiscal transfer is strongly concave in the size of the transfer. Linear methods therefore overstate the efficacy of transfers. The source of this nonlinearity is not policy-function curvature per se, but how the asset distribution moves with the transfer size. Aggregate consumption is an integral over a distribution, and a transfer moves the borrowing-constraint threshold through that distribution. The relevant nonlinear term is the change in the fraction of constrained households: the MPC drop at the borrowing-constraint threshold times the mass of households near that threshold.

The distributional endogenous-grid method (DEGM) of Bayer et al. (2026) provides a representation of the distribution function that captures how it shifts with transfers, including the constraint-crossing share. DEGM writes the Kolmogorov-forward equation (KFE) as a distribution-function update at the cash-on-hand points that map into each asset gridpoint. A transfer shifts cash-on-hand. Holding policies fixed, the DEGM update immediately gives the distribution of end-of-period assets, including the mass that chooses the borrowing constraint.

Using this insight, DEGM gives a fully nonlinear and differentiable distribution update based on the pre-images 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 Young method (Young, 2010), which represents continuous choices as lotteries over a discrete state space and misses nonlinear terms. This makes DEGM well suited to address the concern that Bianchi and Kaplan (2026) raise. Bayer et al. (2026) use DEGM to compute a third-order perturbation solution. Here we first derive the importance of the constraint-crossing term for 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 the second-round income effects in the initial period(s), again with the nonlinear DEGM update. Inside each updated period, output is found from a scalar fixed point, but the rest of the computation remains the linear state-space solution. We call the resulting method the Nonlinear DEGM Update (NDU). We show that 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 (we report this version in Appendix C). In state space, the distribution remains in the system. The linear distribution update can then be replaced by the exact DEGM update for one, or a few, periods, and the correction affects the whole future path. Updating a few more periods exactly further improves the non-impact part of the path.

Quantitatively, our novel method fully captures the nonlinear response documented by Bianchi and Kaplan (2026), at a cost close to that of a linear solution. With a transfer of 10 percent of annual GDP, the nonlinear first-quarter output response is 6.96 percent, the linear response is 20.05 percent, and our Nonlinear DEGM update yields 7.01 percent. Updating only the distribution of the first period matches the impact response, though there are small deviations in later periods. Extending the window by a few periods eliminates these deviations.

We also highlight that the central object in the Bianchi-Kaplan critique is the fraction of households crossing the borrowing-constraint threshold as a function of transfer size. The more curved this constraint-crossing curve is, the more nonlinear is the aggregate effect of a transfer. In the Bianchi-Kaplan calibration, the curve is very steep and very curved. In their model, about 44 percent of households have less than two weeks of income in liquid savings. A transfer of 1 percent of annual GDP moves 13 percentage points of households across this liquidity threshold; a 5 percent transfer moves 42 percentage points. Given this observation, we move to the data and measure the crossing curve as an empirical object. In the SCF, the crossing curve is initially similarly steep but less curved and therefore flattens out less quickly, reaching a crossing share of 40 percentage points only at a transfer of 15 percent.

The paper proceeds as follows. Sections 2 and 3 spell out the distributional argument: the CDF form of the KFE update, the constraint-crossing terms generated by the moved cutoff, and a minimal piecewise-linear example. Section 3.2 reports the partial equilibrium results. Section 4 shows the state-space representation and Section 5 develops NDU. Section 6 reports the quantitative results and computational cost. Appendix C contains the sequence-space version.

2 The Distribution-Function Update

Consider the first period of a transfer experiment. Let x denote beginning-of-period cash-on-hand before the transfer. Let $\theta \in \Theta$ collect the idiosyncratic income component and other household types, and let $\mu(d\theta)$ be the population measure over these states. Conditional on θ, $F_\theta(q)$ is the share of households with cash-on-hand below q, and f_θ is the corresponding density away from mass points. A transfer of size Δ shifts cash-on-hand to

$$x_\Delta = x + \Delta. \tag{1}$$

The conditioning on θ is important. A conditional distribution can have mass points, for example from households at the borrowing limit. The aggregate object integrates over the type-income measure μ. As in Bayer et al. (2026), a continuous income component mixes these conditional distributions, so a cutoff that hits a mass point for one type is a measure-zero event in the aggregate integral.

Let $g_\theta(x)$ denote the end-of-period asset policy. Normalize the borrowing constraint to $\underline{a}$. The policy has the form

$$g_\theta(x) = \begin{cases} \underline{a}, & x \leq \bar{x}_\theta, \\ g^u_\theta(x), & x > \bar{x}_\theta, \end{cases} \tag{2}$$

where $\bar{x}_\theta$ is the cash-on-hand threshold at which the borrowing constraint stops binding, and g^u_θ is strictly increasing on the slack branch. We write $g^u_{\theta,x}$ for the derivative of the slack asset policy with respect to cash-on-hand. Consumption is $c_\theta(x) = x - g_\theta(x)$, and the unconstrained MPC is

$$m^u_\theta(x) = 1 - g^u_{\theta,x}(x). \tag{3}$$

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

$$\widetilde{F}_{\theta,\Delta}(a') = \Pr\left(g_\theta(x+\Delta) \le a'\right). \tag{4}$$

For asset values above the borrowing limit, monotonicity of g_θ^u lets us evaluate this distribution by going back to the cash-on-hand level that would choose a':

$$\widetilde{F}_{\theta,\Delta}(a') = F_\theta\left((g_\theta^u)^{-1}(a') - \Delta\right). \tag{5}$$

At the borrowing constraint, the asset value is fixed at $\underline{a}$, but the mass assigned to that value changes with the transfer:

$$\boxed{\widetilde{F}_{\theta,\Delta}(\underline{a}) = F_\theta\left(\bar{x}_\theta - \Delta\right).} \tag{6}$$

Equation (6) is the key step. As shown in Bayer et al. (2026), this is the DEGM forward step: the KFE is represented by the graph of the post-policy CDF at the cash-on-hand points that map into the asset grid. In this representation, the graph contains the borrowing-limit point

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

and the unconstrained-branch points

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

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, so Taylor expansions apply to this object. Second, the update is *closed-form*: evaluating the exact nonlinear transition for any (Y, Δ) is a single pass of CDF interpolations at these pre-images, no more expensive than the linear transition itself. This is what makes NDU(K) cheap.

3 Constraint-Crossing Terms

Aggregate first-period consumption can be written as

$$C(\Delta) = \int_\Theta \int c_\theta(x+\Delta)\, dF_\theta(x)\, \mu(d\theta). \tag{7}$$

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

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

Thus aggregate derivatives include derivatives of the updated distribution of saving, not only derivatives of individual consumption policies.

Proposition 1 (Constraint-crossing terms)**.** *Let $C_\theta(\Delta)$ denote the conditional type-θ contribution to* (7)*. 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, and suppose g^u_θ is sufficiently smooth on the slack branch. The superscript + denotes the right-hand limit from the slack side of the kink. Then, for regular θ,*

$$C'_\theta(\Delta) = \left[F_\theta(\bar{x}_\theta - \Delta) + \int_{\bar{x}_\theta-\Delta}^{\infty} m^u_\theta(x+\Delta) f_\theta(x)\, dx\right], \tag{9}$$

and

$$C''_\theta(\Delta) = \left[-\left(1 - m^u_\theta(\bar{x}^+_\theta)\right) f_\theta(\bar{x}_\theta - \Delta) + \int_{\bar{x}_\theta-\Delta}^{\infty} m^u_{\theta,x}(x+\Delta) f_\theta(x)\, dx\right]. \tag{10}$$

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:

$$C_\theta(\Delta) = \int_{-\infty}^{q_\theta(\Delta)} (x + \Delta - \underline{a})\, dF_\theta(x) + \int_{q_\theta(\Delta)}^{\infty} c^u_\theta(x+\Delta)\, dF_\theta(x). \tag{11}$$

The first integral is the constrained branch; the second is the slack branch. When we differentiate, the moving lower and upper limits produce boundary terms. Those terms cancel in the first derivative because consumption is continuous at the kink. What remains is the constrained mass times the constrained MPC, plus the average slack MPC, which gives (9). Differentiating (9) with Leibniz' rule gives (10). The new boundary term is the MPC jump at the kink, $1 - m^u_\theta(\bar{x}^+_\theta)$, 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 the borrowing constraint is therefore

$$-\left(1 - m^u_\theta(\bar{x}^+_\theta)\right) f_\theta(\bar{x}_\theta - \Delta). \tag{12}$$

It is generally nonzero even though the set exactly at the kink has measure zero. The reason is that the cutoff itself moves through a positive density of households. This is the precise sense in which aggregate consumption can be differentiable even though the individual policy has a kink.

Remark 1. *For a fixed household that remains constrained after an infinitesimal transfer, all higher derivatives of consumption are zero. This pointwise fact does not imply that aggregate higher-order derivatives are zero. The set of constrained households changes with Δ, and* (6) *differentiates that set.*

Figure 1 shows why the Bianchi-Kaplan calibration is a demanding case for any low-order expansion of the constraint-crossing share. Because the object in (12) has no direct counterpart in the data, the figure shows the crossing of a half-month-of-income to liquid-wealth threshold $\bar{x}$ in both model and data, following Kaplan and Violante (2014). Panel (a) plots the liquid-wealth distribution around the threshold, $F(\bar{x} + \delta) - F(\bar{x})$, that generates the crossing share; Panel (b) plots the crossing share. In the model the crossing share is already 5 percent at

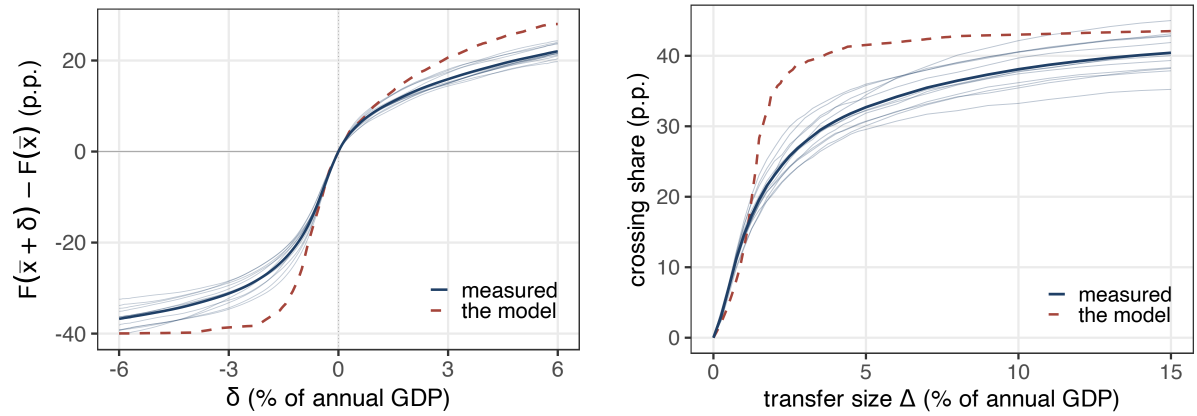

(a) Mass near the threshold

(b) Crossing share, consumption-corrected

Figure 1: **The constraint-crossing share in the SCF, read through the model.** The model crossing share is computed directly from the $N_A = 100$ steady-state DEGM CDF. The data use the SCF liquid-wealth distribution (Aladangady et al., 2023) around a household-specific half-month-income threshold, with transfer sizes expressed as percent of annual GDP per household.

a transfer of 0.5 percent of annual GDP, rises to 35 percent at 2 percent, and then flattens, approaching 44.5 percent by 20 percent. The steepness and curvature come directly from the shape in Panel (a): the model piles its constrained mass tightly against the borrowing limit (the near-vertical segments of the model curve), so a transfer that reaches this concentration sweeps a large fraction across at once, after which the curve flattens.[1]

The SCF crossing curve—wave-averaged over 1989–2022 and read through the model's consumption policy, as Appendix E details—has the same order of magnitude but a different shape, and the two curves cross. At small transfers the measured share is *higher*: 7.05 versus 4.78 percent at $\Delta = 0.5$ percent of annual GDP, and 14.59 versus 12.71 percent at $\Delta = 1$, because the data place more mass immediately below the threshold than the model does (Panel (a)). As the transfer grows the model overtakes: by $\Delta = 2$ it has swept its concentrated near-limit mass across (34.97 percent) while the smoothly spread empirical mass has moved only 23.07 percent. Finally, the two approach but do not meet—the measured curve flattens near 40 percent, about four percentage points below the model's $\sim$44 percent. The gap is the 17.6 percent of SCF households with *negative* liquid wealth, i.e. net revolving debt, which the model cannot represent because its borrowing limit is zero. The mapping assigns them the model's MPC at the borrowing limit (Appendix E); starting well below the threshold, roughly four percent of all households never accumulate enough to cross even at a 20 percent transfer, whereas the model's constrained mass sits at the limit itself, entirely within a large transfer's reach.

The data show a comparably steep curve at the transfer sizes where low-order expansions first run into trouble. But its exact shape, and its response to actual transfers, remain empirical questions. The methods below do not expand the crossing share at all: they evaluate the CDF update (6)–(5) exactly, at the realized shock.

[1]The exact constraint-crossing share in the model—the object in (12), measured at $\bar{x}_\theta$ rather than at the half-month threshold—is even more curved.

3.1 A Minimal Example of Distributional Nonlinearity

The distributional origin of the higher-order terms is transparent when every individual policy is piecewise linear. Set $\omega_\theta = 1$ and suppose

$$g_\theta(x) = \begin{cases} \underline{a}, & x \le \bar{x}_\theta, \\ \underline{a} + s_\theta(x - \bar{x}_\theta), & x > \bar{x}_\theta, \end{cases} \qquad 0 < s_\theta < 1. \tag{13}$$

The slack MPC is constant, $m_\theta = 1 - s_\theta$. Each household has a linear consumption policy on either side of the kink. Nevertheless the aggregate response is

$$C_\theta(\Delta) - C_\theta(0) = m_\theta \Delta + (1 - m_\theta) \int_0^\Delta F_\theta(\bar{x}_\theta - u)\, du. \tag{14}$$

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

$$\begin{aligned} C_\theta(\Delta) - C_\theta(0) = {} & [m_\theta + (1 - m_\theta) F_\theta(\bar{x}_\theta)]\, \Delta \\ & - \frac{1}{2}(1 - m_\theta) f_\theta(\bar{x}_\theta)\Delta^2 + \frac{1}{6}(1 - m_\theta) f_\theta'(\bar{x}_\theta)\Delta^3 + O(\Delta^4). \end{aligned} \tag{15}$$

Thus even with piecewise-linear individual policies, aggregate consumption has nonzero higher-order terms whenever the distribution has density or density slope at the moving threshold. The higher-order terms are not policy curvature; they are derivatives of the CDF. With heterogeneous types, the aggregate expansion is the integral of (15) over $\mu(d\theta)$.

3.2 A Partial-Equilibrium Test

Before adding any general-equilibrium feedback, we test the distribution update in the fixed-price experiment that Bianchi and Kaplan (2026) use to isolate the mechanism. Prices are held at their steady-state values and a one-time transfer of size Δ arrives unanticipated. Because nothing about the household's continuation problem changes—prices are fixed forever and the transfer is one-time—the impact policies are *exactly* the steady-state policies c_θ and g_θ. The entire impact response is therefore distributional, and the two expressions for aggregate consumption from the beginning of this section can both be evaluated without any approximation:

$$C(\Delta) = \int_\Theta \int c_\theta(x + \Delta)\, dF_\theta(x)\, \mu(d\theta) \tag{16}$$

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

$$= \mathcal{X}(\Delta) - \int_\Theta \int a'\, d\widetilde{F}_{\theta,\Delta}(a')\, \mu(d\theta), \tag{17}$$

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 is (7), and it is how the nonlinear baseline computes the response: integrate the nonlinear consumption function, evaluated at shifted cash-on-hand, over the stationary distribution. The second line substitutes the budget identity $c_\theta(x) = x - g_\theta(x)$. The third line is the change of variables $a' = g_\theta(x + \Delta)$: the distribution of end-of-period assets is exactly the shifted distribution $\widetilde{F}_{\theta,\Delta}$ delivered by the closed-form update (5)–(6), so this line recovers (8). This is

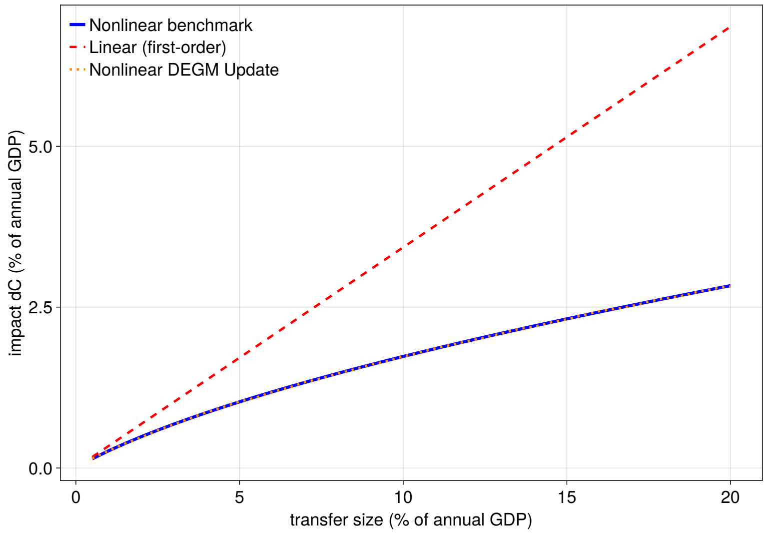

Figure 2: Partial equilibrium: impact consumption response to a one-time transfer

Notes: Fixed prices, steady-state policies. Blue: direct integration (16)—the steady-state consumption policy at shifted cash-on-hand, integrated over the stationary distribution. Orange dotted: the change-of-variables form (17)—the steady-state savings policy integrated over the shifted distribution (6)–(5), with consumption backed out of the resource constraint. Red dashed: first-order (policy-only) approximation, i.e. the aggregate marginal MPC of 0.343 times the transfer. Consumption response in percent of annual GDP; transfer size in percent of annual GDP. Baseline configuration ($N_A = 100$, log-log grid).

how the DEGM CDF update computes the response: integrate the steady-state savings policy over the shifted distribution, then back consumption out of the aggregate resource constraint. In exact arithmetic (16) and (17) are the same number. Numerically, however, they are built from different objects—(16) interpolates the consumption policy and holds the distribution fixed, while (17) interpolates the updated CDF on a fixed asset grid—so their agreement is a test of the numerical implementation.

Figure 2 reports the impact consumption response. The linear (first-order) response is the aggregate marginal propensity to consume times the transfer: the impact multiplier dC_1/Δ is constant at 0.34. The nonlinear response is strongly concave: the multiplier falls to 0.17 at a 10 percent transfer and 0.14 at 20 percent, so the linear approximation overstates the impact response by 94 percent at 10 percent of annual GDP and by 136 percent at 20 percent. The two equivalent integrals (16) and (17) are indistinguishable in the figure across the entire range, with a maximal gap of 0.02 percent of annual GDP: the discretized CDF update preserves the change of variables essentially exactly. Appendix B reports the corresponding table and shows that the second-order policy-only approximation (SOPS), which holds the constrained set fixed and therefore omits the constraint-crossing term (12), barely improves on the first-order one.

While the two aggregations (16) and (17) are equivalent, the one based on the shifted distribution has two advantages: First, the integral $\int a' d\tilde{F}_{\theta,\Delta}(a')$ is just the first moment of the shifted end-of-period distribution, which is an easier object to solve for than the integral of consumption over the distribution (we use integration by parts, see below). Second, the nonlinear consumption function in (16) is a household-level object; a linearized aggregate model has no state for it. The shifted distribution in (17), on the other hand, is an aggregate object which the state-space solution of Section 4 carries in its state vector. Therefore, the exactly-shifted distribution can be evaluated once and handed off to the linear model as a state.

4 A State-Space Solution

We solve the sticky-wage HANK economy of Bianchi and Kaplan (2026) (calibration in Appendix A) with a state-space method in which the cross-sectional distribution enters the state vector directly (Reiter, 2009). This choice is deliberate: with the distribution as a state, its linear update can be replaced by the exact update for the quarters in which the transfer creates a large movement.

States and controls. Let $h = (\text{discount-factor type}, z)$ index the combined permanent-type and persistent-income state, with transition matrix $\Pi_{hh'}$ and N_H values. The state vector x_t stacks

- the distribution coordinates X_t: deviations of the beginning-of-period conditional CDFs $F_{t,h}(a)$ from their steady-state values, at every interior asset grid node—the distribution at *full* grid resolution.[2]
- the lagged real rate $r_t^{lag} = r_{t-1} - \bar{r}$, 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 the consumption-policy deviations dc_t (one per asset node and state h) and the aggregates $(B_t, Y_t, \pi_t, i_t, r_t, T_t)$.

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. Timing: $F_{t,h}$ is the beginning-of-period CDF of assets carried into period t. A household in state h with transitory draw k and beginning assets a has cash-on-hand $x = R_t a + \iota_{h,k}(Y_t, T_t)$, where $\iota_{h,k}(Y,T) = (1-\tau)w_{h,k}Y + T$, and saves $a' = x - c(x)$. 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, and savings are monotone in x, so

$$a' \leq a_j \quad \iff \quad a \leq \frac{x_{j,h} - \iota_{h,k}}{R_t}. \tag{18}$$

The end-of-period savings CDF conditional on the current state is then a *closed-form formula*: the interpolated beginning-of-period CDF evaluated at the exact beginning-of-period asset level on the right-hand side of (18),

$$\tilde{F}_{j,h} = \sum_k w_k \, \hat{F}_{t,h}\left(\frac{x_{j,h} - \iota_{h,k}}{R_t}\right), \qquad F_{t+1,j,h'} = \sum_h \Pi_{hh'} \, \tilde{F}_{j,h} \;\equiv\; \mathcal{T}(F_t;\, x,\, \iota,\, R_t)_{j,h'}\,, \tag{19}$$

where $\hat{F}_{t,h}$ is a monotone-cubic interpolant of column h and w_k are the transitory-draw weights. Aggregate assets follow from the CDF by integration by parts, $A[F] = \sum_h \left(a_N F_{N,h} - \int_0^{a_N} F_h(a)\, da\right)$, and because Π is stochastic and the relabeling moves no assets, $A[F_{t+1}] = A[\tilde{F}_t]$.

[2] 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.

The stacked equilibrium system. An equilibrium path $\{x_t, y_t\}$ satisfies the system of equilibrium conditions

$$\mathbb{E}_t\, \mathcal{F}(x_t,\, y_t,\, x_{t+1},\, y_{t+1}) = 0, \tag{20}$$

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

$$\mathcal{F} = \begin{pmatrix} F_{t+1} - \mathcal{T}(F_t;\, \bar{x} + dc_t,\ \iota(Y_t, T_t),\ R_t) & \text{(distribution: KFE rows)} \\ dc_t - \mathcal{E}(dc_{t+1},\, Y_{t+1},\, T_{t+1},\, r_t) & \text{(household policy: Euler rows)} \\ Y_t - G - \Big[R_t\, A[F_t] + (1-\tau)\omega Y_t + T_t - A\big[\mathcal{T}(F_t; \cdot)\big]\Big] & \text{(aggregate savings row)} \\ \text{Phillips, Taylor, Fisher, fiscal rule, debt identity} & \text{(macro rows)} \end{pmatrix}, \tag{21}$$

where $\mathcal{E}$ is the household block's backward Euler operator (one EGM step off next period's policy, evaluated at next period's income and today's Euler rate). The third row deserves emphasis because it is the row the NDU scalar solve operates on. It is goods-market clearing written through the *aggregate-savings* side of the household budget identity: consumption is beginning-of-period resources minus end-of-period saving, and end-of-period saving is the asset aggregate of the transitioned distribution, $A[\mathcal{T}(\cdot)]$. The nonlinear transition operator $\mathcal{T}$ therefore enters the system in exactly two places: as the distribution's law of motion (first row) and inside the aggregate-savings condition (third row). Everything the change in the constrained share does to aggregates, it does through these two rows.

Linearization. The linear solution replaces every row of (21) by its first-order expansion around the steady state and solves the resulting linear system, giving the decision rule and law of motion

$$y_t = g_x x_t, \qquad x_{t+1} = h_x x_t. \tag{22}$$

In particular, the linearized aggregate-savings row reads

$$dY_t \ = \ \bar{R}\, A'[F_{ss}]\, dX_t + \bar{A}\, dr_t^{lag} + (1-\tau)\omega\, dY_t + dT_t \ - \ d\, A[\mathcal{T}], \tag{23}$$

where $d\, A[\mathcal{T}]$ is the first-order change in aggregate saving: the derivative of $A[\mathcal{T}(F;\, \bar{x} + dc,\, \iota,\, R)]$ with respect to each of its arguments, evaluated at the steady state and applied to the deviations $(dX_t,\, dc_t,\, dY_t,\, dr_t^{lag},\, dT_t)$. This is where the linear solution replaces the moved cutoff $F(\bar{x} - \Delta)$ by its tangent. The purely linear solution—(22) applied from $t = 1$ with $x_1 = \Delta\, e_\varepsilon$—reproduces the Bianchi-Kaplan critique: it overstates the 5 percent-transfer impact response by 126 percent (Table 1). The reason is precisely the tangent in (23), while the true curve flattens quickly (Figure 1).

5 The Nonlinear DEGM Update

NDU(1) 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 (21) and replaces two parts by their exact nonlinear counterparts:

- the *distribution rows*: instead of the linearized law of motion (the first block of (21) expanded to first order), the period-2 distribution is the exact image $F_2 = \mathcal{T}(F_{ss};\, \bar{x}+dc_1,\, \iota(Y_1, T_1),\, \bar{R})$;
- the *aggregate-savings row*: instead of the linearized goods-market condition (23), 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. 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(1) 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 pre-determined return $R_1 = \bar{R}$ is locked in, and only the shock state changes. The period-1 consumption-policy deviation is read off the linear decision rule,

$$dc_1 = g_x^{(c)} x_1 = \Delta \cdot g_x^{(c)} e_\varepsilon. \tag{24}$$

Because g_x solves the full forward-looking system, dc_1 contains the linear path's response to future prices. 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 $\bar{x} + dc_1$,

$$A_2(Y_1) = A\Big[\mathcal{T}\Big(F_{ss};\, \bar{x} + dc_1,\; \iota(Y_1, \bar{T} + \Delta),\; \bar{R}\Big)\Big]\,, \tag{25}$$

consumption follows from the aggregate budget identity, $C_1(Y_1) = \bar{R}\bar{A} + (1-\tau)\omega Y_1 + \bar{T} + \Delta - A_2(Y_1)$, and impact output solves the scalar goods-market condition

$$\boxed{Y_1 - G = C_1(Y_1).} \tag{26}$$

Equation (26) is the aggregate-savings row of (21), restored to its exact form: comparing with the linearized row (23), 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, in which the entire nonlinearity from the constraint-crossing share—the object Figure 1 shows a Taylor expansion cannot capture well in this calibration—is evaluated through $\mathcal{T}$.

Proposition 2 (Contraction). *Consider the function $\Phi(Y_1) = G + C_1(Y_1)$ with C_1 defined through* (25)*, thresholds and return fixed inside the loop. If the monotone-cubic interpolants $\hat{F}_h$*

are nondecreasing and the shifted consumption nodes remain nondecreasing, then

$$\Phi'(Y_1) \in [\,0,\ (1-\tau)\,\omega\,]\,,$$

so with $\tau > 0$, $\omega = 1$, Φ is a global contraction with modulus at most $1-\tau$: the fixed point of (26) *exists, is unique, and Picard iteration converges geometrically.*

Proof sketch. Income enters the transition only through the affine pre-images $u_{j,h,k}(Y_1) = (x_{j,h} - \iota_{h,k}(Y_1))/\bar{R}$, with $\partial u/\partial Y_1 = -(1-\tau)w_{h,k}/\bar{R}$. *Upper bound:* differentiating (19), $\partial G_{j,h}/\partial Y_1 = -\frac{1-\tau}{\bar{R}} \sum_k w_k w_{h,k} \hat{f}_h(u_{j,h,k}) \leq 0$ since $\hat{f}_h = \hat{F}'_h \geq 0$; higher income shifts the savings CDF right (first-order stochastic dominance), so $A'_2 \geq 0$ and $\Phi' \leq (1-\tau)\omega$. *Lower bound:* adjacent pre-images are at least $\Delta a_j/\bar{R}$ apart (using $x = c + a$ on the endogenous grid and monotone consumption nodes), so changing variables from asset nodes to pre-images, each column's density integrates to at most its mass, giving $A'_2 \leq (1-\tau)\omega$ and $\Phi' \geq 0$. □

By a change of variable, the derivative of the contraction is $\Phi'(Y_1) = (1-\tau)\,\mathbb{E}^*[\,w \cdot \mathrm{MPC}\,]$, the income-weighted aggregate MPC. Each iteration of the map Φ adds one round of induced spending; the fixed point sums the within-period Keynesian-cross multiplier.

Step 3: Prices and continuation. Given Y_1^*, the remaining period-1 prices come from the model's equations—the wage Phillips curve, Taylor rule, and Fisher equation—with only expected inflation π_2 read from the linear continuation at the continuation state. Because that state carries r_1 in its lagged-rate coordinate, (r_1, π_2) are pinned down jointly by a small, closed-form two-variable linear system.[3] Debt B_1 follows from the government budget identity. The replaced distribution row delivers the beginning-of-period-2 distribution as the exact image $F_2 = \mathcal{T}(F_{ss}; \bar{x} + dc_1, \iota(Y_1^*, \bar{T} + \Delta), \bar{R})$, and the state

$$x_2 = [\,F_2 - F_{ss};\ \ r_1 - \bar{r};\ \ 0\,]$$

is passed to the linear continuation (22) 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.

Table 1 reports the impact responses. The linear solution overstates the nonlinear response by 126 percent at a 5 percent transfer and by 188 percent at 10 percent, the same concavity documented by Bianchi and Kaplan (2026). NDU for one period removes essentially all of it: the impact gap is at most 6.3 basis points across all sizes through 10 percent of annual GDP.

[3]The 2x2 linear system reads

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

where $\widehat{g}$ is the wage-Phillips gap, the other parameters stem from the Phillips curve $\pi_1 = \kappa\widehat{g} + \bar{\beta}\,\pi_2$, the Taylor rule $i_1 = \bar{r} + \phi_\pi \pi_1$, and the (linearized) Fisher equation $r_1 - \bar{r} = (i_1 - \bar{r}) - \bar{R}\,\pi_2$, and $x_2(r_1)$ denotes the period-2 state, which carries r_1 in its lagged-rate coordinate and hence depends on r_1 (affinely).

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

Transfer size	Nonlinear	Linear	NDU(1)	NDU(6)
0.5% GDP	0.677	1.003	0.664	0.664
1.0% GDP	1.208	2.005	1.176	1.176
2.0% GDP	2.179	4.010	2.116	2.116
5.0% GDP	4.428	10.026	4.376	4.376
10.0% GDP	6.964	20.051	7.010	7.010

Notes: Sticky-wage HANK economy, state-space solution, $N_A = 100$, 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.

What the first period update of the distribution does *not* fix is visible in Figure 3: from quarter 2 on, the distribution shift is again treated linearly. At large shocks the path rises slightly above the benchmark over quarters 1–6, peaking at 0.26 (5 percent) to 0.55 (10 percent) percentage points, before the error vanishes again. This motivates NDU(K), for K periods.

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 $t = 1, \ldots, K$, run the same two-part solve as Section 5, but start from the previous quarter's solved objects rather than from the steady state:

1. *Inputs.* Beginning-of-period CDF F_t (with $F_1 = F_{ss}$); the gross return $R_t = \bar{R} + (r_{t-1} - \bar{r})$, locked in 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.* $dc_t = g_x^{(c)} x_t$ with $x_t = [\,F_t - F_{ss};\ r_{t-1} - \bar{r};\ \varepsilon_t\,]$: the linear decision rule, evaluated at the *exactly transitioned* distribution. This is the natural generalization of (24).

3. *Exact scalar clearing.* Solve Y_t from the goods-market condition with saving $A[\mathcal{T}(F_t; \bar{x} + dc_t, \iota(Y_t, T_t), R_t)]$. Proposition 2 uses the same argument; only the starting distribution changes. Each quarter is therefore the same contraction with modulus $\leq 1 - \tau$.

4. *Advance exactly.* $F_{t+1} = \mathcal{T}(F_t; \bar{x} + dc_t, \iota(Y_t^*, T_t), R_t)$; debt from the budget identity; r_t from the price closure; carry forward.

After quarter K the state $x_{K+1} = [\,F_{K+1} - F_{ss};\ r_K - \bar{r};\ 0\,]$ is passed to (22) 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 small K-vector by iteration. We start from the linear path, run the window once, replace each expected-inflation entry with the inflation rate just implied by the run, and repeat until the entries stop changing. Expected inflation at the window boundary, π_{K+1}, is read from the linear continuation at the continuation state. The number of runs needed is small (5 at $K = 2$, 13 at $K = 4$; Table 3).

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. In this model, that remaining error is quantitatively small.[4]

6 Quantitative Results

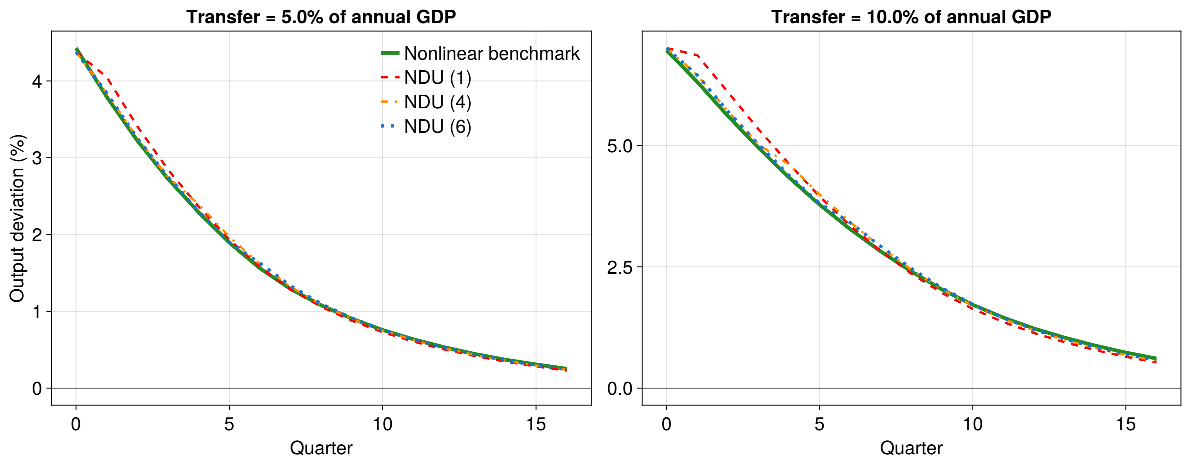

Figure 3: Output IRFs: nonlinear benchmark and Nonlinear DEGM Update

Notes: Sticky-wage HANK economy, state-space solution, $N_A = 100$, log-log grid, $T = 150$. Green: nonlinear DEGM benchmark. Red dashed: NDU(1), exact period 1 with linear continuation from quarter 2. Orange dash-dotted and blue dotted: NDU(4) and NDU(6). Output deviations in percent of steady-state output.

All state-space results use the baseline configuration: a 100-point log-log asset grid on $[0, 4000]$[5], plain (untransformed) CDF state coordinates, and the first-order window policy. The nonlinear benchmark is a fully nonlinear transition path solve (horizon $T = 150$) under the same DEGM distribution update. For consistency, the linear and nonlinear columns in the main tables use the same DEGM/CDF representation.

Figure 3 shows the transition paths at 5 and 10 percent transfers. NDU(1) (red dashed) matches on impact but is above the benchmark over quarters 1–6: with the distribution handed to the linear continuation after one quarter, the continuation applies steady-state (tangent) MPC behavior to a distribution still far from steady state, causing consumption to rise too early.

[4]One can also update household expectations inside the window. We leave this extension for future work.
[5]Bayer et al. (2026) show that DEGM converges much faster in the distribution than the lottery method.

Widening the NDU window applies the exact distributional transition in those quarters, and the path moves close to the benchmark.

Table 2 quantifies the remaining early-path error. It shows the maximum error over the first six quarters for different transfer sizes and NDU window lengths. Even NDU(1) nearly matches the fully nonlinear solution. For a transfer of 10 percent of GDP, the maximum error implies a difference of between 0.55 and 0.14 percentage points to the nonlinear solution as the window length increases.

Table 2: Peak early-path error by window length

	Window length K (exact quarters)						
Transfer size	1	2	3	4	**6**	8	12
1% GDP	0.087	0.039	0.024	0.016	**0.005**	0.005	0.005
5% GDP	0.257	0.188	0.131	0.102	**0.049**	0.049	0.049
10% GDP	0.551	0.442	0.340	0.266	**0.138**	0.138	0.138

Notes: Maximum over quarters 1–6 of the NDU(K) output deviation minus the nonlinear benchmark, in percentage points. Impact ($t = 1$ in the table's quarter-0 convention) is identical across K.

Table 3 reports the timing. NDU(6) costs about five seconds per shock size in the $N_A = 100$ run. It is an order of magnitude cheaper than nonlinear path solving, while matching the nonlinear path very well. This difference will be larger in a model with richer heterogeneity. The curse of dimensionality will affect the nonlinear solution here, while NDU only requires updating the distribution function and solving for a scalar.

Table 3: Cost of NDU(K) (5% transfer, $N_A = 100$)

	Window solve	Inflation iterations
NDU(1)	2.21 s	—
NDU(2)	2.46 s	5
NDU(4)	3.68 s	13
NDU(6)	5.20 s	20
Nonlinear benchmark (with linear solution as initial guess)	70.70 s	

Notes: Times from the $N_A = 100$ run. "Inflation iterations" counts window passes; 20 is the configured cap. Nonlinear times vary with Newton convergence and include cases that stop at the roughly 10^{-6} discretization-noise floor. Therefore, for the nonlinear benchmark we report the average over the times needed to solve for shock sizes 0.5%, 1%, 5%, and 10%.

7 Conclusion

Bianchi and Kaplan (2026) show that large fiscal transfers are a challenge for linear methods. 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 borrowing-constraint threshold through that distribution. The central object is therefore the fraction of households that cross the threshold. This is precisely the object captured by the DEGM representation of the KFE, which turns the nonlinear distribution movement into a closed-form update.

This update is easiest to use in a state-space solution. With the distribution as a state, NDU(K) replaces the linear transition by the exact DEGM update over a short window. In each quarter of the window, this only requires one CDF update and one scalar goods-market solve. In the Bianchi-Kaplan experiment, NDU matches the fully nonlinear solution on impact, tracks the transition path to about 0.14 percentage points, and costs seconds per experiment on top of a one-time linearization.

The crossing curve also gives the hand-to-mouth question (Aguiar et al., 2025) direct policy content. What matters for transfer size dependence is not only how many households are recorded at the constraint, but how much mass lies close enough to cross it and how large the MPC drop is when they cross. The SCF liquid-wealth distribution points in the same direction: a large share of households sits close enough to the threshold to cross after modest transfers.

A Calibration and Computation

A period is a quarter; quarterly GDP is normalized to one, so annual GDP is four. The annual real interest rate is two percent, the borrowing limit is zero, and 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 of 15 percent of GDP, as in Bianchi and Kaplan (2026).

The idiosyncratic income process follows Kaplan and Violante (2022): $\log y = z + \varepsilon$, where the persistent component updates with quarterly probability 0.25 to $\phi_z z + \eta$, $\phi_z = 0.9878$, $\sigma_\eta^2 = 0.0439$, and the transitory component is drawn each quarter, equal to zero with probability 0.75 and otherwise normal with variance $\sigma_\varepsilon^2 = 0.6376$. The persistent component is discretized with 11 Rouwenhorst states. The transitory component uses five Gauss-Hermite nodes, matching the second and fourth central moments of the normal component. Seven equally spaced discount factors are calibrated so that mean wealth equals annual GDP and the average quarterly MPC out of a \$1,000 windfall equals 0.25.

The macro block follows the sticky-wage HANK economy in Bianchi and Kaplan (2026). Hours equal output, $N_t = Y_t$. The wage Phillips curve is

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

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,

$$T_0 = \bar{T} + \Delta, \qquad T_t = \bar{T} - \phi_B \bar{r}(B_{t-1} - \bar{B}), \quad t \geq 1, \qquad \phi_B = 8,$$

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

State-space implementation. The asset grid for the state-space solution has $N_A = 100$ points on $[0, 4000]$, uniform in $\log(1 + \log(1 + a))$ ("log-log"), which concentrates nodes where most household mass lives. The combined idiosyncratic state h has $N_H = 7 \times 11 = 77$ values; the state vector stacks 77×99 interior CDF coordinates plus the lagged rate and shock states. The policy control is DCT-compressed, retaining 1,024 of 7,700 coefficients at the baseline tolerance. The monotone-cubic CDF interpolant uses a Fritsch-Carlson construction with a smooth gate at sign changes of the secant slopes, which keeps the transition differentiable for the automatic-differentiation pass that builds the system Jacobians. The nonlinear benchmark solves the full sequence-space transition under the identical DEGM transport with a safeguarded Newton method, warm-started from the linearized model path; its residual noise floor at finite N_A is around 10^{-6}.

B Partial-Equilibrium Constraint-Crossing Evidence

Table 4 complements the partial-equilibrium figure of Section 3.2 with policy-only approximations. Prices, continuation values, and household policies are held fixed at the steady state. The policy-only first- and second-order approximations (FOPS, SOPS; Appendix D) hold the constrained set fixed, so the constraint-crossing term (12) is absent: going to second order barely improves on first order. The DEGM CDF update evaluates the moved cutoff (6) directly and tracks the nonlinear response throughout.

Table 4: Partial-equilibrium impact consumption response

Transfer size	Nonlinear	FOPS	SOPS	DEGM CDF update
0.5% GDP	0.1430	0.1768	0.1761	0.1477
1.0% GDP	0.2676	0.3536	0.3508	0.2677
2.0% GDP	0.4879	0.7072	0.6960	0.4769
5.0% GDP	1.0285	1.7679	1.6979	1.0324
10.0% GDP	1.7334	3.5358	3.2557	1.7315
20.0% GDP	2.8334	7.0717	5.9511	2.8227

Notes: Responses in percent of annual GDP. FOPS and SOPS are policy-only Taylor approximations with the constrained set held fixed, computed by implicit differentiation of the Euler equation (Appendix D) on the same steady-state DEGM objects as the other columns. The DEGM CDF update pushes the steady-state distribution through the steady-state policy using the shifted distribution $\tilde{F}_\Delta$. Baseline configuration ($N_A = 100$, log-log grid), the same as Figure 2.

C 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 path of output $Y = (Y_1, \ldots, Y_T)$, $T = 150$. Given Y and Δ, the macro block computes prices and fiscal variables; the household block maps (Y, Δ) 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 correction is a nonlinear saving term. Starting from the steady-state distribution and policy, apply one DEGM forward step with period-1 cash-on-hand shifted by income Y_1 and the transfer Δ; let $\tilde{A}_1(Y_1, \Delta)$ be the implied aggregate end-of-period saving, $\tilde{A}_1^{lin}$ its first-order expansion, and

$$\gamma_A(Y_1, \Delta) = \tilde{A}_1(Y_1, \Delta) - \tilde{A}_1^{lin}(Y_1, \Delta) \tag{27}$$

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

$$(J_Y - I)(Y - \bar{Y}) + J_\Delta \Delta - e_1 \gamma_A(Y_1, \Delta) = 0. \tag{28}$$

Because the right-hand side depends on Y only through the scalar 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

$$Y_1 = v_1(\Delta) + w_1 \gamma_A(Y_1, \Delta), \tag{29}$$

solved by safeguarded Newton/bisection in 4–9 iterations (21 in one knife-edge case), residual below 10^{-9}, with no household re-solve.

Table 5: Sequence space: general-equilibrium first-quarter output response

Transfer size	Nonlinear	Linear	Sequence-space correction	NDU(1)
1.0% GDP	1.208	2.005	1.229	1.176
5.0% GDP	4.428	10.026	4.671	4.376
10.0% GDP	6.964	20.051	7.806	7.010

Notes: Sticky-wage HANK economy, horizon $T = 150$, baseline configuration ($N_A = 100$, log-log grid). Output responses in percent deviation from steady-state output. Linear is the linearization of the DEGM model. The sequence-space correction uses the additive savings correction in (28). NDU(1) is the impact response of the main text's method (identical for every window length K by construction).

Table 5 reports the impact result. The linear solution overstates the nonlinear impact response by 126 percent at a 5 percent transfer and by 188 percent at 10 percent; the sequence-space correction cuts this to 5.5 and 12.1 percent. The remaining gap has a structural source that no refinement of this formulation removes cheaply: in sequence space the distribution is substituted out of the system, so the one-period correction captures the nonlinearity of the period-1 update evaluated at the steady-state distribution, but the same borrowing-constraint clip keeps acting at dates $t \geq 2$ on distributions already shifted by the shock. Correcting those later dates would require carrying the shifted distribution itself—which is exactly what the state-space formulation does.

The comparison with the main text is the point of this appendix. On impact, the sequence-space correction reduces the 5-percent overshoot from 126 to 5.5 percent (at 10 percent: from 188 to 12.1), while NDU(1) brings the impact error down to 1.2 percent at 5 percent and 0.7 percent at 10. Only the state-space method can repair the transition path, because there the corrected distribution remains in the system as a state. The difference is not the DEGM update—both use the identical operator—but where it lives: as a residual patch on a market-clearing condition, or as the law of motion of an explicit distributional state whose equilibrium row can be replaced exactly.

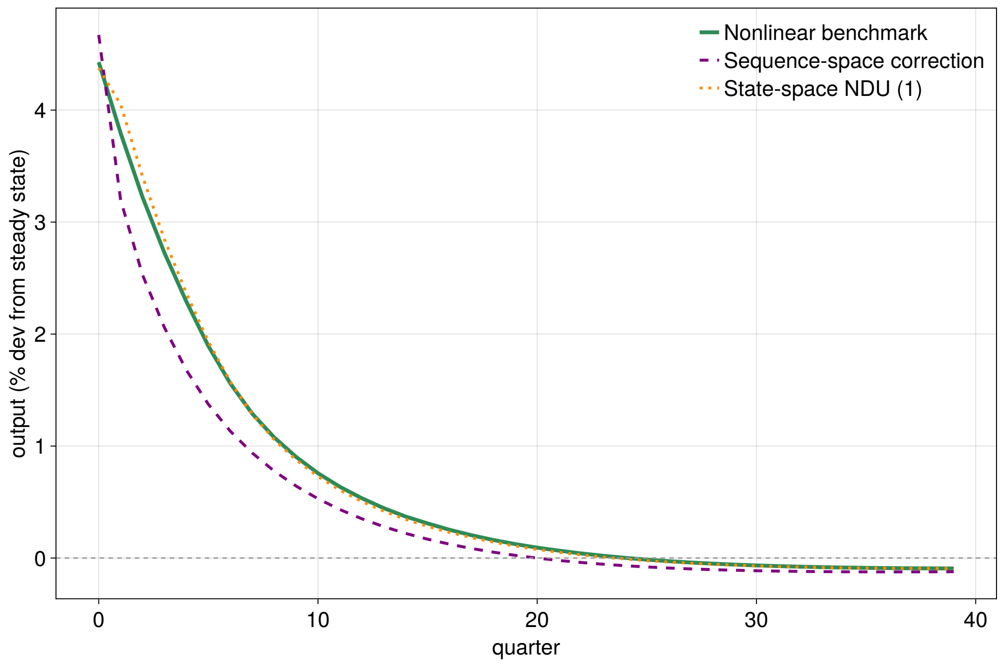

Figure 4: Output IRF: sequence-space correction versus NDU(1)

Notes: One-time transfer of 5 percent of annual GDP. Green: nonlinear benchmark. Purple dashed: the sequence-space correction (29). Orange dotted: NDU(1). Output in percent deviation from steady state; baseline configuration ($N_A = 100$, log-log grid).

D Policy-Only Approximations

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:

$$m(x,y) = \frac{\Lambda}{u''(c) + \Lambda}, \qquad \Lambda \equiv \beta R^2\, \mathbb{E}\left[u''(c')\, m' \mid y\right], \tag{30}$$

$$m_x(x,y) = \frac{[R(1-m)]^2 \Psi - u'''(c)\, m^2}{u''(c) + \Lambda}, \qquad \Psi \equiv \beta R\, \mathbb{E}\left[u'''(c')(m')^2 + u''(c')\, m'_x \mid y\right], \tag{31}$$

with $m = 1$ and $m_x = 0$ on the constrained branch. The aggregate coefficients are expectations over the fixed steady-state distribution. The constrained set is held fixed, so the boundary term in (12) is absent by construction.

E Reading the Crossing Curve through the Model

The empirical curve in Figure 1 is built to measure the same object as the model curve: the share of households that a transfer moves across the half-month-of-income threshold, once each household consumes the share of the transfer implied by the model's consumption policy. This appendix records the construction.

Data. 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 the survey weights.

Liquid wealth follows the Kaplan-Violante definition (Kaplan and Violante, 2014): checking, saving, money-market and call accounts; directly held stocks, bonds, mutual funds, and savings bonds; net of revolving credit-card debt. The threshold is household-specific and set to half a month of the household's own income, and transfers are expressed as a percent of annual GDP per household, so the units match the model's.

Mapping to the income state. The model's policy is indexed by a persistent income state z. We assign each household to a z by its rank in the *residualized* income distribution within its wave: we regress log income on a third-order polynomial in age and on indicators for education, marital status, number of children, race, and sex of the household head, and rank households by the residual against the cumulative stationary probabilities of the eleven Rouwenhorst states (Appendix A). Ranking on the residual rather than on raw income strips predictable life-cycle and demographic variation, so the rank proxies the model's idiosyncratic state. The choice matters little for the result, because the model's MPC is nearly invariant to z once wealth is expressed in months of the household's own income.

The consumption response and the crossing. Each household receives the transfer and consumes the share of it implied by the model's steady-state consumption policy $c(a, z)$, interpolated between its asset nodes with the same Fritsch-Carlson monotone-cubic construction the model uses internally (Appendix A). A household with liquid wealth a and transfer Δ consumes $c(a + \Delta, z) - c(a, z)$ of it and saves the rest; it counts as crossing if its liquid wealth was at or below the cutoff and, after saving, exceeds it. The model curve applies the same consumption correction—on neither side does the household save the whole transfer—which is what makes the two curves in Panel (b) of Figure 1 directly comparable.

Households outside the model's support. About 17.6 percent of SCF households (wave-averaged) hold *negative* liquid wealth, i.e. net revolving debt, which the model cannot represent because its borrowing limit is zero. We keep them in the sample and assign them the model's MPC at the borrowing limit (the value at $a = 0$), held flat below zero, so they save a fixed fraction of any transfer. Because they start well below the threshold, a subset never accumulates enough to cross even at the largest transfers; this is why the measured curve saturates about four percentage points below its own constrained share and below the model, whose constrained mass sits at the limit itself.

Discount-factor heterogeneity. The model has seven discount-factor types; the baseline curve weights them by their population shares. A variant that instead weights each household's type by the model's stationary type distribution conditional on assets and income is nearly indistinguishable from the baseline.

References

Aguiar, M., Bils, M., and Boar, C. (2025). Who are the hand-to-mouth? *Review of Economic Studies*, 92(3):1293–1340.

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.

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.

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., Weiss, M., and Winkelmann, Y. (2026). An endogenous gridpoint method for distributional dynamics. *Journal of Monetary Economics*, page 103895.

Bianchi, J. and Kaplan, G. (2026). How Small Is Small? Non-Linearities in Heterogeneous Agent Models. *NBER WP 35311*.

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

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

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

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.