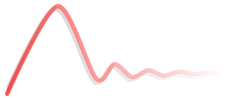

Dynare Working Papers Series
https://www.dynare.org/wp/

# Rich Heterogeneity in Dynare 7: A Practical Description

Normann Rion

Working Paper no. 88

July 2026

CEPREMAP
CENTRE POUR LA RECHERCHE ECONOMIQUE ET SES APPLICATIONS
48, boulevard Jourdan — 75014 Paris — France
https://www.cepremap.fr

# Rich Heterogeneity in Dynare 7: A Practical Description

Normann Rion[1,2]

[1]Dynare Team
[2]CY Cergy Paris Université (THEMA)

**Abstract**

Dynare 7 ships a toolkit for solving heterogeneous-agent models: a `.mod` syntax declaring a heterogeneity dimension alongside the standard aggregate blocks, commands that load or compute the stationary equilibrium, and a first-order solver whose output is the sequence-space Jacobian of the aggregate equilibrium path with respect to the aggregate innovations, the array `dr.G`. This document is the toolkit's mathematical companion, and it reads the methods it implements as one object: the state-space treatment of the heterogeneous block due to Bhandari et al. [2023], the sequence-space assembly of Auclert et al. [2021], and the sequence-space determinacy theory of Auclert et al. [2025], recast through the theory of linear time-invariant systems and Toeplitz operators. After stating the class of models, the syntax, and the steady-state algorithms, I develop the first-order theory behind the solver, the backward operator of the heterogeneous block, the transfer operator moving the cross-section, and their assembly into the sequence-space Jacobian, and read the state-space recursion of Bhandari et al. [2023] and the sequence-space Jacobian of Auclert et al. [2021] as two readings of one Fréchet derivative, with the borrowing-constraint atom priced by a weak-derivative and finite-element reading that never forces a derivative onto a Dirac at first order. The same continuous objects carry two extensions the toolkit does not yet ship. The first is existence and determinacy: the criterion of Auclert et al. [2025], established there on the discretized block, is shown to hold for the continuous operators, and the winding-number verdict read off the discretized symbol is certified to be the continuous model's by a finite-sample bound. The second is a perturbation cascade that carries the second-order heterogeneous construction of Bhandari et al. [2023] to arbitrary order in the sequence space, the one Wiener–Hopf factorization computed at the first order serving every order.

## Introduction

This document is the mathematical companion of the heterogeneous-agent toolkit shipped with Dynare 7. The toolkit combines the state-space treatment of the heterogeneous block due to Bhandari et al. [2023] with the sequence-space assembly of Auclert et al. [2021], and this document records, in one self-contained place, what the toolkit computes, under which assumptions, and by which algorithms. Where it goes past a record, and the polished results make it worth stating plainly, it is a mathematical reinterpretation: it reads the methods of Bhandari et al. [2023] and Auclert et al. [2021], together with the sequence-space determinacy theory of Auclert et al. [2025], as readings of one object, the Fréchet derivative of the equilibrium map at the steady state, and it links that object to concepts the theory of linear time-invariant systems and Toeplitz operators has long owned. Where an object it needs already has a name in another literature, I use that name and cite the source rather than rebuild it.

Three papers meet here, all from the literature that linearizes heterogeneous-agent models around a steady state, and this document rereads them through a second, mathematical literature: linear time-invariant systems and Toeplitz operators [Böttcher and Silbermann, 1999, Gohberg et al., 1990]. The heterogeneous block is that of Bhandari et al. [2023]: the household problem is differentiated through an implicit function theorem into a backward operator $\mathfrak{r}$, the cross-section moves under a transfer operator $\Lambda$, and the borrowing-constraint kink is disciplined so that, at first order, no generalized function ever appears; I recast their histogram analysis in a weak-derivative and finite-element language, which prices exactly what the $P^1$ histogram sees and misses. The aggregate assembly is that of Auclert et al. [2021]: the equilibrium is closed in the sequence space, and the object the solver returns, the array `dr.G`, is their sequence-space Jacobian read as the derivative of the aggregate path with respect to the innovations. Read together, the two are one Fréchet derivative in two coordinate systems, the equivalence Bhandari et al. [2023] establish, here obtained from the operator side.

The mathematics then answers two questions the truncated arrays cannot answer from inside. The first is determinacy: the stacked equilibrium operator is quasi-Toeplitz, its symbol is a transfer function, and existence and uniqueness are read from a winding number, the criterion Auclert et al. [2025] give for the discretized block and this document establishes for the continuous operators, with the winding number a discretized solve returns certified to be the continuous model's. The second is the order of approximation: the same operator, factored once at the first order, closes every higher order of the perturbation, the order-$k$ response solving a linear system in that one factorization with a right-hand side built from the orders below, so the second-order heterogeneous construction of Bhandari et al. [2023] extends to arbitrary order in the sequence space, the order-$k$ statement, the symbol form, and the constraint-layer corrections being what the reading adds.

Throughout, one worked example carries every abstract object: the Krusell and Smith [1998] economy, in the calibration of the `sequence-jacobian` toolkit of Auclert et al. [2021], with total factor productivity following an AR(1) in deviation $Z_t$. Its `.mod` file, its steady-state construction, and its solved dynamics appear in full, so that every mathematical statement can be matched against a file the reader can run.

The document is organized as follows. Section 1 states the class of models, two blocks $F$ and $G$ tied by cross-sectional integrals, and maps it onto the Dynare 7 syntax: the `heterogeneity_dimension` declaration, the `heterogeneity=` options, the `SUM` operator, and the $\perp$-clauses for occasionally binding constraints. Section 2 rewrites the class in recursive form and defines the steady state. Section 3 documents the two steady-state commands, the discretized objects they consume, and the algorithms behind `heterogeneity_compute_steady_state`, time iteration, forward iteration, and the calibration solve. Section 4 is the core: it perturbs the recursive equilibrium to first order, the backward operator and the policy response on the heterogeneous side, the transfer operator and its weak derivative on the distribution side, their assembly into the sequence-space Jacobian, and the finite-element reading of what `heterogeneity_solve` actually iterates, down to `dr.G` and the commands that consume it. Section 5 develops what the same objects buy beyond the solver: the quasi-Toeplitz structure, the winding-number test for existence and determinacy, and the order-$k$ cascade, shock kernels and risk corrections alike, all solved by the one factorization. The appendices collect the helpers, the worked `initial_guess` construction, a functional-analysis primer written for economists, the residual derivatives, the decay classes, a code dictionary, and the reconstruction of the cross-sectional responses.

The assumptions are stated explicitly, and the list of those not derived from primitives is short but not empty. The geometric-stability clauses of Assumption 2, the mixing of the household dynamics and the discounting of the backward map, are stated, not derived from primitives; so are the boundary clauses of Section 5.2 and the quantitative decay rates of Appendix E that the higher orders consume, although their qualitative frame is established in the text and their leading case is derived in that appendix. The determinacy test and the order-$k$ cascade are mathematics the framework delivers on arrays Dynare 7 already forms, but they are not yet in the code: the toolkit ships the first order.

Two audiences are intended, and they can read differently. A Dynare user who wants to write, solve, and simulate a heterogeneous-agent model needs Sections 1 through 3 and, in Section 4, the pipeline and simulation paragraphs around `dr.G`, and can stop there. A reader after the methods, why the constraint never forces a Dirac at first order, why the truncated Jacobian cannot certify determinacy but its symbol can, what an order-$k$ solve costs, will find the full arguments in Sections 4 and 5 and the appendices, written so that a PhD economist can check every step.

# 1 The problem and the Dynare 7 syntax to describe it

**Abstract framework.** In this section, I describe the class of heterogeneous-agent problems Dynare 7 is designed to solve. To do so, I essentially follow the notations in Bhandari et al. [2023] with a few departures. Consider a continuum of heterogeneous agents indexed by $i \in \mathcal{I}$ in a discrete-time environment. For the sake of clarity, symbols representing heterogeneity-related variables are lower-case red, and standard aggregate variables are upper-case blue. In period $t$, heterogeneous agents choose endogenous variables $x_{i,t}$ based on their time-$t$ information set, which includes pre-determined variables $a_{i,t-1} \subseteq x_{i,t-1}$ inherited from period $t-1$, exogenous shocks $\theta_{i,t}$, and the subset of aggregate variables $Y_t$ that are relevant to heterogeneous agents. Between vectors, the symbol $\subseteq$ reads entrywise: the left vector collects a subset of the right one's entries. Expectations are rational, and the agent-level optimality conditions take the form

$$F(\theta_{i,t}, a_{i,t-1}, x_{i,t}, \mathbb{E}_{i,t} x_{i,t+1}, Y_t) = 0, \quad \forall i \in \mathcal{I},\ t \geq 0.$$

In the same manner, the remaining aggregate equilibrium conditions can be written as

$$G\left(\int y_{i,t}\,di,\; A_{t-1},\; X_t,\; \mathbb{E}_t X_{t+1},\; \mathcal{E}_t\right) = 0, \quad \forall t \geq 0,$$

where $X_t$ are aggregate endogenous variables, $A_{t-1} \subseteq X_{t-1}$ are aggregate pre-determined variables, and $\mathcal{E}_t$ are aggregate i.i.d. Gaussian innovations. The subset of aggregate variables that is relevant to heterogeneous agents satisfies $Y_t \subseteq [A_{t-1}, X_t, \mathbb{E}_t X_{t+1}, \mathcal{E}_t]$. Similarly, $y_{i,t} \subseteq x_{i,t}$ denotes the subset of heterogeneous variables that enter $G$ through cross-sectional integrals. I further assume that each idiosyncratic process $\{\theta_{i,t}\}_{t\geq 0}$ is Markovian in time, that the shocks $\{\theta_{i,t}\}_{i\in\mathcal{I}}$ are i.i.d. across agents, and that the two collections $\{\theta_{i,t}\}_{i,t}$ and $\{\mathcal{E}_t\}_{t\geq 0}$ are mutually independent.

**A running Krusell–Smith example.** To make this concrete, I work through the entire setup on a Krusell–Smith economy that serves as the running example for the remainder of the paper. Consider a continuum of infinitely-lived households indexed by $i \in [0,1]$, who supply labor inelastically and accumulate capital $k_{i,t}$ to self-insure against idiosyncratic labor-productivity risk under incomplete markets. A representative competitive firm operates a Cobb–Douglas technology in capital and labor, and aggregate uncertainty enters through total factor productivity $Z_{ss} + Z_t$, the deviation $Z_t$ from the mean level $Z_{ss}$ following an AR(1) process. I follow Krusell and Smith [1998] with one departure: the idiosyncratic productivity process $e_{i,t}$ is independent of the aggregate disturbance.

A household enters period $t$ with capital $k_{i,t-1}$ inherited from the previous period, observes its current log productivity $e_{i,t}$, and solves

$$V(k_{i,t-1}, e_{i,t}) = \max_{c_{i,t}, k_{i,t}} \left\{ \frac{{c_{i,t}}^{1-1/\eta}}{1-1/\eta} + \beta\, \mathbb{E}_{i,t} V(k_{i,t}, e_{i,t+1}) \right\}$$

subject to a borrowing constraint $k_{i,t} \geq 0$, the AR(1) law of motion

$$e_{i,t} = \rho^e\, e_{i,t-1} + \epsilon^e_{i,t}, \quad \epsilon^e_{i,t} \sim \mathcal{N}(0, \sigma_e^2),$$

and the budget constraint

$$c_{i,t} + k_{i,t} = (1 + r_t)\, k_{i,t-1} + w_t\, \exp(e_{i,t}),$$

where $\eta > 0$ is the intertemporal elasticity of substitution, $\beta \in (0,1)$ the discount factor, $\rho^e \in (-1,1)$ the persistence of log productivity, and $r_t$, $w_t$ the rental rate of capital and the wage rate. The associated optimality conditions are

$$\begin{aligned}
&\nu_{i,t} \equiv \partial V_{i,t}/\partial k = (1 + r_t)\, {c_{i,t}}^{-1/\eta}, \\
&k_{i,t} \geq 0 \perp {c_{i,t}}^{-1/\eta} - \beta\, \mathbb{E}_{i,t} \nu_{i,t+1} \geq 0, \\
&c_{i,t} + k_{i,t} = (1 + r_t)\, k_{i,t-1} + w_t\, \exp(e_{i,t}),
\end{aligned}$$

where $\nu_{i,t}$ is the marginal value of capital. These three equations are the heterogeneous block $F$; they pin down $x_{i,t} = [c_{i,t}, k_{i,t}, \nu_{i,t}]'$ given the predetermined capital stock $a_{i,t-1} = [k_{i,t-1}]'$, the idiosyncratic shock $\theta_{i,t} = [e_{i,t}]'$, and the relevant aggregate prices $Y_t = [r_t, w_t]'$.

Competitive factor pricing and capital market clearing combine to

$$\begin{aligned}
Y^a_t &= \left(Z_{ss} + Z_t\right) {K_{t-1}}^{\alpha}, \\
r_t &= \alpha \left(Z_{ss} + Z_t\right) {K_{t-1}}^{\alpha-1} - \delta, \\
w_t &= (1-\alpha) \left(Z_{ss} + Z_t\right) {K_{t-1}}^{\alpha}, \\
K_t &= \int k_{i,t}\, di, \\
Z_t &= \rho^Z\, Z_{t-1} + \epsilon^Z_t,
\end{aligned}$$

where $Y^a_t$ denotes aggregate output (renamed to keep it distinct from the relevant-aggregates notation $Y_t$), $Z_{ss}$ is the mean level of TFP, a parameter; labor, supplied inelastically, is normalized to one and drops from the displays, while the listing below carries it explicitly as $\mathtt{L}$; $\rho^Z \in (-1,1)$ is the persistence of the TFP deviation, and $\epsilon^Z_t$ the i.i.d. Gaussian innovation feeding it. These five equations form the aggregate block $G$; they determine $X_t = [Z_t, K_t, r_t, w_t, Y^a_t]'$ given the predetermined state $A_{t-1} = [K_{t-1}, Z_{t-1}]'$ and the aggregate innovation $\mathcal{E}_t = [\epsilon^Z_t]'$.

**Writing the model in a `.mod` file.** How does a Dynare `.mod` file describe such a problem? The aggregate endogenous variables $X_t$, innovations $\mathcal{E}_t$, and equations $G$ are declared using the standard `var`, `varexo`, and `model` syntax. For the variables and equations featuring heterogeneity, we adjoin the `var`, `varexo`, `model`, and `shocks` keywords with the same option `heterogeneity=NAME`, which binds the declaration or block to the heterogeneity dimension $\mathcal{I}$ declared at the top of the file by the `heterogeneity_dimension NAME;` statement. The current release of Dynare 7 supports a single heterogeneity dimension per `.mod` file; admitting several dimensions concurrently is a planned extension. Declarations not adorned with the option live in the aggregate scope, exactly as in any standard Dynare model. Note that the aggregate equations include cross-sectional integrals $\int y_{i,t}\, di$ that link the aggregate state to the distribution of heterogeneous variables. These integrals are invoked through the `SUM(var)` construct, in which *`var`* is a single heterogeneous endogenous variable declared along the relevant dimension. The macro `SUM` denotes a continuum integral over agents, not a finite summation, and it may appear only inside an aggregate equation. The following code shows how the user can write the version of the Krusell–Smith model described above in a `.mod` file.

```
// Declare heterogeneity dimension
heterogeneity_dimension households;

// Household-level variables
var(heterogeneity=households)
    c  (long_name = 'Consumption')
    a  (long_name = 'Assets')
    Va (long_name = 'Derivative of the value function w.r.t assets')
;

// Household-level shocks
varexo(heterogeneity=households)
    e  (long_name = 'Idiosyncratic productivity shock')
;

// Aggregate variables
var
    Y (long_name = 'Aggregate output')
    r (long_name = 'Rate of return on capital net of depreciation')
    w (long_name = 'Wage rate')
    K (long_name = 'Aggregate capital')
    Z (long_name = 'Aggregate productivity, deviation from Z_ss')
;

// Innovation to the productivity AR(1)
varexo eps_Z (long_name = 'Innovation to aggregate productivity');

// Parameters
parameters
    L     (long_name = 'Labor')
    alpha (long_name = 'Share of capital in production function')
    beta  (long_name = 'Subjective discount rate of households')
    delta (long_name = 'Capital depreciation rate')
    eis   (long_name = 'Elasticity of intertemporal substitution')
    rho_Z (long_name = 'Aggregate TFP persistence')
    Z_ss  (long_name = 'Aggregate TFP average value')
;

L = 1;
Z_ss = 0.8816460975214567;
alpha = 0.11;
beta = 0.9819527880123727;
delta = 0.025;
eis = 1;
rho_Z = 0.80;

// Household optimization problem
model(heterogeneity=households);
    [name='Euler equation with borrowing constraint']
    c^(-1/eis) = beta*Va(+1)   a>=0;

    [name='Budget constraint']
    c + a = (1+r)*a(-1) + w*e;

    [name='Envelope condition']
```

```
    Va = (1+r)*c^(-1/eis);
end;

// Aggregate equilibrium conditions
model;
    [name='Production function']
    Y = (Z_ss + Z) * K(-1)^alpha * L^(1 - alpha);

    [name='Capital rental rate']
    r = alpha * (Z_ss + Z) * (K(-1) / L)^(alpha - 1) - delta;

    [name='Wage rate']
    w = (1 - alpha) * (Z_ss + Z) * (K(-1) / L)^alpha;

    [name='Capital market clearing']
    K = SUM(a);

    [name='Productivity process']
    Z = rho_Z * Z(-1) + eps_Z;
end;

// Aggregate shock specification
shocks;
    var eps_Z; stderr 0.01;
end;
```

**Notation in the listing.** A few notational bridges between the mathematical setup and the listing are worth pausing on before turning to the restrictions that the preprocessor imposes. The marginal value of capital $\nu_{i,t}$ is declared as `Va` (the variable's long name spells out that it is the derivative of the value function with respect to assets). The household's capital stock $k_{i,t}$ is declared simply as `a`: Dynare 7 uses the generic term *assets* across all such models, so that the same syntax serves models with bonds, debt, or multi-asset portfolios. The elasticity of intertemporal substitution $\eta$ is declared as `eis`, aggregate output $Y_t^a$ as `Y`, the rental rate $r_t$ as `r`, and the wage $w_t$ as `w`; at the calibrated $\eta = 1$ the period utility reads as its log limit, $\log c_{i,t}$, while the first-order conditions in the listing, written directly in $c^{-1/\eta}$, need no case distinction. The TFP deviation $Z_t$ keeps its name and enters the production function and factor prices through the level $Z_{ss} + Z_t$; its AR(1) law of motion is an aggregate model equation, with the persistence $\rho^Z$ declared as `rho_Z` and the innovation $\epsilon_t^Z$ as `eps_Z`. One convention switch between the mathematics and the listing is silent and worth naming: the mathematical budget constraint carries $\exp(e_{i,t})$ with $e$ log-productivity, whereas the listing's budget reads `w*e` with no exponential, because the `rouwenhorst` helper of Appendix A.1 returns its grid already exponentiated and renormalized to unit mean, so the declared `e` is the productivity *level*. The calibration is that of the Krusell–Smith notebook of the `sequence-jacobian` toolkit of Auclert et al. [2021] ($\eta = 1$, $\delta = 0.025$, $\alpha = 0.11$, $\rho^e = 0.966$, $\sigma_e = 0.5$, $\rho^Z = 0.80$, $L = 1$), with $\beta$ the calibrated value that clears the asset market and $Z_{ss}$ normalizing steady-state output to one (Appendix B). The heterogeneity dimension is named `households` and is referenced by the `heterogeneity=households` option on the household-level `var`, `varexo`, and `model` blocks, while the cross-sectional integral $\int k_{i,t}\, di$ that ties aggregate capital to the household distribution is written `SUM(a)` in the capital-market-clearing equation.

**Equation labels.** Each equation in the listing is preceded by an optional tag of the form `[name='...']`. This is standard Dynare practice rather than something specific to Dynare 7's heterogeneous-agent features: the label is reproduced verbatim in residual reports, in parser error messages, and in model-introspection output, so that problems can be traced back to a named equation rather than to a numerical index. I follow this convention throughout the running example as a matter of good housekeeping. Other equation tags exist in the Dynare grammar (notably `[static]`, `[dynamic]`, and the OccBin regime tags for occasionally-binding constraints), but they play no role in Dynare 7's heterogeneous-agent features at this stage.

**Mixed complementarity.** The Euler equation in the listing attaches the borrowing constraint $k_{i,t} \geq 0$ to the equation by means of a $\perp$-clause. The convention adopted in heterogeneous-agent model blocks is identical to the one already documented for Dynare's `lmmcp` solver for perfect-foresight models: any equation can be paired with a one- or two-sided inequality on a single endogenous variable, separated from the equation by a perpendicular symbol that is accepted either as the Unicode glyph ⊥ or as the

ASCII triplet `_|_`. The bound condition takes one of the forms `var >= lb`, `var <= ub`, or `lb <= var <= ub`, where `lb` and `ub` are arbitrary expressions in the model's parameters. As a brief reminder, writing $X$ for the constrained variable, $F$ for the equation's residual (which Dynare computes as left-hand side minus right-hand side), and $LB$, $UB$ for the two bounds when present, the construct specifies a mixed complementarity problem of the form

$$\begin{aligned} X = LB &\;\Rightarrow\; F(X) \geq 0, \\ LB < X < UB &\;\Rightarrow\; F(X) = 0, \\ X = UB &\;\Rightarrow\; F(X) \leq 0. \end{aligned}$$

Two writing conventions follow. First, the complementarity clause must be attached to the equation that actually determines $X$; attaching it to an unrelated equation poses a different problem to the solver. Second, the sign of the residual must conform to the MCP form above: when the bound is one-sided lower, the equation must be written so that left-hand side minus right-hand side is non-negative at the binding point, and conversely for a one-sided upper bound. This sign requirement bites only at the steady state, where the solver reads whether each agent's bound binds or stays slack off the complementarity clause and a misaligned sign selects the wrong case, breaking the solve. Once the binding pattern is fixed by the steady state, the first-order dynamics operate on the interior condition $F = 0$ alone, which is insensitive to the residual's sign, so a sign error left in the model file does not surface at the dynamics stage.

**Implicit conditional expectations.** The convention adopted here, with the conditional expectations placed on the individual forward-looking variables inside $F$ and $G$ rather than on $F$ and $G$ themselves, is the opposite of standard Dynare's, where the expectation operator sits outside the equation and the perturbation algorithm operates on

$$\mathbb{E}_{i,t} F(\theta_{i,t}, a_{i,t-1}, x_{i,t}, x_{i,t+1}, A_{t-1}, X_t, X_{t+1}, \mathcal{E}_t) = 0, \quad \mathbb{E}_t G\left(\int y_{i,t}\, di,\; A_{t-1},\; X_t,\; X_{t+1},\; \mathcal{E}_t\right) = 0,$$

following Collard and Juillard [2001a,b], Adjemian et al. [2011]. Writing the heterogeneous block in this form expands the relevant-aggregates slot $Y_t \subseteq [A_{t-1}, X_t, \mathbb{E}_t X_{t+1}, \mathcal{E}_t]$ into its constituents; the aggregate expectation $\mathbb{E}_t$ that pre-multiplied $X_{t+1}$ inside $Y_t$ is then absorbed into the outer agent-level expectation $\mathbb{E}_{i,t}$, since under rational expectations the agent observes the aggregate information set and $\mathbb{E}_{i,t}\mathbb{E}_t = \mathbb{E}_{i,t}$ when acting on aggregate variables. The two writings are equivalent up to the introduction of auxiliary variables, so adopting one or the other is without loss of generality. Linearity of the conditional expectation makes the equivalence transparent whenever forward-looking variables enter $F$ separably, that is, through affine combinations of the form $\sum_k \alpha_k \mathbb{E}_{i,t} x_{i,t+1}^{(k)}$: each lead-(+1) term can equivalently be read as living inside or outside the expectation. When forward-looking variables enter $F$ through a non-linear sub-expression $\phi(x_{i,t+1})$, the equivalence is recovered through an auxiliary variable: the preprocessor introduces $z_{i,t} \equiv \phi(x_{i,t})$, the same sub-expression at the current period, appends the dummy equation $z_{i,t} = \phi(x_{i,t})$, and substitutes its lead $z_{i,t+1}$ for $\phi(x_{i,t+1})$ in $F$, where the forward term then enters affinely. From the user's perspective, this rewriting is automatic: the Dynare 7 preprocessor scans each heterogeneous equation, isolates non-affine forward-looking sub-expressions, and creates the corresponding auxiliary variables and equations behind the scenes.

It is worth separating the two kinds of forward-looking term that the inner expectation governs, because Dynare 7 treats them differently once the dynamics are taken. A heterogeneous lead such as $\mathbb{E}_{i,t} x_{i,t+1}$ remains a genuine conditional expectation: it integrates the next-period idiosyncratic shock against the Markov kernel of $\theta_{i,t}$, so idiosyncratic risk is priced. An aggregate lead such as $\mathbb{E}_t X_{t+1}$, by contrast, is treated as a perfect-foresight (certainty-equivalent) term: the first-order dynamics compute the linearized response to a one-time aggregate (MIT) shock, along which the aggregate path is deterministic, so the aggregate lead enters as a plain (+1) with no expectation operator attached, the certainty equivalence that Proposition 11 establishes at order one. The two collapse to the same constant at the steady state, where aggregate variables and their conditional expectations coincide, which is why the distinction surfaces only once the dynamics are computed.

**Non-separability.** The preprocessor imposes a non-separability restriction on the heterogeneous block: a sub-expression that entangles a heterogeneous endogenous lead $x_{i,t+1}$ and a heterogeneous endogenous lag $x_{i,t-1}$ inside a non-separable operator is rejected at parse time. Addition and subtraction always

separate a lead from a lag; multiplication separates them whenever one factor carries no lead; division does so only when the lead sits in the numerator and the denominator carries no lead, a lead inside a denominator being treated as entangled, like a power; every other operator (a power, log, exp, min, max) entangles them. Concretely, expressions of the form

$$\log(x_{i,t-1} + x_{i,t+1})\,, \qquad \exp(x_{i,t-1} \cdot x_{i,t+1})\,, \qquad (x_{i,t-1} + x_{i,t+1})^{\alpha}$$

are not accepted, whereas the separable combinations

$$\beta\, x_{i,t+1} + (1 + r_t)\, x_{i,t-1}, \qquad x_{i,t-1} \cdot c_{i,t+1}, \qquad (x_{i,t-1})^{\alpha} x_{i,t+1}$$

are admissible: the last two are separable because the lag enters as a lead-free multiplicative coefficient, however non-linear that coefficient is in $x_{i,t-1}$. The technical reason is that the preprocessor recovers the equivalence by isolating a purely forward sub-expression $\phi(x_{i,t+1})$ and substituting the lead $z_{i,t+1}$ of the current-period auxiliary variable $z_{i,t} \equiv \phi(x_{i,t})$, the scheme of the previous paragraph; a non-separable coupling admits no such forward-only factor, and would require internally a second pass of expectations and an iterative computation of the resulting fixed point, which Dynare 7 does not yet implement. Aggregate variables are not subject to this restriction: standard Dynare admits arbitrary non-linear couplings of aggregate leads and lags. Dynare's preprocessor reports back the offending sub-expression to the user with its equation number and source line.

**Block squareness.** Each block is required to be square. The preprocessor enforces $\dim(F) = \dim(x_{i,t})$ at every heterogeneity dimension and $\dim(G) = \dim(X_t)$ at the aggregate level, with the few auxiliary variables that the preprocessor introduces counted on both sides of the equality. The `.mod` file must therefore already declare exactly as many equations as endogenous variables in each block.

A side effect of this rigidity is that redundant equations cannot be carried alongside the model for verification purposes. The aggregate resource constraint of the Krusell–Smith economy, for instance, is implied by the household budget constraint, capital-market clearing, and competitive factor pricing through Walras' law, and including it as an extra equation would be a useful runtime sanity check; the current preprocessor rejects this construction. Relaxing the squareness check to admit diagnostic equations flagged as such by the user is a natural improvement for future versions of Dynare 7.

**Specification of the idiosyncratic shocks.** Each heterogeneous exogenous variable is required to follow a Markov process, so that its discretization reduces to a grid of nodes and a transition kernel (the $\pi$ of Section 2). Two modeling restrictions follow from the independence assumptions made in the abstract framework. A shock that is common to all agents cannot be declared as a heterogeneous exogenous variable: being shared across the continuum, it belongs in the aggregate block as an entry of $X_t$ driven by an innovation in $\mathcal{E}_t$. And the idiosyncratic process may not be correlated with the aggregate state: any such dependence has to be routed through the relevant-aggregates slot $Y_t$ that enters $F$, not through the law of motion of $\theta_{i,t}$. Given its frequent occurrence in the literature, the latter case should be taken care of in a future version of Dynare.

**Admissible leads and lags.** The two blocks are treated asymmetrically with respect to the leads and lags they admit. On the heterogeneous side, the preprocessor restricts these at parse time: heterogeneous endogenous variables are admitted with lag in $\{-1, 0\}$ and lead in $\{0, +1\}$ only, and heterogeneous exogenous variables are admitted contemporaneously, with lag $= 0$. Any deeper lag or further lead on a heterogeneous variable is rejected. This restriction mirrors the abstract framework: the solver iterates on the time slices $(a_{i,t-1}, x_{i,t}, \mathbb{E}_{i,t} x_{i,t+1})$, and admitting wider leads or lags would presuppose a richer state for the heterogeneous problem. If the user's economic model genuinely features, say, a two-period-ahead heterogeneous lead, the practical workaround is to introduce a one-period lead auxiliary variable $\tilde{x}_{i,t} \equiv x_{i,t+1}$ by hand and rewrite the equation in terms of $\tilde{x}_{i,t+1}$. On the aggregate side, the standard Dynare strategy applies unchanged: aggregate endogenous variables with $|\text{lag}| \geq 2$ and aggregate exogenous variables with any non-zero lead or lag are substituted out automatically into chained auxiliary variables. The user thus enjoys the usual Dynare freedom on $X_t$ and $\mathcal{E}_t$, but is held to a narrow range of leads and lags on $x_{i,t}$ and $\theta_{i,t}$. The narrow window is load-bearing beyond parsing: the risk corrections of Section 5.2 close in one step precisely because no conditional expectation of the model reaches further than one period ahead (Proposition 13).

**What the preprocessor generates.** Once the `.mod` file is parsed, the preprocessor emits a MATLAB package `+`$M$ next to the model file, where $M$ is the model name. The routines bundled in this package fall into two groups, one per equilibrium block.

On the heterogeneous side, `dynamic_het1_resid.m` (with its temporary-term helper `dynamic_het1_resid_tt.m`) computes the residuals of the heterogeneous block $F$, evaluated at user-supplied heterogeneous states $(\theta_{i,t}, a_{i,t-1}, x_{i,t}, \mathbb{E}_{i,t}x_{i,t+1})$ and at the relevant aggregate slot $Y_t$. The first- and second-order derivatives of these residuals with respect to $x_{i,t}$, $\mathbb{E}_{i,t}x_{i,t+1}$, and $Y_t$ live in `dynamic_het1_g1.m` and `dynamic_het1_g2.m`, each with its `_tt` helper.

Two further files round out the heterogeneous-side contract: `dynamic_het1_set_auxiliary_variables.m` assigns values to auxiliary variables that the preprocessor introduces (chiefly the complementarity slacks discussed below), and `dynamic_het1_complementarity_conditions.m` returns the lower and upper bounds attached to each heterogeneous endogenous variable through $\perp$-clauses. The numeric `het1` suffix is the heterogeneity-dimension index, which is always 1 in the current release of Dynare 7 but anticipates the planned extension to multiple dimensions concurrently.

On the aggregate side, the same conventions apply minus the `het1` decoration: `dynamic_resid.m`, `dynamic_g1.m`, `dynamic_g2.m`, and `set_auxiliary_variables.m` cover the aggregate block $G$. The aggregate-side routines accept an additional argument `yagg` that delivers the cross-sectional integrals appearing in $G$. Static counterparts, used at steady-state pre-evaluation, are emitted under the names `static_resid.m`, `static_g1.m`, and so on.

**Auxiliary variables and complementarity slacks.** The mixed-complementarity construct documented above is implemented by the preprocessor through an auxiliary-variable expansion that I now spell out on the running example. The user-level Euler equation with borrowing constraint,

$$c_{i,t}^{-1/\eta} - \beta\,\mathbb{E}_{i,t}\nu_{i,t+1} = 0 \quad \perp \quad k_{i,t} \geq 0,$$

is rewritten into a system that the solver can handle without case-splitting. The preprocessor introduces an auxiliary variable that I denote $\mu_{i,t}$ here (its actual internal name follows the pattern `MULT_L_`*`var`* for a lower bound and `MULT_U_`*`var`* for an upper one, so `MULT_L_a` in the Krusell–Smith case), and emits, in `dynamic_het1_resid.m`, two equations in place of the original Euler-plus-clause:

$$\begin{aligned}\left(c_{i,t}^{-1/\eta} - \beta\,\mathbb{E}_{i,t}\nu_{i,t+1}\right) - \mu_{i,t} &= 0,\\ k_{i,t}\cdot\mu_{i,t} &= 0.\end{aligned}$$

The second equation is the slackness condition: at any solution, either the constraint binds ($k_{i,t} = 0$) or the multiplier is zero ($\mu_{i,t} = 0$), and the original Euler is recovered wherever the constraint is slack. The numerical value of $\mu_{i,t}$ is assigned in `dynamic_het1_set_auxiliary_variables.m` by a closed-form expression of the form $\mu_{i,t} = \left(c_{i,t}^{-1/\eta} - \beta\,\mathbb{E}_{i,t}\nu_{i,t+1}\right)\cdot\mathbf{1}_{\{k_{i,t}\leq 0\}}$, so that the slack vanishes where the constraint is slack and equals the Euler residual where it binds. The pair ($F = 0$, slackness) together with the bounds returned by `dynamic_het1_complementarity_conditions.m` (which records the entry $\texttt{lb}(\cdot) = 0$ at the asset variable's index in the Krusell–Smith case) constitute the input to the Fischer–Burmeister handling that closes the loop at solver time.

**The `SUM` operator and the `yagg` argument.** Cross-sectional aggregation through the `SUM` macro is similarly expanded at preprocessor time. The user equation

$$K_t = \texttt{SUM}(k_{i,t})$$

becomes two equations in `dynamic_resid.m`. The preprocessor introduces a fresh aggregate auxiliary variable, which I denote $S_t$ here (its internal name follows the pattern `SUM_`$\langle var\rangle$, here `SUM_a`), and emits

$$\begin{aligned}K_t - S_t &= 0,\\ S_t - \texttt{yagg}(\cdot) &= 0.\end{aligned}$$

The first of these is the user's original equation with the `SUM` slot replaced by the auxiliary variable's name; the second binds the auxiliary to the cross-sectional integral fed into the routine at solver time through the `yagg` argument. At runtime, Dynare 7 computes $\int k_{i,t}\,d\Omega_t(e_{i,t}, k_{i,t-1})$, the converged household policy integrated against the discretized joint cross-sectional distribution $\Omega_t$ of (current shock, inherited capital) that Section 2 defines, then passes it as the appropriate slot of `yagg` into the aggregate residual

function. The construct `SUM` is therefore not a function evaluated at runtime; it is a macro expanded at preprocessor time into an auxiliary variable plus a binding equation, and the cross-sectional integral itself lives entirely in `yagg`.

## 2 Recursive form

The time-indexed presentation of Section 1 is convenient for stating the model, but the numerical algorithms operate on the recursive representation of the equilibrium. I follow Bhandari et al. [2023], dropping time subscripts and using primes to denote next-period quantities. The aggregate state is the triple

$$\mathcal{Z} \equiv (\mathcal{E},\, A,\, \Omega)\,,$$

where $\mathcal{E}$ is the current aggregate innovation, $A$ is the predetermined-aggregate vector (the subset of aggregate endogenous variables carried into the current period from the previous one), and $\Omega$ is the cross-sectional distribution of $(\theta, a)$ over $\Theta \times \mathcal{A}$, with $\Theta$ the shock space and $\mathcal{A}$ the predetermined-state space. I adopt the timing convention that $(\theta, a)$ denotes the agent's state at the start of the current period *after* the current-period idiosyncratic shock has been observed: $\theta$ is the current shock (the random variable $\theta_{i,t}$ of Section 1) and $a$ is the predetermined state inherited from the previous period ($= a_{i,t-1}$ in time-indexed notation). Consequently $\Omega$ is the joint distribution of (current shock, previous predetermined state) at the start of the current period. The equilibrium is encoded by three policy functions

$$\bar{x} : \Theta \times \mathcal{A} \times \mathfrak{Z} \to \mathbb{R}^{n_x}, \qquad \bar{X} : \mathfrak{Z} \to \mathbb{R}^{n_X}, \qquad \bar{\Omega} : \mathfrak{Z} \to \mathcal{M}(\Theta \times \mathcal{A})\,,$$

where $\mathfrak{Z}$ is the space of aggregate-state triples $(\mathcal{E}, A, \Omega)$ the economy can visit and $\mathcal{M}(\Theta \times \mathcal{A})$ the space of probability measures on $\Theta \times \mathcal{A}$. The function $\bar{x}(\theta, a, \mathcal{Z})$ returns the optimal value of the heterogeneous endogenous variable for an agent in state $(\theta, a)$ when the aggregate state is $\mathcal{Z}$; analogously $\bar{a}(\theta, a, \mathcal{Z})$ is the policy for the next-period predetermined state. The bar distinguishes the policy, a decision rule defined on the state space, from the time-indexed random variable: along an equilibrium path, $x_{i,t} = \bar{x}(\theta_{i,t}, a_{i,t-1}, \mathcal{Z})$. The aggregate state evolves according to

$$\mathcal{Z}' \equiv \big(\mathcal{E}',\, P\,\bar{X}(\mathcal{Z}),\, \bar{\Omega}(\mathcal{Z})\big)\,,$$

where $\mathcal{E}'$ is the next-period aggregate innovation, $P$ is the constant selection matrix that picks out the predetermined-aggregate entries of $\bar{X}(\mathcal{Z})$ (a $\dim(A) \times n_X$ matrix of zeros and ones whose rows are an orthonormal subset of the canonical basis of $\mathbb{R}^{n_X}$), and $\bar{\Omega}(\mathcal{Z})$ is the next-period cross-sectional distribution.

In this representation, the equilibrium conditions read

$$\begin{aligned}
&F\big(\theta,\, a,\, \bar{x}(\theta, a, \mathcal{Z}),\, \mathbb{E}[\bar{x}(\theta', \bar{a}(\theta, a, \mathcal{Z}), \mathcal{Z}') \,|\, \theta, \mathcal{Z}]\,,\, \bar{Y}(\mathcal{Z})\big) = 0 \quad \forall (\theta, a) \in \Theta \times \mathcal{A},\\
&G\big(\textstyle\int \bar{y}(\theta, a, \mathcal{Z})\, d\Omega(\theta, a),\, A,\, \bar{X}(\mathcal{Z}),\, \mathbb{E}\big[\bar{X}(\mathcal{Z}') \,\big|\, \mathcal{Z}\big]\,,\, \mathcal{E}\big) = 0,\\
&\bar{\Omega}(\mathcal{Z})(B) = \int \mathbf{1}\{(\theta', \bar{a}(\theta, a, \mathcal{Z})) \in B\}\; d\pi(\theta' \mid \theta)\, d\Omega(\theta, a),
\end{aligned}$$

where $\bar{Y}(\mathcal{Z}) \subseteq \big[A,\, \bar{X}(\mathcal{Z}),\, \mathbb{E}[\bar{X}(\mathcal{Z}') \mid \mathcal{Z}],\, \mathcal{E}\big]$ is the recursive-form image of the relevant-aggregates slot $Y_t$ from Section 1, $\pi$ is the Markov kernel governing $\theta_{i,t}$, $\theta' \sim \pi(\cdot \mid \theta)$ is the next-period idiosyncratic shock, and $B$ ranges over the Borel subsets of $\Theta \times \mathcal{A}$. The conditional expectations are taken jointly over $(\theta', \mathcal{E}')$, with $\mathcal{Z}'$ derived from $\mathcal{Z}$ and $\mathcal{E}'$ by the law of motion stated above. In the distribution equation, $\bar{a}$ acts on the *current* shock $\theta$ (consistently with the F equation that says the agent observes $\theta$ at the start of the period), and the next-period shock $\theta'$ enters the pushforward only through the indicator over the next-period state. Note that $\bar{\Omega}(\mathcal{Z})$ enters $\mathcal{Z}'$, so the distribution itself is an endogenous state.

The steady state of the model laid out in Section 1 is the situation in which the aggregate innovation $\mathcal{E}_t$ is switched off, aggregate variables are held constant over time, and the cross-sectional distribution of heterogeneous agents is invariant under the policy-induced transition. It is the natural reference point around which the dynamic solution is later constructed, and computing it is the first numerical step of solving any heterogeneous-agent model. At a steady state, the aggregate innovation vanishes, $\mathcal{E} = 0$, and

the aggregate state is constant in time at a value I denote $\mathcal{Z}^*$. The general policy functions then drop their aggregate-state argument: I introduce the shorthands

$$\bar{x}^*(\theta,a) \equiv \bar{x}(\theta,a,\mathcal{Z}^*), \qquad \bar{a}^*(\theta,a) \equiv \bar{a}(\theta,a,\mathcal{Z}^*), \qquad \bar{y}^*(\theta,a) \equiv \bar{y}(\theta,a,\mathcal{Z}^*),$$
$$\bar{X}^* \equiv \bar{X}(\mathcal{Z}^*), \qquad \bar{Y}^* \equiv \bar{Y}(\mathcal{Z}^*), \qquad \bar{\Omega}^* \equiv \bar{\Omega}(\mathcal{Z}^*),$$

where the star marks "general policy function evaluated at the steady-state aggregate state". With this convention $\bar{X}^*$ and $\bar{Y}^*$ are real vectors and $\bar{\Omega}^*$ is the stationary cross-sectional measure on $\Theta \times \mathcal{A}$. The recursive system at the steady state then becomes the time-invariant one

$$F\big(\theta,\, a,\, \bar{x}^*(\theta,a),\, \mathbb{E}_{\theta'|\theta}[\bar{x}^*(\theta',\bar{a}^*(\theta,a))]\,,\, \bar{Y}^*\big) = 0 \quad \forall(\theta,a) \in \Theta \times \mathcal{A},$$
$$G\big(\textstyle\int \bar{y}^*(\theta,a)\, d\bar{\Omega}^*(\theta,a),\, P\bar{X}^*,\, \bar{X}^*,\, \bar{X}^*,\, 0\big) = 0,$$
$$\bar{\Omega}^*(B) = \int \mathbf{1}\{(\theta',\bar{a}^*(\theta,a)) \in B\}\; d\pi(\theta' \mid \theta)\, d\bar{\Omega}^*(\theta,a),$$

where the three occurrences of $\bar{X}^*$ in $G$ fill the slots for $A$, $\bar{X}(\mathcal{Z})$, and $\mathbb{E}[\bar{X}(\mathcal{Z}') \mid \mathcal{Z}]$ (all equal at the steady state since aggregate variables and their conditional expectations coincide with their constant steady-state values), the conditional expectation on the $F$ equation collapses to the idiosyncratic-only $\mathbb{E}_{\theta'|\theta}$ (the aggregate state is fixed and $\mathcal{E}' = 0$), and the distribution law again uses $\bar{a}^*(\theta,a)$ for the same timing reason as in the general form. The unknowns are the two policy functions $\bar{x}^*$, $\bar{a}^*$, the stationary measure $\bar{\Omega}^*$, and the aggregate vector $\bar{X}^*$. The two policy equations are infinite-dimensional, the aggregate equation is finite-dimensional, and the last equation is the fixed-point condition that selects the invariant measure of the policy-induced transition. The next section is concerned with how Dynare represents and solves this system numerically.

# 3 Steady state loading and computation

Dynare 7 offers two ways into a steady state: the user may provide a pre-computed one via `heterogeneity_load_steady_s` or ask Dynare 7 to compute it from an initial guess via `heterogeneity_compute_steady_state`. The two commands consume structurally identical objects, so the rest of this section first restates the equilibrium in recursive form, then catalogues the discrete objects the two commands operate on, and finally documents the commands themselves. I then contrast the way Dynare and the `sequence-jacobian` package compute the steady state.

## 3.1 Discretized steady state

The discretized steady-state stems from replacing each continuous object with a discrete counterpart suitable for finite computation. I write every finite-dimensional discretization in bold, on top of the lowercase-red / uppercase-blue color rule of Section 1.

**Shock grids.** Each heterogeneous exogenous variable $e$ is discretized as a finite Markov chain: a grid of nodes $\boldsymbol{\theta}_{[e]}$ with $k$-th node $\boldsymbol{\theta}_{[e],k}$, and a row-stochastic transition matrix $\boldsymbol{\pi}_{[e]}$ whose entry $(\boldsymbol{\pi}_{[e]})_{k\ell}$ is the probability of moving from node $k$ to node $\ell$. With several heterogeneous shocks the per-variable grids and matrices combine into the joint shock grid $\boldsymbol{\theta} = \prod_e \boldsymbol{\theta}_{[e]}$ and the joint kernel $\boldsymbol{\pi} = \bigotimes_e \boldsymbol{\pi}_{[e]}$. The `rouwenhorst` helper of Appendix A.1 produces one such pair for the common AR(1) case.

**Predetermined-state grids.** Each predetermined state $v$ (capital $k$ in Krusell–Smith) carries two grids: a policy grid $\boldsymbol{a}^p_{[v]}$, on which the time iteration solves for the policy, and a distribution grid $\boldsymbol{a}^d_{[v]}$, on which the histogram of the cross-section lives. These grids carry a further $p$ or $d$ superscript distinguishing the policy grid from the distribution grid; the superscript propagates to individual nodes $\boldsymbol{a}^p_{[v],j}$ and $\boldsymbol{a}^d_{[v],j}$. Allowing $\boldsymbol{a}^p_{[v]}$ and $\boldsymbol{a}^d_{[v]}$ to differ is a runtime optimization. Time iteration solves a nonlinear system at every node of the joint shock-and-policy-state tensor, so densifying any $\boldsymbol{a}^p_{[v]}$ multiplies the solve count along that axis; the forward iteration of the distribution sweeps the analogous joint shock-and-distribution-state tensor through a sparse matrix-vector product, so densifying any $\boldsymbol{a}^d_{[v]}$ is cheap. The typical idiom is therefore to keep $\boldsymbol{a}^p_{[v]}$ coarse, for time-iteration speed, and $\boldsymbol{a}^d_{[v]}$ dense, for the accuracy of the cross-sectional integrals against which $G$ is evaluated.

**Arrays on the joint mesh.** On the joint policy mesh $\boldsymbol{\theta} \times \boldsymbol{a}^p$ Dynare 7 tabulates a policy array $\boldsymbol{x}_{[v]}$ for each heterogeneous endogenous variable $v$; on the joint distribution mesh $\boldsymbol{\theta} \times \boldsymbol{a}^d$ it tabulates the cross-section histogram $\boldsymbol{\Omega}$, whose entries are non-negative and sum to one. The aggregate endogenous variables collapse to the finite vector $\bar{X}^*$, and a continuum integral against the distribution becomes the finite sum $\int \bar{y}^* \, d\bar{\Omega}^* \approx \sum \boldsymbol{\Omega} \cdot \boldsymbol{y}$, the policy arrays $\boldsymbol{y}$ first read on the distribution mesh through their interpolants (the prolongation of Appendix C) whenever the two grids differ.

**Input/Outputs of steady-state commands** Both steady-state commands consume these discrete objects as the fields of a single MATLAB structure, supplied either as a workspace variable or saved to a `.mat` file. The fields mirror the objects above one for one:

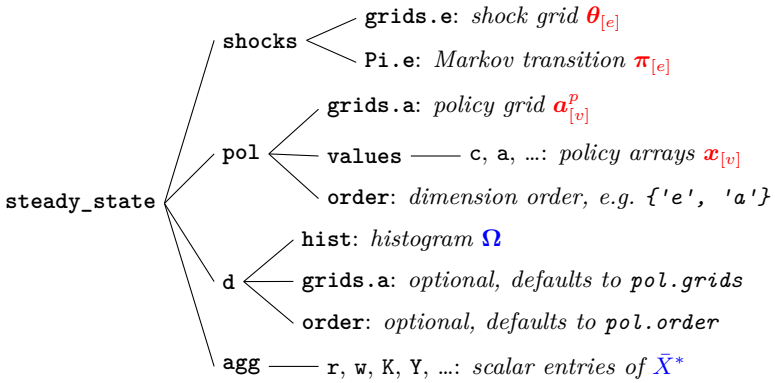


Figure 1: The MATLAB steady-state structure consumed and returned by both commands. The field values shown are those of the Krusell–Smith example.

The cell array `pol.order` fixes the dimension layout of the policy and histogram arrays; the distribution grids `d.grids` default to the policy grids `pol.grids` when omitted.

The structure is read by either of

```
heterogeneity_load_steady_state(filename = FILE, variable = NAME, [options...]);
heterogeneity_compute_steady_state(filename = FILE, variable = NAME, [options...]);
```

which take the structure named `NAME` from the `.mat` file `FILE`, or from the base workspace when `filename` is omitted. The command `heterogeneity_load_steady_state` takes the supplied arrays as a finished steady state: it checks the aggregate-block residual against a tolerance `tolf` (default $10^{-6}$),

$$\left\| G\left(\sum \boldsymbol{\Omega} \cdot \boldsymbol{y},\, P\bar{X}^*,\, \bar{X}^*,\, \bar{X}^*,\, 0\right) \right\|_\infty < \texttt{tolf},$$

trusts the heterogeneous block without re-solving it, and writes the validated steady state to `oo_.heterogeneity.steady_`. The command `heterogeneity_compute_steady_state` instead treats the same structure as an initial guess and solves the discretized steady-state system from it, as detailed in Subsection 3.2.

In both cases the steady state is returned in `oo_.heterogeneity.steady_state` as a structure with the same fields as Figure 1, now holding the validated or converged policy arrays, the stationary histogram, and the aggregate scalars. The steady-state structure therefore plays a dual role, as the input the commands consume and the output they return, so a steady state produced by one run can be fed straight back into a later one.

## 3.2 Computing the steady state

When no steady state is supplied ready-made, `heterogeneity_compute_steady_state` produces one from the initial guess of Figure 1. It is convenient to phrase what the command does as solving a small calibration problem.

**Supplying the initial guess.** The command reads the structure of Figure 1 as an initial guess, with two differences: `d.hist` may be omitted, since the forward iteration recomputes the histogram, and the optional `free_parameters` field flags the parameters to be calibrated, each with an initial value and optional bounds. The clean, documented way to supply it is to build the structure in a standalone script, save it to a `.mat` file, and pass that file through the `filename` option. It can also be assembled inline in

the `.mod` file itself and handed over as a workspace variable through the `variable` option, which keeps the model and its initial guess in a single source file; this route is convenient but a bit hackish, since the Dynare preprocessor forces parameter assignments to sit inline while the `initial_guess` structure must be built inside `verbatim;` blocks. Appendix B works that inline Krusell–Smith construction in full.

**Calibration as a residual-map zero.** Partition the inputs and the aggregate block. A handful of model parameters $\boldsymbol{\beta} \in \mathbb{R}^{n_\beta}$ are peeled off the calibrated set and declared *free* (e.g. the discount factor $\beta$ in Krusell–Smith), and as many aggregate equations $H \subseteq G$ are picked as *targets* (e.g. capital-market clearing $K = \int k_i \, di$ in Krusell–Smith). For any value of $\boldsymbol{\beta}$, the remaining equations have a unique steady state $(\bar{x}^*, \bar{\Omega}^*, \bar{X}^*)$, and the *calibration residual map*

$$R : \mathbb{R}^{n_\beta} \to \mathbb{R}^{n_\tau}, \qquad R(\boldsymbol{\beta}) \equiv H\left( \int \bar{y}^* \, d\bar{\Omega}^*, \, P\bar{X}^*, \, \bar{X}^*, \, \bar{X}^*, \, 0 \right),$$

returns the residual of the target equations at that steady state. Computing the steady state is then solving $R(\boldsymbol{\beta}^*) = 0$, a square system since Dynare 7 requires as many free parameters as target equations, $n_\beta = n_\tau$. The targets default to the aggregate equations that contain a `SUM` operator and can be overridden with `calibration_target_equations = ['name1', 'name2']`, listed by their `[name='...']` tags.

## 3.3 The algorithm in detail

Each evaluation of $R(\boldsymbol{\beta})$ runs three steps at the current parameters: a time iteration that converges the policy arrays $\boldsymbol{x}$, a forward iteration that converges the histogram $\boldsymbol{\Omega}$, and an aggregation that forms the cross-sectional integrals and returns the residual. When a `free_parameters` field is supplied, an outer Broyden step wraps these three and adjusts $\boldsymbol{\beta}$ until $\|R(\boldsymbol{\beta})\|_\infty <$ `calibration_tolf`, overwriting the corresponding entries of `M_.params` on exit; when it is absent, the three steps run once at the parameters held in `M_.params` and the aggregate residuals are reported unchanged (residual-only mode). None of these is specific to Dynare 7, so I take them in turn below, giving the canonical reference for each and naming the governing options inline with their defaults in parentheses.

**Step 1: time iteration.** The policy is found by *time iteration*: holding next period's policy fixed, one solves the agents' current-period optimality conditions for the current policy, and iterates the resulting map to its fixed point [Coleman, 1990]. A sweep carries the current iterate $\bar{x}^{(n)}$ to a half-step update $\bar{x}^{(n+1/2)}$ by solving, independently at each node $(\theta, a)$ of the policy mesh $\boldsymbol{\theta} \times \boldsymbol{a}^p$,

$$F\left( \theta, a, \bar{x}^{(n+1/2)}(\theta, a), \mathbb{E}_{\theta'|\theta}\left[ \bar{x}^{(n)}\left( \theta', \bar{a}^{(n+1/2)}(\theta, a) \right) \right], \bar{Y}^* \right) = 0,$$

where the continuation policy $\bar{x}^{(n)}$ is reconstructed off the grid by multilinear interpolation on $\boldsymbol{a}^p$, leaving the current-node policy as the only unknown. A $\perp$ clause makes its node a nonlinear complementarity problem rather than a plain equation; Dynare 7 reformulates it through the Fischer–Burmeister function, which collapses a clause $u \geq 0 \perp w \geq 0$ into the single smooth equation

$$\phi_{FB}(u, w) = u + w - \sqrt{u^2 + w^2 + \varepsilon} = 0, \qquad \varepsilon = 10^{-14},$$

the same complementarity handling as Dynare's `lmmcp` solver. The node problem is then a smooth square system, solved by a derivative-based trust-region method; the small slack the regularization leaves at a binding constraint is removed afterwards by snapping any coordinate that has settled within $\sqrt{\varepsilon} \approx 10^{-7}$ of its bound exactly onto it. Successive sweeps are relaxed,

$$\bar{x}^{(n+1)} = (1 - \lambda)\, \bar{x}^{(n)} + \lambda\, \bar{x}^{(n+1/2)},$$

with relaxation weight $\lambda$ set by `time_iteration_learning_rate` (default 1, undamped), and the iteration stops once $\|\bar{x}^{(n+1/2)} - \bar{x}^{(n)}\|_\infty$ drops below `time_iteration_tol` ($10^{-8}$), within a budget of `time_iteration_max_iter` sweeps (1000). An early-stopping guard aborts if the update norm rises for `time_iteration_early_stopping` successive sweeps (3), and the per-node solver exposes its own tolerances and step bounds through the `time_iteration_solver_*` options.

**Step 2: forward iteration.** With the policy fixed, the cross-section is the invariant distribution of the Markov chain the policy induces on the discretized state, reached by iterating that chain forward, the non-stochastic simulation of Young [2010] (the histogram method of Reiter, 2009). One step pushes the histogram $\boldsymbol{\Omega}$ through two maps: a lottery on the predetermined state, which splits the mass at $(\theta, a)$ between the two $\boldsymbol{a}^d$ nodes bracketing $\bar{a}^*(\theta, a)$ by linear interpolation, so the next-period state need not land on a grid node, followed by the shock kernel $\boldsymbol{\pi}$ on the exogenous axes. Starting from the stationary shock marginal of $\boldsymbol{\pi}$ against a uniform distribution on $\boldsymbol{a}^d$, the iteration runs until $\|\boldsymbol{\Omega}^{(n+1)} - \boldsymbol{\Omega}^{(n)}\|_1$ falls below `forward_tol` ($10^{-10}$), tested every `forward_check_every` steps (100) and capped at `forward_max_iter` ($10^4$).

**Step 3: aggregation.** The cross-sectional integrals in $G$ are now finite sums against the converged histogram, $\int \bar{y}^* \, d\bar{\Omega}^* \approx \sum \boldsymbol{\Omega} \cdot \boldsymbol{y}$ with $\boldsymbol{y}$ prolonged onto the distribution mesh, one per `SUM` operator; substituting them into the aggregate block and reading off the target equations returns the residual $R(\boldsymbol{\beta})$.

**Step 4: outer Broyden.** When free parameters are present, these three steps are wrapped in a quasi-Newton solve: Broyden's method updates $\boldsymbol{\beta}$ from the residual alone, building up an approximate Jacobian as it goes, until $\|R(\boldsymbol{\beta})\|_\infty$ falls below `calibration_tolf` ($10^{-4}$) or `calibration_max_iter` (50) is reached, rejecting any step that would leave the `lower_bound` / `upper_bound` box. This is exactly the computation of a stationary incomplete-markets equilibrium [Huggett, 1993, Aiyagari, 1994]: the free discount factor, equivalently the steady-state interest rate, is moved until the asset market clears.

# 4 Dynamics

The steady state of Section 3 is the point around which the model's response to aggregate shocks is built, and this section perturbs the recursive equilibrium of Section 2 around it to first order. The perturbation method for the heterogeneous block is that of Bhandari et al. [2023], including the treatment of the borrowing-constraint kink and the histogram analysis of their Appendix C; the sequence-space reading of the aggregate block is that of Auclert et al. [2021]. This section reads both through one object: I show, from the operator side, the equivalence Bhandari et al. [2023] establish, that their state-space recursion and the sequence-space Jacobian of Auclert et al. [2021] are two readings of the Fréchet derivative of the equilibrium map at the steady state, and I recast that histogram analysis in a minimal weak-derivative and finite-element language so that the borrowing-constraint atom never forces a derivative onto a Dirac.

The two methods are complementary. The state-space construction of Bhandari et al. [2023] is the more general: it does not depend on the choice of projection basis, but it builds large sparse matrices, whose sparse-by-sparse products can be costly. The sequence-space code of Auclert et al. [2021] specializes to a $P^1$ basis and gains the speed that specialization allows, never forming those matrices explicitly. Dynare 7 aims at the best of both worlds, taking the heterogeneous block from the former (which admits an aggregate lag in the heterogeneous equations) and the aggregate-block assembly from the latter; this section shows the two are first-order compatible, which is what makes the combined approach legitimate.

A single vocabulary serves both literatures. I keep the steady-state macros of Section 1, lowercase red for heterogeneous objects and uppercase blue for aggregates, and introduce each method's own symbols once, where they first appear, before dropping them in favour of the canonical column. Appendix F collects the full correspondence between the symbols of this section, those of Bhandari et al. [2023] and Auclert et al. [2021], and Dynare 7's code field names, for the reader who wants to read across the three codebases. The displacement operator $\mathsf{M}$ is a sans-serif M, kept distinct from the measure space $\mathcal{M}(\Theta \times \mathcal{A})$ of Section 2.

## 4.1 First-order representation

Switch the aggregate innovation back on and read the recursive law of motion of Section 2 as a map from one aggregate state to the next, $\bar{\mathcal{Z}}(\mathcal{Z}) = \big[\mathcal{E},\, P\bar{X}(\mathcal{Z}),\, \bar{\Omega}(\mathcal{Z})\big]$. I scale the aggregate innovations by a scalar perturbation parameter $\sigma$, so that the deterministic steady state $\mathcal{Z}^*$ sits at $\sigma = 0$ and the stochastic economy at $\sigma = 1$, and expand the equilibrium to first order in $\sigma$. The first-order response is the derivative in $\sigma$, equivalently the Fréchet derivative of $\bar{\mathcal{Z}}$, $\bar{X}$, and $\bar{\Omega}$ along the perturbed state path.

Collecting the impulse responses gives the moving-average representation of the aggregate path

$$X_t \;=\; \bar{X}^* + \sum_{s=0}^{t} \widehat{X}_s\, \mathcal{E}_{t-s} + O(\sigma^2)\,, \qquad \widehat{\mathcal{Z}}_0 \equiv [1,\, 0,\, (\theta, a) \mapsto 0]^\top\,, \quad \widehat{\mathcal{Z}}_t \equiv D_{\mathcal{Z}}\bar{\mathcal{Z}} \cdot \widehat{\mathcal{Z}}_{t-1}, \quad \widehat{X}_s \equiv D_{\mathcal{Z}}\bar{X} \cdot \widehat{\mathcal{Z}}_s.$$

Here $D_{\mathcal{Z}}\bar{\mathcal{Z}}$ and $D_{\mathcal{Z}}\bar{X}$ are the Fréchet derivatives of those maps with respect to the aggregate state, taken at the steady state $\mathcal{Z}^*$: each is the bounded linear map that carries a perturbation of the state to the corresponding first-order response. A Fréchet derivative is always read at a base point, and unless I state otherwise every Fréchet derivative in this section is taken at the steady state $\mathcal{Z}^*$. Appendix C defines the notion, and the Gâteaux derivative it refines, in detail. $\widehat{X}_s$ is the aggregate impulse response $s$ periods after a unit innovation. This recursion is the perturbation of the equilibrium map at the steady state: Bhandari et al. [2023, Lemma 1] differentiate the heterogeneous-agent equilibrium in just this way to read off the first-order aggregate response, and casting such a perturbation solution as the moving-average, or Volterra, convolution shown above is the representation Lan and Meyer-Gohde [2013] brought to dynamic stochastic general equilibrium models. The $O(\sigma^2)$ remainder is not a nuisance term: Section 5.2 gives it a computable form, the higher kernels of the same expansion.

What the convolution needs is the response of the aggregate path $\widehat{X}_t$, and the rest of this section builds it in the order the problem dictates (Figure 2). Differentiating the aggregate equilibrium condition $G$ shows that the aggregate response $\widehat{X}_t$ depends on two heterogeneous objects, the response of the individual policy $\hat{x}_t$ and the response of the distribution $\hat{\Omega}_t$. Differentiating the household optimality condition $F$ links $\hat{x}_t$ to the relevant aggregates $\widehat{Y}_t$; differentiating the distribution's law of motion links $\hat{\Omega}_t$ to $\hat{x}_t$. Stacking these relations over the horizon closes a sparse linear system in the aggregate responses $\{\widehat{X}_t\}_{t\geq 0}$. The subsections that follow compute the two pieces, the individual policy response $\hat{x}_t$ and the distribution response $\hat{\Omega}_t$, with the functional-analysis tools that make sense of them and the bridges to the state-space and sequence-space readings, and the last assembles them; the finite-element specialization Dynare 7 actually computes follows in Section 4.5.

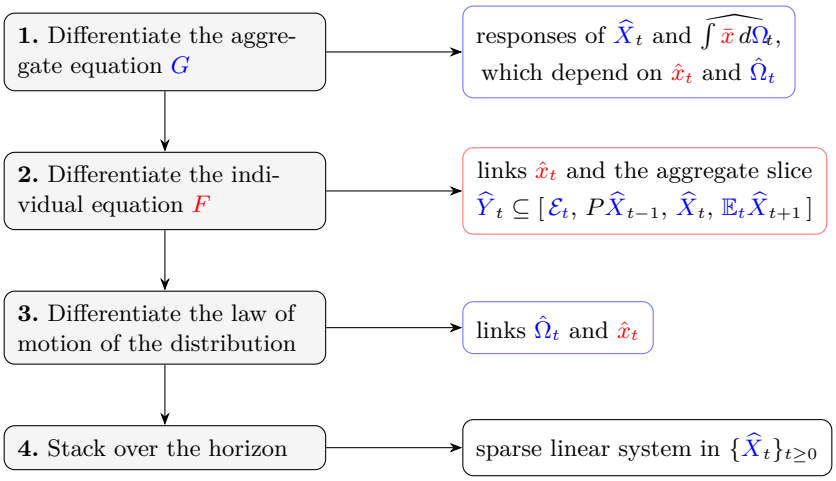


Figure 2: The first-order cascade, in the order the problem dictates, following the steps of Bhandari et al. [2023]. A one-time innovation perturbs the equilibrium map at the steady state. Differentiating the aggregate equation $G$ (§4.4) makes the aggregate responses and the summed heterogeneous response depend on the individual policy response $\hat{x}_t$ and the distribution response $\hat{\Omega}_t$; differentiating the individual equation $F$ (§4.2) fixes $\hat{x}_t$, and differentiating the distribution's law of motion (§4.3) fixes $\hat{\Omega}_t$; stacking over the horizon closes the sparse system, assembled in §4.4.

## 4.2 The heterogeneous block: the backward operator and the policy response

Collect the pointwise optimality conditions $F = 0$ over all $(\theta, a)$ into one residual that maps a triple of functions to a function,

$$\mathcal{F}\big(\bar{x}, \bar{x}^+, \bar{Y}\big)\langle \theta, a, \mathcal{Z}\rangle \;=\; F\big(\theta,\, a,\, \bar{x}(\theta, a, \mathcal{Z}),\, \mathbb{E}\big[\bar{x}^+ \,\big|\, \theta, a, \mathcal{Z}\big]\,,\, \bar{Y}(\mathcal{Z})\big) \;=\; 0,$$

where $\bar{x}, \bar{x}^+ \in \mathcal{X}$ are the current and next-period policy functions, $\bar{Y} \in \mathcal{Y}$ is the relevant-aggregates function of Section 2, the composition $\bar{x}^+\langle \theta, \theta', a, \mathcal{Z}, \mathcal{E}'\rangle = \bar{x}^+(\theta', \bar{a}(\theta, a, \mathcal{Z}), \mathcal{Z}')$, with $\theta'$ and $\mathcal{E}'$ the next-period

draws the expectation integrates, carries the policy into next period through its own predetermined-state output, and $\mathcal{F} : \mathcal{X}^2 \times \mathcal{Y} \to \mathcal{H}$ acts between spaces of functions on the joint state, each carried with the supremum norm: the policy deviations live in $\mathcal{X} = L^\infty$, the space every bound of this section is stated in. The conditional expectation is the one written out in Section 2.

The backward operator solves $\mathcal{F} = 0$ for the current policy in terms of the next-period policy and the relevant aggregates; the implicit function theorem furnishes it, and its derivatives, at the steady state. One standing hypothesis governs the whole section: the unfolded household problem, the square generalized equation obtained from the $\perp$-clauses as in Section 1, is *strongly regular* at every state in the sense of Robinson [1980]; Proposition 2 below unpacks what it delivers.

**Proposition 1.** *Let $J \equiv D_{\bar{x}}\mathcal{F}\big(\bar{x}^*, \bar{x}^*, \bar{Y}^*\big) : \mathcal{X} \to \mathcal{H}$. If $J$ is an isomorphism, then on a neighbourhood of the steady state there is a unique map $\mathfrak{x}$ with $\mathcal{F}\big(\mathfrak{x}(\bar{x}^+, \bar{Y}),\, \bar{x}^+,\, \bar{Y}\big) = 0$, and*

$$D_{\bar{Y}}\mathfrak{x} \;=\; -\,J^{-1}\,D_{\bar{Y}}\mathcal{F}, \qquad D_{\bar{x}^+}\mathfrak{x} \;=\; -\,J^{-1}\,D_{\bar{x}^+}\mathcal{F}.$$

*Proof.* The implicit function theorem [Abraham et al., 1988, Thm. 2.5.7] applies to $\mathcal{F}$ at the steady state, where $\mathcal{F}(\bar{x}^*, \bar{x}^*, \bar{Y}^*) = 0$: it furnishes the locally unique $\mathfrak{x}$ and, on differentiating the identity $\mathcal{F}(\mathfrak{x}, \bar{x}^+, \bar{Y}) = 0$ in $\bar{Y}$ and in $\bar{x}^+$ and solving, the two derivative formulas. It remains to verify the theorem's two hypotheses, that $\mathcal{F}$ is continuously Fréchet differentiable and that $J$ is an isomorphism, since the complementarity constraint makes neither automatic.

The differentiability is the subtler of the two, because the constraint forces the choice between the two senses of derivative of Appendix C: the Gâteaux (directional) derivative of Definition 8 and the stronger Fréchet derivative of Definition 9, and it is the Fréchet derivative that the theorem requires. Composing the residual with the kinked next-period policy $\bar{x}^+$ makes $\mathcal{F}$ only Gâteaux differentiable on the kink set, where the one-sided derivatives of $\bar{x}^+$ disagree and no single linear map approximates uniformly. What supplies the stronger derivative is the economics, not more analysis: at the steady state strict complementarity, the property that at each state off the $\bar{\Omega}^*$-null binding boundary $\Gamma$ of Definition 1 below one member of the complementarity pair, the multiplier or the bound's gap, stays strictly positive, fixes each node's regime, so no agent crosses the kink under a small aggregate shock, and on each fixed regime $\mathcal{F}$ is continuously differentiable, so by Proposition 14 its directional derivative is a Fréchet one with which it coincides. The policy responses below may then be read interchangeably as directional derivatives of the backward map or as its Fréchet derivative applied in a direction.

This regime-fixing is what makes both hypotheses hold despite the steady-state policy being only Lipschitz: with the constraint unfolded, the standing strong-regularity hypothesis, through Proposition 2 below, establishes both that $\mathcal{F}$ is continuously Fréchet differentiable and that $J$ is an isomorphism. The spaces here are Banach (the policy lives in $L^\infty$), so a continuous linear isomorphism is automatically a homeomorphism, hence boundedly invertible [Abraham et al., 1988, Thm. 2.2.16]; that bounded $J^{-1}$ is the property the policy-response recursion uses. □

Proposition 1 leaves the derivatives of $\mathcal{F}$ implicit; computing them turns its abstract formulas into the usable recursion below. Differentiating $\mathcal{F}$ through the inner arguments of $F$,

$$\begin{aligned} D_{\bar{x}}\mathcal{F}\cdot d\bar{x} \;&=\; \big(F_x + F_{x^+}\,\mathbb{E}\big[D_a\bar{x}^+ \,\big|\, \theta, a, \mathcal{Z}\big]\,p\big)\,d\bar{x}, \\ D_{\bar{x}^+}\mathcal{F}\cdot d\bar{x}^+ \;&=\; F_{x^+}\,\mathbb{E}\big[d\bar{x}^+ \,\big|\, \theta, a, \mathcal{Z}\big]\,, \\ D_{\bar{Y}}\mathcal{F}\cdot d\bar{Y} \;&=\; F_Y\,d\bar{Y}, \end{aligned}$$

where $F_x, F_{x^+}, F_Y$ are the partial derivatives of the pointwise residual and $p$ selects the predetermined-state component $\bar{a}$ out of $\bar{x}$ (lowercase, following Bhandari et al., 2023, and distinct from the aggregate selection $P$ of Section 2); Appendix D derives the three identities.

The map $\mathfrak{x}$ is the conceptual hinge between the state-space method of Bhandari et al. [2023] and the sequence-space method of Auclert et al. [2021]. It is exactly the backward step that Auclert et al. [2021] require the user to supply by hand, the rule that returns this period's policy from next period's policy and the current aggregates, and whose differentiation is the step their computational appendix treats with the most care [Auclert et al., 2021, App. C.1]. The state-space method of Bhandari et al. [2023] never asks for it, recovering the same derivatives implicitly from the household optimality conditions through $J^{-1}$.

A presentational difference is worth stating once. Auclert et al. [2021] work throughout with fully discretized objects, finite vectors and matrices on the grid, whereas here I keep the functionals that underlie those discrete objects and discretize only at the end. The two are the same construction read at two levels of resolution, set side by side across the three codebases in Appendix F.

**The policy response to an announced shock.** Announce, at date 0, a one-time perturbation $d\bar{Y}$ to the relevant aggregates at a single future date $T$, transitory in that the economy rests at the steady state again from date $T+1$ on, and write $\times_s$, in the cross notation of Bhandari et al. [2023], for the Fréchet coefficient of the policy response $s$ periods before the perturbation lands, so the response itself is $\times_s \cdot d\bar{Y}$. This coefficient is the object both methods compute, the state-space recursion of Bhandari et al. [2023] and the announcement experiments of Auclert et al. [2021] being two routes to it; in the latter, households learn at date 0 of a shock announced for $T$ that never materializes. The operator inverted in Proposition 1 is

$$J \;=\; F_x + F_{x^+}\,\mathbb{E}[D_a \bar{x}^* \,|\, \theta, a]\, p,$$

its conditional expectation the one of Section 2, taken at the steady state. Along a one-time aggregate shock the aggregate path is deterministic, so it enters as perfect foresight and carries no expectation, and the only expectation left is over next period's idiosyncratic state $\theta'$,

$$\mathbb{E}[g \,|\, \theta, a] \;=\; \int g(\theta', \bar{a}^*(\theta, a))\; d\pi(\theta' \mid \theta),$$

the same Koopman operator $\mathcal{K}$ that Section 4.3 builds.

**Lemma 1** (The policy-response coefficients)**.** *Under the hypothesis of Proposition 1, the policy-response coefficients are generated by the backward recursion*

$$\begin{aligned} \times_0 &= -J^{-1} F_Y, \\ \times_s &= -J^{-1} F_{x^+}\,\mathbb{E}[\times_{s-1} \,|\, \theta, a] \quad (s \geq 1). \end{aligned}$$

*Read in the sequence space this is the backward step of Auclert et al. [2021]: the period-$(T-s)$ policy response to the announced perturbation equals $\times_s \cdot d\bar{Y}$.*

*Proof.* Along the perturbed path the date-$t$ policy solves the household system $\mathcal{F}(\bar{x}_t, \bar{x}_{t+1}, \bar{Y}_t) = 0$, the time-invariant recursive form of Section 2, so by Proposition 1 it is the value of the backward operator, $\bar{x}_t = \mathfrak{x}(\bar{x}_{t+1}, \bar{Y}_t)$; that operator carries no date index, since $\mathcal{F}$ carries none. Write $d\bar{x}_t = \bar{x}_t - \bar{x}^*$ for the policy deviation and linearize at the steady state, where $\mathfrak{x}(\bar{x}^*, \bar{Y}^*) = \bar{x}^*$, to get the constant-coefficient backward map

$$d\bar{x}_t \;=\; D_{\bar{x}^+}\mathfrak{x} \cdot d\bar{x}_{t+1} \;+\; D_{\bar{Y}}\mathfrak{x} \cdot d\bar{Y}_t.$$

The perturbation lands only at $T$, so $d\bar{Y}_t$ equals $d\bar{Y}$ at $t = T$ and vanishes otherwise; and because the household problem is forward-looking, once the aggregates are back at the steady state the policy is too, $d\bar{x}_t = 0$ for $t > T$. Read the map backward from $T$. At the landing date the next-period policy is still at the steady state, so the only change is through the aggregates, $d\bar{x}_T = D_{\bar{Y}}\mathfrak{x} \cdot d\bar{Y}$; one period earlier the aggregates are unchanged and the only change is the anticipated next-period response, $d\bar{x}_{T-1} = D_{\bar{x}^+}\mathfrak{x} \cdot d\bar{x}_T$; and iterating, $d\bar{x}_{T-s} = D_{\bar{x}^+}\mathfrak{x} \cdot d\bar{x}_{T-s+1}$. Factoring the common direction $d\bar{Y}$ out of the response, $d\bar{x}_{T-s} = \times_s \cdot d\bar{Y}$, leaves the coefficient recursion $\times_0 = D_{\bar{Y}}\mathfrak{x}$, $\times_s = D_{\bar{x}^+}\mathfrak{x} \cdot \times_{s-1}$. Substituting the two derivatives of Proposition 1, $D_{\bar{Y}}\mathfrak{x} = -J^{-1}F_Y$ and $D_{\bar{x}^+}\mathfrak{x} = -J^{-1}F_{x^+}\,\mathbb{E}[\cdot \mid \theta, a]$, with the aggregate part of the expectation dropped under perfect foresight, gives the stated recursion.

The announcement reading is the same iteration under another name. Because the perturbation is transitory nothing happens after $T$, and by time-invariance the doubly-indexed family of responses, written $y_t^s$ in the notation of Auclert et al. [2021] for the date-$t$ response to a shock anticipated for date $s$, depends only on the gap $s-t$ to the shock: shifting both indices together leaves it unchanged, $y_t^s = y_{t-1}^{s-1}$, so the single sequence $\{\times_s\}$, indexed by that gap, carries all of them. □

This is the recursion Bhandari et al. [2023, Lemma 3] obtain for the same coefficient, so the two methods compute one object. Bhandari et al. [2023] recover $\times_s$ analytically through $J^{-1}$, and Auclert et al. [2021] read the same coefficient off a one-sided finite difference of the user-supplied backward step, a perturbation of size $h$ giving $\times_0 \cdot d\bar{Y} \approx (\mathfrak{x}(\bar{x}^*, \bar{Y}^* + h\,d\bar{Y}) - \mathfrak{x}(\bar{x}^*, \bar{Y}^*))/h$ and the higher coefficients by feeding $\times_{s-1}$ back in. The two evaluations agree as the grid is refined and $h \to 0$; the residual is the finite-difference and discretization error, not a conceptual gap. Two practical differences are worth flagging. First, the announcement route of Auclert et al. [2021] is exact only because the household problem is forward-looking: they note that an aggregate *lag* in the heterogeneous block breaks it, whereas the implicit-function reading of Bhandari et al. [2023], which Dynare 7 follows, admits such a lag. Second, the analytic route inverts $J$, so its accuracy is governed by the conditioning of $D_{\bar{x}}\mathcal{F}$, while the finite-difference route propagates derivatives through the backward step without an inverse but at the cost of supplying that step by hand.

**The overall response.** The lemma delivers the coefficients; the response to an arbitrary path of relevant aggregates follows from them by one unrolling. Iterating the constant-coefficient backward map of the proof, $d\bar{x}_t = D_{\bar{x}^+}\mathfrak{x} \cdot d\bar{x}_{t+1} + D_{\bar{Y}}\mathfrak{x} \cdot d\bar{Y}_t$, forward under the terminal condition $d\bar{x}_t \to 0$ that forward-lookingness imposes gives

$$d\bar{x}_t \;=\; \sum_{s\geq 0} (D_{\bar{x}^+}\mathfrak{x})^s \, D_{\bar{Y}}\mathfrak{x} \; d\bar{Y}_{t+s} \;=\; \sum_{s\geq 0} \times_s d\bar{Y}_{t+s},$$

the second equality by the coefficient recursion $\times_s = (D_{\bar{x}^+}\mathfrak{x})^s D_{\bar{Y}}\mathfrak{x}$ of the lemma. The response is the convolution of the relevant-aggregate path with the coefficients, and the sum runs over the *future* dates $t+s$: it is anti-causal, and it is so precisely because each backward step carries next period's policy, that is, because the household problem is forward-looking. The lemma's one-time perturbation is the special case of this convolution in which $d\bar{Y}_\tau$ equals $d\bar{Y}$ at a single future date $\tau = T$ and vanishes otherwise; the whole response path is then read straight off the coefficients,

$$d\bar{x}_\tau \;=\; \begin{cases} \times_{T-\tau} d\bar{Y}, & \tau \leq T, \\ 0, & \tau > T, \end{cases}$$

each date $\tau$ reacting through the coefficient at its own horizon $T-\tau$ to the shock, and none reacting once the shock has passed. This shock is genuine, merely transitory: the lemma reads its coefficients off this entire path, and nothing here is fake. Fakeness enters only to isolate one coefficient, as Auclert et al. [2021] do. To read $\times_{T-t}$ on its own, one announces the shock at date $t$, so that the date-$t$ policy reacts by $\times_{T-t} d\bar{Y}$; the announcement, new at $t$, leaves the earlier dates unmoved, and revealing at the end of the period that the shock was fake sets $d\bar{x}_\tau = 0$ for $\tau \geq t+1$, so the lone surviving reaction is the coefficient itself.

**The constraint at first order.** The recursion runs over every state uniformly; the complementarity constraint enters not through a binding boundary I must find, the set $\Gamma$ defined next, but through the way Dynare 7 unfolds it, and settling that is what makes the inversion of $J$ legitimate.

**Definition 1** (Regimes and the binding boundary)**.** Unfold the bounds as Section 1 does and label them $b = 1, \ldots, B$. A *regime* is a region of $\Theta \times \mathcal{A}$ on which the active set

$$A(\theta, a) \;=\; \big\{\, b \;:\; \text{bound } b \text{ binds at } (\theta, a) \,\big\}$$

is constant, so there are at most $2^B$ regimes. For a single one-sided bound, with $\mu^*$ the bound's steady-state multiplier and $g^* = \bar{a}^* - \underline{a} \geq 0$ its gap, the two regimes are the *interior regime* $\{(\theta,a) : \mu^* = 0,\ g^* > 0\}$ and the *binding regime* $\{(\theta,a) : g^* = 0,\ \mu^* > 0\}$. The difference of the two members of the complementarity pair,

$$\phi \;\equiv\; \mu^* - g^* \;=\; \begin{cases} \mu^* \;>\; 0 & \text{on the binding regime,} \\ -\,g^* \;<\; 0 & \text{on the interior regime,} \end{cases}$$

sorts the regimes by its sign, and the *binding boundary* is its zero set,

$$\Gamma \;=\; \{\phi = 0\} \;=\; \big\{(\theta,a) \;:\; \mu^* = 0 \;\text{ and }\; \bar{a}^* = \underline{a}\big\},$$

the failure set of strict complementarity: at a point of $\Gamma$ feasibility makes $\mu^*$ and $g^*$ equal and complementarity annihilates their product, so both vanish at once; that the zero set is also the common frontier of the two regimes is part of what Section 5.2 establishes. The function $\phi$ introduces no new object: $\mu$ is the multiplier coordinate the unfolding adds to $\bar{x}$ and $g^*$ is the slack of the $\perp$-clause row of $F$, so $\phi$ is read directly off the complementarity pair the `.mod` file declares, and no shape of the policy, monotone or otherwise, enters the definition. With several bounds each contributes its own boundary $\Gamma_b$, the zero set of its own $\phi_b$, and $\Gamma = \bigcup_b \Gamma_b$, a finite union. Section 5.2 shows that, under two policy-side clauses, the slice of each $\Gamma_b$ in a Markov row $\theta$ is a hypersurface of codimension one in $\mathcal{A}$ (Definition 17), and equips it with the regularity the higher orders consume. At $d_a = 1$, when the slice is a single crossing point, the running example's case, I write

$$\Gamma \cap \big(\{\theta\} \times \mathcal{A}\big) \;=\; \{\theta\} \times \{a_\Gamma(\theta)\},$$

the position at which the multiplier has just fallen to $\mu^*(\theta, a_\Gamma(\theta)) = 0$; this labeled scalar case is the only context in which the symbol $a_\Gamma$ appears.

**Proposition 2.** *Unfold the complementarity constraint as Section 1 does, taking first a single one-sided bound: one coordinate of $\bar{x}$ is held to a one-sided bound through the $\perp$-clause, its single multiplier $\mu$ carried as a further coordinate of $\bar{x}$ and the complementary slackness, the product of $\mu$ with the bound's gap, carried as a row of $\mathcal{F}$, so the steady-state household problem is a square generalized equation; the running example's asset bound $\bar{a}^* \geq \underline{a}$, with slackness $\mu\,(\bar{a}^* - \underline{a}) = 0$, is this case. Assume it is strongly regular at every $(\theta, a)$ in the sense of Robinson [1980], and that the binding boundary $\Gamma$ of Definition 1, where strict complementarity fails, has $\bar{\Omega}^*(\Gamma) = 0$. Then*

1. *off $\Gamma$ the policy response is the classical per-regime derivative, $J$ is a bounded multiplication operator with $J^{-1}$ bounded, and the response stays bounded, $\times_s \in L^\infty(\Theta \times \mathcal{A})$ for every $s$, with jumps across the binding boundary $\Gamma$, the two one-sided limits differing but staying finite, and no Dirac;*
2. *on the binding regime the constrained variable is pinned, $p \times_s = 0$ for every $s$, the anticipation surviving only on the multiplier coordinate;*
3. *every aggregate $\langle \bar{\Omega}^*, \times_s \rangle$ is a well-defined integral, unambiguous under $\bar{\Omega}^*(\Gamma) = 0$.*

*Proof. Solution map.* Unfolded, the household problem is a parametric generalized equation; by strong regularity [Robinson, 1980] its solution $\bar{x}^*$ is single-valued and locally Lipschitz, classically (Fréchet) differentiable wherever strict complementarity holds, that is off $\Gamma$, and directionally differentiable on the $\bar{\Omega}^*$-null set $\Gamma$ [Facchinei and Pang, 2003, §5]. Off $\Gamma$ the active set is locally constant, so each agent solves a smooth square system, the interior regime or the binding regime, and the implicit function theorem applies on that regime.

*$J^{-1}$ is bounded.* $J = F_x + F_{x^+}\,\mathbb{E}[D_a \bar{x}^* \mid \theta, a]\,p$ acts by multiplication: at each state it is the matrix $J(\theta, a)$ on the value there. Its coefficients are bounded, $D_a \bar{x}^*$ being an $L^\infty$ jump and not a Dirac, so $\mathbb{E}[D_a \bar{x}^* \mid \theta, a]$ is an integral of a bounded function; and strong regularity, holding up to $\Gamma$ from either regime, makes $\|J(\theta, a)^{-1}\|$ bounded almost everywhere. So $J^{-1}$ is a bounded operator on $L^\infty$, and $J$ need not vary smoothly in the state, only be invertible at each one.

*The recursion stays bounded.* With $\mathcal{K}$ the conditional-expectation operator of Section 4.3, nonexpansive on $L^\infty$ (an averaging, $\|\mathcal{K}\varphi\|_\infty \leq \|\varphi\|_\infty$), the recursion $\times_0 = -J^{-1} F_Y$, $\times_{s+1} = -J^{-1} F_{x^+}\,\mathcal{K} \times_s$ composes bounded operators, so $\times_s \in L^\infty$ for every $s$. A jump in $J$ across the binding boundary $\Gamma$ gives $\times_s$ a bounded jump, never an impulse, and $\mathcal{K}$ averages the previous response, so nothing sharpens with $s$.

*The constrained variable.* On the binding regime the constrained variable's own row is $\bar{a}^* = \underline{a}$, which carries no lead; differentiating it gives $p \times_s = 0$ at every horizon, since $\underline{a}$ is a parameter. The lead has nowhere to land but the multiplier's row, the Euler residual, so the multiplier coordinate alone responds, with no appeal to continuity and no boundary condition.

*Aggregation.* Each downstream object integrates $\times_s \in L^\infty$ against the finite measure $\bar{\Omega}^*$; the value of $\times_s$ on the null $\Gamma$ enters only through an atom of $\bar{\Omega}^*$ on $\Gamma$, excluded by $\bar{\Omega}^*(\Gamma) = 0$. □

*Two-sided bounds and several constraints.* The single one-sided bound is no loss of generality. An upper bound mirrors a lower one, the coordinate replaced by its negative, so it carries its own non-negative multiplier and its own slackness row, active on a region disjoint from the lower bound's; a two-sided bound is the two together, with the multipliers $\mu^-, \mu^+$ and slacknesses $\mu^-(x - \underline{x}) = 0$, $\mu^+(\overline{x} - x) = 0$. Any number of $\perp$-constraints, on one or several coordinates of $\bar{x}$, stack the same structure independently, and Definition 1 already carries the general vocabulary: the binding boundary $\Gamma$ is the finite union of the pairwise boundaries between adjacent regimes, still $\bar{\Omega}^*$-null, and the per-regime argument runs unchanged on each regime, every multiplier inert wherever its own bound is slack.

The construction needs no generalized function, and this is where it parts from Bhandari et al. [2023]. The knife-edge $\Gamma$ is the set they single out in their Appendix B.1, the only states at which the multiplier and the constrained variable both vanish. Their single Dirac [Bhandari et al., 2023, Appendix B.2, Claim 5], $\bar{x}^*{}_{aa} = \mathring{x}_{aa} + \delta\, x_a^\Delta$, their $\mathring{x}_{aa}$ the per-regime classical part $\{D_a^2\,\bar{x}^*\}$ and their $x_a^\Delta$ the slope jump $[D_a \bar{x}^*]$ of Section 5.2, multiplies the jump in the *derivative* and so appears only in a second derivative; their first-order statement from the same Claim, $x_Z = \mathring{x}_Z$, is exactly the $L^\infty$ response above. I form only the first derivative $\times_s$, and on the distribution side the weak first derivative of Section 4.3, and stop, so the jump in $D_a \bar{x}^*$ stays a bounded coefficient inside $J$. Bhandari et al.'s own code makes the same first-order reduction: on the binding regime it replaces the Euler by the constraint, so the forward block vanishes and the response on the binding regime is $\times_s = 0$ for $s \geq 1$, the boundary value of their Lemma 3, $\times_s(\kappa(\theta), \theta) = 0$, their $\kappa(\theta)$ the position $a_\Gamma(\theta)$ of Definition 1, recovered here as the projection onto a policy that omits the multiplier. At first order the constraint's mark is not on the individual response but on the distribution side, the atom of Section 4.3.

### 4.3 The aggregate block: the transfer operator and its weak derivative

The policy response feeds the aggregate variables through the distribution, and it is on the distribution side that the singular mass the constraint piles up must be handled with care. I introduce only the machinery this needs. The aggregate equations never use the distribution $\bar{\Omega}^*$ pointwise; they use it through the cross-sectional integrals $\int \bar{y}^* \, d\bar{\Omega}^*$ of Section 3.1, that is, through its pairing with functions. So take a test space $V$ of functions on $\Theta \times \mathcal{A}$ on which point evaluation, the rule $\varphi \mapsto \varphi(\theta, a)$ returning a function's value at a point, is bounded, and let $V'$ be its dual; the distribution acts by the pairing $\langle \mu, \varphi \rangle = \int \varphi \, d\mu(\theta, a)$, and a point mass at the constraint is then a legitimate element of $V'$. Which test space $V$ does this is stated next; the main text stays at the level of why it is the right object, with the construction in Appendix C.

**Definition 2** (Layers and forward orbits)**.** For the bound $a^i \geq \underline{a}^i$ on the $i$-th continuous coordinate:

- the set $F_i \equiv \{a \in \mathcal{A} : a^i = \underline{a}^i\}$ collects the positions at which the bound's value is attained, a set of codimension one in $\mathcal{A}$ (Definition 17); an intersection of such sets has codimension the number of bounds it stacks;
- a *layer* is a measure $\sigma^*$ carried by such a set $F_i$, by an intersection of them, or by a set reached from one along the forward orbits below, with a density on its carrier in a fractional class of Appendix C, $W^{1/p,\,p'}(F_i)$ for the codimension-one carrier $F_i$;
- the set $\{(\theta, a) \in \Theta \times \mathcal{A} : \bar{a}^{*i}(\theta, a) = \underline{a}^i\}$ collects the states whose households the policy pins this period.

Three distinct objects, said once: the set $\{\bar{a}^{*i} = \underline{a}^i\}$ is the *source* of the piled-up mass, the set $F_i$ its *destination*, and the binding boundary $\Gamma$ of Definition 1 the regime-switching *threshold* between regimes. Finally, the *forward orbit* of a state,

$$\mathcal{O}(\theta, u) \;\equiv\; \bigcup_{t \geq 0} \big\{(\theta_t, u_t) \;:\; u_{t+1} = \bar{a}^*(\theta_t, u_t),\; \pi(\theta_{t+1} \mid \theta_t) > 0,\; (\theta_0, u_0) = (\theta, u)\big\},$$

collects, date by date, the countably many states its households can reach under the steady-state policy and the positive-probability Markov moves.

**Assumption 1** (Regularity and test space)**.** The predetermined idiosyncratic state splits into a finite Markov component $\theta$ and the $d_a$ continuous predetermined coordinates $a \in \mathcal{A} \subseteq \mathbb{R}^{d_a}$. At the steady state strict complementarity holds off the $\bar{\Omega}^*$-null binding boundary $\Gamma$, so each agent's regime is fixed; the policy $\bar{a}^*$ is continuously differentiable on each regime and Lipschitz throughout, with non-differentiability confined to the binding boundary $\Gamma$ (Proposition 2). The test space is $V = \ell^2(\Theta) \otimes W^{1,p}(\mathcal{A})$ with $p > d_a$. Four further clauses discipline the transport below.

1. The stationary distribution decomposes as

$$\bar{\Omega}^* \;=\; \rho^* \;+\; \sum_{f=1}^{N_\sigma} \sigma_f^* \;:$$

an interior part with a density $\rho^* \in L^{p'}$, and finitely many singular parts $\sigma_f^*$, each a measure carried by $\{\theta_f\} \times F_f$, with $F_f$ of codimension $c_f \in \{1, \ldots, d_a\}$ (Definition 17) and the density on the carrier in the fractional class $W^{c_f/p,\,p'}(F_f)$ of Appendix C. The invariance $\Lambda \bar{\Omega}^* = \bar{\Omega}^*$ disciplines the family before anything is assumed: the Lebesgue decomposition being unique, its singular parts on both sides must balance,

$$\sum_f \sigma_f^* \;=\; \sum_f \Lambda\, \sigma_f^* \;+\; \big(\Lambda \rho^*\big)_{\text{sing}},$$

where each $\Lambda \sigma_f^*$ is singular (the per-regime policy is Lipschitz, so a carrier's image is again a null set) and $(\Lambda \rho^*)_{\text{sing}}$, the mass of the interior part that the pinning regimes pile up, is carried by the intersections of the sets $\{a : a^i = \underline{a}^i\}$ of Definition 2 and by nothing else, clause 2 forbidding interior flats of the policy. The balance forces the carrier family to be closed under the steady-state dynamics, each carrier's image lying in the family's union, and to contain every bound-set intersection onto which the interior part piles up. What the clause assumes is therefore twofold:

*finitely many carriers suffice*, the counterpart of the finitely many free-position mass points of Bhandari et al. [2023, Assumption 2(b)], the forward closure being countable in general; and *every carrier lies in a bound-set intersection or in one of its date-by-date forward images*, so that all singular mass originates at the bounds, the one configuration excluded being a family that feeds itself in the interior without ever touching a bound, which takes an exact calibration coincidence to produce. The codimension-$d_a$ carriers are points, and those $\sigma_f^*$ are the *atoms*, $m_n\,\delta_{(\theta_n,u_n)}$ with masses $m_n$ at states $(\theta_n, u_n)$; with a single predetermined coordinate ($d_a = 1$) they are the only singular parts. The codimension-one carriers are the layers of Definition 2, with the density class $W^{1/p,\,p'}(F_f)$ stated there.

2. On each regime, the block of the policy Jacobian $D_a\bar{a}^*$ along the coordinates the regime leaves free is non-degenerate, its determinant bounded away from zero.

3. The stationary distribution $\bar{\Omega}^*$ puts no mass on the upper boundary of $\mathcal{A}$.

4. The forward orbits of the singular carriers, $\mathcal{O}(\theta_n, u_n)$ for the atoms and the date-by-date images of $\{\theta_f\} \times F_f$ for the carriers of lower codimension, avoid the binding boundary $\Gamma$.

Under Assumption 1 this frame is admissible: $V \hookrightarrow C^0$ (Theorem 1), so the constraint atoms, the corner where all $d_a$ bounds bind at once included, are legitimate elements of the dual $V' = \ell^2(\Theta) \otimes W^{-1,p'}(\mathcal{A})$ (Proposition 16); the kinked policy lies in $V$; and the Koopman operator below carries $V$ into $V$, so its adjoint $\Lambda$ acts on $V'$. That this frame serves every model in the class, and that no tighter, Hilbert space does once two bounds co-bind on a sloped frontier, is the content of Appendix C: the admissibility is Lemma 12, and the ladder of test spaces (Table 1, Figure 4) shows why the Hilbert structure has to go.

The four clauses serve one purpose each. The decomposition of clause 1 restricts, at $d_a = 1$, to the world of Bhandari et al. [2023, Assumption 2], an absolutely continuous part plus finitely many mass points: it separates the mass whose response stays in the dual $V'$ from the finitely many singular pieces whose own motion must be tracked by hand, through the finite span $D$ of dipoles built after Proposition 3; and since $p$ may be taken as large as desired above $d_a$, the conjugate exponent $p'$ sits as close to one as desired, so the density class asked of $\rho^*$ is barely more than integrability. Clause 2 makes the interior transport preserve densities and the chain rule hold almost everywhere, pre-images of null sets being null: Lemma 3 and the proof of Corollary 1 both draw on it. The clause is the price of my insistence that every pairing be an ordinary integral wherever possible, and economically it is cheap: it asks that, on the interior, no direction of today's asset vector loses all influence on tomorrow's. Clause 3 and the vanishing of the constrained response wherever its bound binds (Proposition 2) remove boundary terms: the derivatives below are all built by integration by parts, the test functions of $V$ do not vanish at the boundary of $\mathcal{A}$, and these two facts are what make every boundary term drop, a no-flux property I use silently in every integration by parts of this section and record here once; clause 3 is asymmetric because the lower boundary is where the bounds sit, its mass is exactly clause 1's singular part, and no-flux there is the vanishing constrained response itself, so only the upper boundary needs a clause of its own. Clause 4 keeps the dipole pairings $\langle \dot{\delta}_{(\theta,u)}[v],\ \bar{x}^* \rangle = v \cdot \nabla_a \bar{x}^*(\theta, u)$ well-defined at every horizon: the forward orbits it speaks of are the states an atom's households visit as the steady-state dynamics carry them on, the binding boundary $\Gamma$ collects the surfaces where the policy's gradient is undefined, and the clause asks that the two never meet. Appendix C defines both terms and explains why the requirement is mild: violating it takes one of countably many knife-edge coincidences, and a perturbation of the calibration removes them.

The distribution's law of motion in Section 2 is linear in the distribution, and I realize it as an operator by duality rather than by writing a kernel. In one step the idiosyncratic state $(\theta, a)$ moves to $(\theta', \bar{a}^*(\theta, a))$: the predetermined coordinate is carried deterministically by the policy, and the exogenous shock jumps to $\theta'$ drawn from $\pi(\cdot \mid \theta)$. The conditional-expectation operator sends a test function to its expected value one step ahead,

$$(\mathcal{K}\varphi)(\theta, a) \;=\; \mathbb{E}\big[\, \varphi(\theta', \bar{a}^*(\theta, a)) \,\big|\, \theta, a \,\big] \;=\; \int \varphi(\theta', \bar{a}^*(\theta, a))\; d\pi(\theta' \mid \theta),$$

the conditional expectation Dynare 7 forms at every node.

This conditional expectation is the Koopman operator of the idiosyncratic dynamics, and the name is worth earning, because it can look foreign from either side. A dynamical system can be read through its trajectories, each state carried to the next, or through its observables, the functions of the state

whose values ride along as the state moves. The Koopman operator is the second reading made linear: it carries an observable $\varphi$ to the one whose value at today's state is the expected value of $\varphi$ at tomorrow's. When tomorrow's state is a deterministic map $S$ of today's, that expected value is $\varphi(S(\cdot))$ and the operator is composition with $S$, the form the dynamical-systems literature writes [Lasota and Mackey, 1994, Def. 3.3.1]; here tomorrow's shock is genuinely random, so the same operator is a conditional expectation, and the deterministic composition is only the special case in which the transition is a point mass. To the economist $\mathcal{K}$ is nothing new, it is the expectation already formed at every node; the name marks the decision to treat that expectation as one linear operator on functions rather than a family of trajectory averages. The gain is that $\mathcal{K}$ is linear in $\varphi$ although the dynamics are not, so it has an adjoint; and that adjoint, the lemma below shows, is exactly the forward law of motion of the distribution.

Independently of that, the distribution has its own law of motion. If the cross-section is $\mu$ today, tomorrow's is the distribution of the next state $(\theta', a')$: the policy carries the current $(\theta, a) \sim \mu$ to $a' = \bar{a}^*(\theta, a)$, and the shock moves on to $\theta'$ under its transition $\pi(\cdot \mid \theta)$. Write $\Lambda\mu$ for that distribution; the map $\mu \mapsto \Lambda\mu$ is the transfer operator, the Perron–Frobenius operator of the dynamical-systems dictionary.

**Lemma 2** (Koopman and Perron–Frobenius operators)**.** *The transfer operator $\Lambda$ and the Koopman operator $\mathcal{K}$ are adjoints: for every distribution $\mu \in V'$ and test function $\varphi \in V$,*

$$\langle \Lambda\mu,\ \varphi \rangle \ = \ \langle \mu,\ \mathcal{K}\varphi \rangle,$$

*so $\Lambda = \mathcal{K}^*$. The stationary distribution is the fixed point, $\bar{\Omega}^* = \Lambda\bar{\Omega}^*$.*

*Proof.* This is the law of iterated expectations. Pairing $\Lambda\mu$ against $\varphi$ is the expected value of $\varphi$ at the next state, computed either directly over that state's law $\Lambda\mu$ or by conditioning first on the current state. Writing the second,

$$\langle \Lambda\mu,\ \varphi \rangle \ = \ \iint \varphi\big(\theta', \bar{a}^*(\theta, a)\big)\, d\pi(\theta' \mid \theta)\, d\mu(\theta, a) \ = \ \int (\mathcal{K}\varphi)\, d\mu \ = \ \langle \mu,\ \mathcal{K}\varphi \rangle,$$

the inner integral over $\theta'$ being $(\mathcal{K}\varphi)(\theta, a)$. The policy is only evaluated, inside $\varphi$, and never inverted, so the kinked, mass-collapsing $\bar{a}^*$ raises no difficulty. That $\mathcal{K}$ carries $V$ into $V$ and is bounded is Lemma 12, so the adjoint $\Lambda = \mathcal{K}^*$ is bounded on the whole dual $V'$, moving every distribution there. Finally $\bar{\Omega}^*$ is invariant under the one-step transition, hence $\Lambda\bar{\Omega}^* = \bar{\Omega}^*$. □

On the grid the two operators are one transition matrix read two ways, the expectation contracting it along its rows and the push-forward along its columns; the names Koopman and Perron–Frobenius are those of Lasota and Mackey [1994, Defs. 3.3.1, 3.2.3, (5.7.3)]. The transfer operator depends on the policy only through the predetermined-state map $\bar{a}^*$, and I differentiate it in that direction by pairing against a test function, so the chain rule lands the derivative on the smooth $\varphi$ and never on a measure.

Every derivative on the distribution side reaches the policy through $\mathcal{K}$ alone, so I record once and for all how $\mathcal{K}$ differentiates. Two derivatives can hit the expectation $(\mathcal{K}\psi)(\theta, a) = \int \psi(\theta', \bar{a}^*(\theta, a))\, d\pi(\theta' \mid \theta)$: the state $a$ can move, carrying the landing point $\bar{a}^*(\theta, a)$ with it, or the policy itself can be perturbed, moving the landing point at a fixed state. In both cases the only object differentiated is the test function, evaluated where the agent lands, and the weight in front is the slope of the landing map in the direction that moved.

**Corollary 1** (Differentiating the Koopman operator)**.** *Under Assumption 1, for every Lipschitz test function $\psi$ and every bounded direction $da = (da^j)_j$,*

$$\partial_{a^j}\big(\mathcal{K}\psi\big) \ = \ \sum_i \partial_{a^j}\bar{a}^{*i}\ \mathcal{K}\big(\partial_{a'^i}\psi\big) \quad \textit{almost everywhere,} \qquad \big(D_{\bar{a}^*}\mathcal{K}\cdot da\big)\,\psi \ = \ \sum_j da^j\ \mathcal{K}\big(\partial_{a'^j}\psi\big).$$

Each product on the right is read as zero wherever its weight $\partial_{a^j}\bar{a}^{*i}$, respectively $da^j$, vanishes; that convention settles the formulas on the sets $\{\bar{a}^{*i} = \underline{a}^i\}$, where the pinned coordinates make the composition constant and the gradient of $\psi$ may be undefined at the landing point $\bar{a}^*(\theta, a)$. The proof is in Appendix C, next to Lemma 12, whose composition argument it makes explicit. The restriction to Lipschitz test functions costs nothing: the functions this section ever pushes through $\mathcal{K}$ are the steady-state policy $\bar{x}^*$ and its own $\mathcal{K}$-iterates, all Lipschitz, and the operator identities below extend from this dense class to all of $V$ by continuity.

**Fields, flux, and divergence.** The derivative about to be computed is best read with three words of fluid mechanics, fixed here once. A vector field on the state space assigns an arrow to each state; the field of interest attaches to the household at $(\theta, a)$ the direction and distance $da(\theta, a)$ by which a perturbed policy moves its landing point, weighted by the mass $\bar{\Omega}^*$ sitting there, so the arrows carry mass. The flux of such a mass-carrying field through a piece of surface is the mass its arrows push across it; the divergence at a point is the net outflow per unit volume, what leaves minus what arrives. Mass that is neither created nor destroyed obeys the continuity equation of fluid mechanics: writing $\rho$ for a generic density and $v$ for the velocity field transporting it, the change of the density is minus the divergence of the mass flow $\rho\, v$,

$$\partial_t \rho \;+\; \nabla \cdot (\rho\, v) \;=\; 0,$$

arrows converging on a point piling mass up there, arrows spreading out draining it. The proposition below is the discrete-time face of this identity: perturbing the policy displaces each household's landing point, and the first-order change of the pushed-forward distribution is minus the divergence of the mass-weighted displacement.

**Proposition 3.** *Perturbing the policy moves where each agent lands. Holding the steady state $\bar{\Omega}^*$ fixed, carry the mass-weighted policy direction forward by the transfer operator: the displacement operator has $j$-th component $(\mathsf{M} \cdot da)^j = \Lambda\big(da^j\, \bar{\Omega}^*\big)$, the transfer operator applied to the signed measure $da^j\, \bar{\Omega}^*$, which pairs with a test function $\psi$ by*

$$\big\langle (\mathsf{M} \cdot da)^j,\ \psi \big\rangle \;=\; \iint \psi(\theta', \bar{a}^*(\theta, a))\, da^j(\theta, a)\, d\pi(\theta' \mid \theta)\, d\bar{\Omega}^*(\theta, a).$$

*Then, for every bounded direction $da$ with $da^i = 0$ on $\{\bar{a}^{*i} = \underline{a}^i\}$ for each bound $i$, as every policy response has (Proposition 2), the policy-perturbation of the pushed-forward steady state, at fixed $\bar{\Omega}^*$, is the weak divergence of the displacement,*

$$\big(D_{\bar{a}^*} \Lambda \cdot da\big)\, \bar{\Omega}^* \;=\; -\nabla_{a'} \cdot \big(\mathsf{M} \cdot da\big).$$

*Proof.* The derivative on the left is the weak-* directional derivative: tested against a test function $\varphi$, it is by definition the $h$-derivative at zero of $\big\langle \Lambda[\bar{a}^* + h\, da]\, \bar{\Omega}^*,\ \varphi \big\rangle$; every later use of it is such a pairing, so no stronger notion is ever needed. Fix a Lipschitz $\varphi$. The computation below establishes the identity against every such $\varphi$; the decomposition of Assumption 1 then splits the right side, the interior part extending to all of $V$ by density, its divergence a bounded functional by the regularity discussed after the proof, while the finitely many atom parts remain functionals on the Lipschitz class, the dipole span $D$ built below. Differentiating the adjoint relation of Lemma 2 in the policy moves the derivative onto the Koopman operator, where Corollary 1 evaluates it:

$$\big\langle \big(D_{\bar{a}^*} \Lambda \cdot da\big) \bar{\Omega}^*,\ \varphi \big\rangle \;=\; \big\langle \bar{\Omega}^*,\ \big(D_{\bar{a}^*} \mathcal{K} \cdot da\big)\varphi \big\rangle \;=\; \sum_j \big\langle da^j\, \bar{\Omega}^*,\ \mathcal{K}\big(\partial_{a'^j} \varphi\big) \big\rangle \;=\; \sum_j \big\langle \Lambda\big(da^j \bar{\Omega}^*\big),\ \partial_{a'^j} \varphi \big\rangle.$$

The first equality holds because the pairing with $\bar{\Omega}^*$ and the $h$-limit exchange by dominated convergence, the difference quotient of the Lipschitz $\varphi$ being bounded by its Lipschitz constant times $\|da\|_\infty$; the second is Corollary 1 in the policy direction; the third is the adjoint relation again, now applied to the signed measure $da^j \bar{\Omega}^*$, and it holds for the integrand $\partial_{a'^j} \varphi$ because the adjoint identity is Fubini on the double integral and asks only integrability, not membership in $V$. The right-hand side is $\sum_j \big\langle (\mathsf{M} \cdot da)^j,\ \partial_{a'^j} \varphi \big\rangle$ by the definition of the displacement, and by the definition of the weak divergence [Brenner and Scott, 2008, (1.2.4)], applied componentwise, with the boundary term dropped by the no-flux clause of Assumption 1, it equals $-\big\langle \nabla_{a'} \cdot (\mathsf{M} \cdot da),\ \varphi \big\rangle$. No derivative ever reaches the policy or the measure: the only object differentiated, throughout, is the test function, through $\mathcal{K}$. □

The object $\mathsf{M} \cdot da$ is the mass-weighted displacement of the agents, pushed forward to each destination, and the negative divergence is the transport identity for the perturbed flow. Assumption 1's decomposition of $\bar{\Omega}^*$ splits the displacement into pieces of different natures, and each is handled by the tool built for it. The displacement of the interior part, with components $\Lambda(da^j \rho^*)$, is absolutely continuous: the direction has $da^i = 0$ on each set $\{\bar{a}^{*i} = \underline{a}^i\}$, the only states a push-forward collapses, and off them the policy is non-degenerate on each regime (clause 2), so the image of an $L^{p'}$ density is again an $L^{p'}$ density. Its weak divergence pairs with $\nabla\varphi \in L^p$ as an ordinary Hölder integral, asking no continuity of $\nabla\varphi$: this part lies in $V'$.

The displacement of a singular part $\sigma_f^*$ splits along its carrier. Displaced tangentially, the only directions open where the carrier's own coordinates bind, it stays a measure on its face, its density shifted, and the pairing consumes the tangential gradient of the trace (Theorem 4 and its fractional refinement in Appendix C). Displaced normally, at a row where a carrier coordinate unbinds, the mass leaves the carrier, and per unit mass the $h$-limit is the normal derivative of the carrier's Dirac layer, no element of $V'$: the normal derivative of a $W^{1,p}$ test function carries no trace on the carrier. That is the statement on the first grade $V = V_1$; the higher grades of Section 5.2 pair one degree deeper, one grade per degree. For an atom the carrier is a point and every direction is normal: its displacement is a transported point mass, the weight $m_n\, da(\theta_n, u_n)$ at the destinations $(\theta', \bar{a}^*(\theta_n, u_n))$, and the paragraph below gives it its proper, finite home. One space, $W^{1,p}$ with $p > d_a$, thus serves the aggregation pairing $\langle \bar{\Omega}^*, \varphi \rangle$ through continuity and the interior displacement through integrability; what it cannot serve, by design of no admissible space, is the motion of the singular mass itself, which is bookkeeping on finitely many carriers, not analysis.

**The motion of the singular mass.** The atoms of Assumption 1 are legitimate members of the dual $V'$; it is their *motion* that leaves it, and the Krusell–Smith example, the $d_a = 1$ instance in which every singular carrier is a point, tells the whole story. The steady state parks an atom at the borrowing limit $\underline{a}$: a positive mass of households, spread over the Markov rows $\theta$, holds exactly $\underline{a}$. A row still constrained today does not move: its policy response vanishes on the binding regime (Proposition 2), the atom stays parked at the bound, and a parked atom raises no difficulty, it is a Dirac in $V'$. The interesting row is the one that arrived at the bound yesterday and draws an income today at which the constraint no longer binds: it saves, so the whole point mass $m$ it carries leaves the bound for a single interior landing point $q = \bar{a}^*(\theta, \underline{a})$. Now perturb the policy by $h\, da$. The landing point shifts to $q + h\, v$ with $v = da(\theta, \underline{a})$, so tomorrow's cross-section carries, at each Markov row it mixes into, a point mass at a shifted position; per unit of that mass, the response of the distribution is the $h$-derivative

$$\lim_{h \to 0} \frac{\delta_{q + h\, v} \;-\; \delta_q}{h} \;=\; \dot{\delta}_q[v],$$

two spikes closing in on each other, the mass arriving next to the mass leaving: the dipole of electrostatics. I state the general object as a definition.

**Definition 3** (Dipole)**.** For a state $(\theta, u) \in \Theta \times \mathcal{A}$ and a direction $v \in \mathbb{R}^{d_a}$, the *dipole* $\dot{\delta}_{(\theta,u)}[v]$ is the linear functional

$$\big\langle \dot{\delta}_{(\theta,u)}[v],\ \psi \big\rangle \;=\; v \cdot \nabla_a \psi(\theta, u),$$

up to sign the distributional derivative of the Dirac $\delta_{(\theta,u)}$ [Friedlander and Joshi, 1998, §2.2, eq. (2.2.4)].

Nothing else can appear: a distribution supported at a single point is a finite combination of the point mass and its derivatives [Friedlander and Joshi, 1998, Thm. 3.2.1], and at first order only the mass and its dipoles survive. The dipole is not in $V'$. Reading it against a test function takes the *slope* of the test function at one point, where the Dirac asked only its value; a $W^{1,p}$ test function has a value everywhere, its continuous representative, but a gradient only almost everywhere, so the slope at the single point $q$ is not a meaningful question on $V$, and Appendix C shows that no admissible test space makes it one. So the dipoles are carried by hand, and the bookkeeping is finite. Differentiating the push-forward of one atom in the direction $da$,

$$\frac{d}{dh}\bigg|_0 \Lambda[\bar{a}^* + h\, da]\,\big(m_n\, \delta_{(\theta_n, u_n)}\big) \;=\; m_n \sum_{\theta'} \pi(\theta' \mid \theta_n)\; \dot{\delta}_{(\theta',\, \bar{a}^*(\theta_n, u_n))}\big[\, da(\theta_n, u_n)\,\big],$$

one dipole per destination, at the states where the atom's households land, pointing in the direction of their response: this is the atom part of Proposition 3, and it vanishes at every binding row, where $da = 0$, so only the atoms' unbinding rows, the story's households leaving the bound today, seed dipoles. The transfer operator keeps the family closed: pairing against $\mathcal{K}\psi$ and applying the chain rule of Corollary 1 at a state $(\theta, u)$ off the binding boundary $\Gamma$,

$$\Lambda\, \dot{\delta}_{(\theta,u)}[v] \;=\; \sum_{\theta'} \pi(\theta' \mid \theta)\; \dot{\delta}_{(\theta',\, \bar{a}^*(\theta,u))}\big[\, D_a \bar{a}^*(\theta, u)\, v\,\big],$$

a dipole again, at the next point of the forward orbit, its direction carried forward by the policy Jacobian.

The distribution-side responses therefore live in

$$E \;=\; V' \,\oplus\, D,$$

where $D$ is the register of singular motions: one entry per singular carrier of Assumption 1, finitely many in all, each entry a carrier and a weight. At an atom's orbit point the weight is a mass and a direction, the dipole $w\,\dot\delta_{(\theta,u)}[v]$ just built; on a carrier of positive dimension it is the density of the normal-displacement layer on that carrier, in the same fractional taxonomy $W^{c_f/p,\,p'}$ of Appendix C that clause 1 asks of the steady-state layers. The transfer operator acts on the register carrier by carrier: each carrier goes to its image carrier, a point to the next orbit point, a positive-dimension carrier through the per-regime policy, and the weight is transformed along, by the policy Jacobian at a point and by the change of variables of Lemma 3 on a positive-dimension carrier; each transformed density stays in its fractional class, the per-carrier maps being bi-Lipschitz (Section 5.2, where $D$ is the first grade $D_1$ of Definition 7). The atom entries' dynamics is the explicit dipole recursion just displayed, and their pairings

$$\big\langle w\,\dot\delta_{(\theta,u)}[v],\; g\big\rangle \;=\; w\; v\cdot\nabla_a g(\theta,u)$$

are ordinary numbers for every per-regime smooth $g$, clause 4 of Assumption 1 keeping every dipole at a state where $\bar x^*$ is continuously differentiable. At $d_a = 1$, the Krusell–Smith case, the atoms are the only singular parts (clause 1) and the register reduces to the span

$$D \;=\; \operatorname{span}\Big\{\,\dot\delta_{(\theta,u)}[v] \;:\; (\theta,u)\in\textstyle\bigcup_n \mathcal{O}(\theta_n,u_n),\; v\in\mathbb{R}\,\Big\},$$

the dipoles at the states of the forward orbits $\mathcal{O}(\theta_n,u_n)$ of the atoms of Assumption 1, finite-dimensional at each horizon and invariant under $\Lambda$. The direct sum is two registers, not a repair of the dual. Nothing is adjoined to $V'$: the atoms themselves were always inside it, every density response stays inside it, and $V'$ is exactly the space it was. $D$ is the second register, a finite list of carriers and weights, updated by the explicit transport above and kept outside the functional analysis by design; pairing a weighted dipole with an aggregating function is the literal number $w\,v\cdot\nabla_a\bar x^*(\theta,u)$. The Gelfand triple of Appendix C is untouched: it keeps sorting what is a function from what is a distribution, its pivot serves exactly where it always did, and $D$ simply never enters it. Economically $D$ is the outflow channel: it tracks, to first order, the drift of exactly the mass the constraint had parked. No generalized function is ever created from the policy, and none is differentiated; the singular parts were data in the steady state, and their motion is carried as bookkeeping on finitely many carriers, the same information Bhandari et al. [2023, Assumption 2] track inside their generalized-function calculus.

**Identification with Bhandari et al.'s displacement operator.** Bhandari et al. [2023] write the displacement operator of Proposition 3 as a Dirac kernel against the steady-state distribution,

$$\bar\Lambda(\theta',a',\theta,a) \;=\; \prod_i \delta\Big((\bar a^*)^i(\theta,a)-a'^i\Big)\;\pi(\theta'\mid\theta)$$

$$(\mathcal{M}\cdot da)^j\langle\theta',a'\rangle \;=\; \iint \bar\Lambda(\theta',a',\theta,a)\,da^j(\theta,a)\,d\bar\Omega^*(\theta,a),$$

a distribution in the destination $(\theta',a')$ that acquires meaning only paired against a test function. Pairing against $\psi\in V$ collapses the $\delta$ through the destination integral, and the policy is then only ever evaluated, never differentiated,

$$\big\langle(\mathcal{M}\cdot da)^j,\;\psi\big\rangle \;=\; \iint \psi(\theta',\bar a^*(\theta,a))\,da^j(\theta,a)\,d\pi(\theta'\mid\theta)\,d\bar\Omega^*(\theta,a),$$

which is exactly the pairing that defines $\mathsf{M}$ in Proposition 3. So the displacement operator $\mathsf{M}$ is Bhandari et al.'s $\mathcal{M}$, read so that the Dirac kernel is integrated against the test function up front and never has to be differentiated into a $\delta'$.

**The velocity operator.** A displacement, once deposited, does not stay put: in later periods the steady-state dynamics carry it along. What is carried is a field, one signed measure per predetermined coordinate, and Bhandari et al. [2023, §3.3] transport it with a second operator, written, like the displacement operator, as the Dirac kernel $\bar\Lambda$ integrated against a slope-weighted field,

$$\big(\mathcal{L}^{(a)}\cdot\omega\big)^i\langle\theta',a'\rangle \;=\; \iint \bar\Lambda(\theta',a',\theta,a)\sum_j \partial_{a^j}\bar a^{*\,i}(\theta,a)\;\omega^j(\theta,a)\,d\theta\,da.$$

Pairing against a test function $\psi$ collapses the kernel through the destination integral, exactly as it did for $\mathsf{M}$,

$$\big\langle \big(\mathcal{L}^{(a)} \cdot \omega\big)^i,\ \psi \big\rangle \ = \ \iint \psi\big(\theta', \bar{a}^*(\theta,a)\big) \sum_j \partial_{a^j} \bar{a}^{*\,i}(\theta,a)\, \omega^j(\theta,a)\, d\pi(\theta' \mid \theta)\, d\theta\, da,$$

so that, read on fields,

$$\big(\mathcal{L}^{(a)}\, \omega\big)^i \ = \ \Lambda\Big( \sum_j \partial_{a^j} \bar{a}^{*\,i}\, \omega^j \Big), \qquad i = 1, \dots, d_a :$$

weight the field by the policy slopes, the share of a displacement along coordinate $j$ today that survives into coordinate $i$ of tomorrow's position, then push forward. The weight is what separates Bhandari et al.'s two operators: $\mathcal{M}$ weights by the mass and moves it, $\mathcal{L}^{(a)}$ weights by the slopes $\partial_{a^j} \bar{a}^{*\,i}$ and lets an existing displacement drift.

The slope weight has a regularity payoff that the bare push-forward lacks.

**Lemma 3** (Transport regularity)**.** *Under Assumption 1, let the field $\omega$ have a density in $L^{p'}$, respectively in $L^\infty$. Then $\mathcal{L}^{(a)}\omega$ consists of an interior part with a density in the same class and, when $d_a \geq 2$, finitely many layers seeded by the regimes that pin one coordinate, each carried by its coordinate's set $\{a^i = \underline{a}^i\}$; no iterate $(\mathcal{L}^{(a)})^t\omega$ ever develops a point atom. With a single predetermined coordinate no layers arise, the transported density is the composition of the old density with the per-regime inverses of the policy, and the $L^\infty$ bound does not grow. The proof is in Appendix C.*

The mechanism deserves one sentence, because it is the honest content of the word velocity. Pushing a density through the policy $\bar{a}^*$ distorts volumes by the Jacobian of $\bar{a}^*$, and where households compress toward the constraint that distortion blows the density up; but with one predetermined coordinate the slope weight $\partial_a \bar{a}^*$ *is* that Jacobian, so in the change of variables the two cancel exactly, and the transported density is the old density read at the pre-image $g(\theta, a') \equiv \bar{a}^*(\theta, \cdot)^{-1}(a')$, the inverse of the policy on the regime. Where the policy is flat, the weight and the Jacobian vanish together, so the mass a bare push-forward would pile into an atom is annihilated before it moves: $\mathcal{L}^{(a)}$ transports displacement without distorting it, which is what a velocity field does. With several coordinates the weight is the matrix $D_a \bar{a}^*$ while the volume distortion is its determinant, so the exact cancellation relaxes to a bound through clause 2 of Assumption 1; and a regime that pins one coordinate, sending mass exactly onto the set $\{a^i = \underline{a}^i\}$ while the free coordinates keep moving, deposits a layer there, the kind of object clause 1 already admits in the steady state itself.

The operator enters through one identity, the algebraic engine of Bhandari et al.'s state-space bookkeeping.

**Lemma 4** (The transfer operator commutes with the divergence)**.** *Under Assumption 1, for every field $\omega$ with an $L^{p'}$ density,*

$$\Lambda^t \left(\nabla_a \cdot \omega\right) \ = \ \nabla_{a'} \cdot \Big( \big(\mathcal{L}^{(a)}\big)^t\, \omega \Big) \qquad \text{for every } t \geq 0,$$

*where the weak divergence is defined against the test space by $\langle \nabla_a \cdot \omega,\ \psi \rangle = -\sum_j \langle \omega^j,\ \partial_{a^j} \psi \rangle$, the boundary term absent by the no-flux property of Assumption 1.*

*Proof.* One step suffices: Lemma 3 keeps every iterate $(\mathcal{L}^{(a)})^k\omega$ inside the class the pairings need. Fix a smooth $\psi$; smooth functions are dense in $V$, and each part of both sides is a bounded functional, the interior densities on $V$ and the layers through their traces (Appendix C), so proving the identity on smooth $\psi$ proves it on the class every later use draws from. Then

$$\begin{aligned}
\big\langle \Lambda\big(\nabla_a \cdot \omega\big),\ \psi \big\rangle \ &= \ \big\langle \nabla_a \cdot \omega,\ \mathcal{K}\psi \big\rangle \ = \ -\sum_j \big\langle \omega^j,\ \partial_{a^j}\big(\mathcal{K}\psi\big) \big\rangle \ = \ -\sum_j \sum_i \big\langle \partial_{a^j} \bar{a}^{*\,i}\, \omega^j,\ \mathcal{K}\big(\partial_{a'^i} \psi\big) \big\rangle \\
&= \ -\sum_i \Big\langle \Lambda\Big( \sum_j \partial_{a^j} \bar{a}^{*\,i}\, \omega^j \Big),\ \partial_{a'^i} \psi \Big\rangle \ = \ -\sum_i \big\langle \big(\mathcal{L}^{(a)}\omega\big)^i,\ \partial_{a'^i} \psi \big\rangle \ = \ \big\langle \nabla_{a'} \cdot \big(\mathcal{L}^{(a)}\omega\big),\ \psi \big\rangle.
\end{aligned}$$

The steps are, in order: the adjoint relation of Lemma 2; the definition of the weak divergence, the pairing an honest integral part by part, the interior density against $\partial_{a^j}(\mathcal{K}\psi) \in L^p$ and each layer against the one-sided trace of the same function; the state-direction chain rule of Corollary 1; the adjoint relation again, applied to the weighted field; the definition of $\mathcal{L}^{(a)}$; and the weak divergence read backwards, its boundary term again dropped by no-flux. Nothing was differentiated along the way except the test function. □

Transporting the divergence of a field is therefore the same as weighting the field by the policy slopes, transporting it, and taking the divergence at the destination. Read economically: moving a cloud of displaced households forward and asking where their mass piles up is the same as letting the displacement itself drift with the dynamics and asking at the end.

*The created layers.* With one predetermined coordinate the two lemmas above are complete: no layers arise, and every iterate is a plain density. With several coordinates the velocity operator itself creates layers, because a regime that pins one coordinate and leaves the others free pushes its interior mass onto the pinned coordinate's set $\{a^i = \underline{a}^i\}$ through the free block of the policy Jacobian (proof of Lemma 3, Appendix C). What the iterates of Lemma 4 need of these created layers is derived, not asserted, and it is one class weaker than what clause 1 asks of the steady-state carriers. The created density is the fiber integral of the interior density through the collapse, the co-area of the change of variables that powers Lemma 3, hence an $L^{p'}$ function on its carrier by Minkowski's inequality, transported thereafter by the bi-Lipschitz per-carrier maps of Section 5.2's register, which preserve $L^{p'}$; and an $L^{p'}$ weight is all a degree-zero layer's pairings consume, the trace of a test function lying in $L^p$ of the carrier and Hölder closing the integral. The fractional class $W^{1/p,\,p'}$ is not derivable from $\rho^* \in L^{p'}$ alone, nothing along the collapsed fibers being there to inherit, and it is not needed: clause 1 asks it of the *steady-state* carriers, whose iterated-trace pairings at codimension two and above do consume it. The law of motion and the Jacobian below consume none of this; only the velocity-field bookkeeping does, and the remark that follows prices it.

*A comparative remark.* The measure-level construction of this section needs none of the machinery of the last two lemmas: the law of motion below and the Jacobian of the next subsection use only the boundedness of the transfer operator on $V'$ and the finite dipole span $D$, whatever the number of predetermined coordinates $d_a$. The velocity field is Bhandari et al.'s bookkeeping, and the transport and commutation lemmas price it: the fields must keep a density in every part, interior and face alike, which is what the non-degeneracy clause and the fractional face classes of Assumption 1 buy. Neither demand is Bhandari et al.'s own: tracking generalized functions throughout, they never need the transported fields to be densities at all, and the two clauses are what keeping every pairing an ordinary integral costs. That the velocity formulation is the more demanding one is a result in itself; the identifications with Bhandari et al. below hold wherever its demands are met, and the objects they compute agree with the measure-level ones throughout.

**The law of motion of the distribution response.** The pieces assemble into the first-order law the cross-section obeys. Along a perturbed aggregate path the distribution moves for two reasons: the inherited cross-section has already moved, and the policy pushing it forward is moving too. The first force propagates by the transfer operator, the second deposits the displacement of Proposition 3.

**Proposition 4** (Law of motion of the distribution response)**.** *Along a perturbed aggregate path with predetermined-state policy responses $\{\hat{a}_t\}_{t\geq 0}$, the first-order response of the cross-section obeys, in $E = V' \oplus D$,*

$$\hat{\Omega}_{t+1} \;=\; \Lambda\,\hat{\Omega}_t \;-\; \nabla_{a'}\cdot\big(\mathsf{M}\cdot\hat{a}_t\big), \qquad \hat{\Omega}_0 \;=\; 0,$$

*so that, unrolled from the steady state,*

$$\hat{\Omega}_t \;=\; -\sum_{\tau=0}^{t-1} \Lambda^{t-1-\tau}\,\nabla_{a'}\cdot\big(\mathsf{M}\cdot\hat{a}_\tau\big).$$

*Proof.* The law of motion of Section 2 reads $\bar{\Omega}_{t+1} = \Lambda[\bar{a}_t]\,\bar{\Omega}_t$, the transfer operator built from the date-$t$ policy. Differentiating along the path at the steady state, where $\Lambda[\bar{a}^*]\,\bar{\Omega}^* = \bar{\Omega}^*$, the derivative acts once on each argument,

$$\hat{\Omega}_{t+1} \;=\; \big(D_{\bar{a}^*}\Lambda\cdot\hat{a}_t\big)\,\bar{\Omega}^* \;+\; \Lambda\,\hat{\Omega}_t,$$

and the first term is the weak divergence of Proposition 3 in the direction $\hat{a}_t$, its interior part in $V'$ and its singular part seeding $D$, the dipoles at $d_a = 1$. The initial condition is predeterminedness: the date-0 cross-section was inherited from a past at the steady state, so $\hat{\Omega}_0 = 0$. The date-1 response is then the pure displacement $\hat{\Omega}_1 = -\nabla_{a'}\cdot(\mathsf{M}\cdot\hat{a}_0)$, and iterating the recursion gives the sum, the transfer operator acting boundedly on the $V'$ summand and by the explicit dipole transport on $D$. □

**Identification with Bhandari et al.'s Lemma 4.** Bhandari et al. [2023] state the same law in a velocity-field bookkeeping. In their several-state form [Bhandari et al., 2023, supplementary material, App. B.1, Lemma 4(FO MV)] they evolve an auxiliary field $\hat{\omega}_t$ by

$$\hat{\omega}_{t+1} \;=\; \mathcal{L}^{(a)}\,\hat{\omega}_t \;-\; \mathsf{M}\cdot\hat{a}_t, \qquad \hat{\omega}_0 \;=\; 0,$$

and recover the response of the distribution function $\hat{\Omega}_t$ by one divergence, $\frac{d}{da}\frac{d}{d\theta}\hat{\Omega}_t = \nabla_a\cdot\hat{\omega}_t$. Their main text [Bhandari et al., 2023, Lemma $4^{HA}$, eq. (34)] is the one-state case: there the field is the single function $\frac{d}{d\theta}\hat{\Omega}_t$, the response of the joint distribution function with its $\theta$-cumulation differentiated away, and the divergence is the plain derivative in $a$. The bridge to Proposition 4 is one application of Lemma 4. If $\hat{\Omega}_t = \nabla_{a'}\cdot\hat{\omega}_t$ at some date, then

$$\hat{\Omega}_{t+1} \;=\; \Lambda\big(\nabla_{a'}\cdot\hat{\omega}_t\big) \;-\; \nabla_{a'}\cdot\big(\mathsf{M}\cdot\hat{a}_t\big) \;=\; \nabla_{a'}\cdot\big(\mathcal{L}^{(a)}\hat{\omega}_t \;-\; \mathsf{M}\cdot\hat{a}_t\big) \;=\; \nabla_{a'}\cdot\hat{\omega}_{t+1},$$

an induction that starts at $\hat{\Omega}_0 = 0 = \nabla_{a'}\cdot\hat{\omega}_0$: at every date the measure response is the divergence of their field, sign and timing included. The identification carries the singular mass along with it. Their steady-state input is the full $\omega^*$ of their Assumption 2, mass points included, so their field $\hat{\omega}_t$ carries transported point-mass terms $v\,\delta_{(\theta,u)}$, and the divergence of such a term is exactly the dipole, $-\nabla_a\cdot\big(v\,\delta_{(\theta,u)}\big) = \dot{\delta}_{(\theta,u)}[v]$: what this section splits into $V'$ and $D$, their generalized-function calculus carries as one object. On those atom terms the commutation identity holds verbatim, the slope-weighted push of a weighted point mass and the transport of its dipole being the two displays after Definition 3 read through one divergence. The economic reading is theirs: the $\mathcal{L}^{(a)}$ term is the mechanical force, households displaced yesterday drifting on through the unchanged steady-state policy, and the source is the behavioral force, today's change of saving behavior $\hat{a}_t$ depositing fresh displacement [Bhandari et al., 2023, §3.3].

Proposition 4 is the continuous, basis-free law. Its $P^1$ specialization, the histogram update that Dynare 7 actually iterates, is constructed in Section 4.5, where the divergence and its sign are absorbed into the derivatives of the hat functions.

## 4.4 Assembling the general-equilibrium Jacobian

The two heterogeneous responses now combine into the object the aggregate block consumes. The heterogeneous block enters the aggregate equations only through cross-sectional integrals, and differentiating one of them along a perturbed path mixes the two responses: the policy moves inside the integral, and the distribution it is integrated against moves too. No announcement experiment is needed; the sequence-space Jacobian falls out of differentiating along an arbitrary path of relevant aggregates.

**Proposition 5** (The sequence-space Jacobian)**.** *Along a perturbed path of relevant aggregates $\{\widehat{Y}_s\}_{s\geq 0}$, the response of a cross-sectional aggregate is the convolution*

$$\left(\widehat{\int \bar{x}\,d\bar{\Omega}}\right)_t \;=\; \sum_{s\geq 0} J_{t,s}\,\widehat{Y}_s, \qquad J_{t,s} \;=\; \big\langle \hat{\Omega}_{t,s},\; \bar{x}^*\big\rangle \;+\; \big\langle \bar{\Omega}^*,\; \times_{s-t}\big\rangle,$$

*with the convention $\times_k = 0$ for $k<0$, and with distribution-response coefficients $\hat{\Omega}_{t,s}\in E$ generated by*

$$\hat{\Omega}_{0,s} \;=\; 0, \qquad \hat{\Omega}_{t+1,s} \;=\; \Lambda\,\hat{\Omega}_{t,s} \;-\; \nabla_{a'}\cdot\big(\mathsf{M}\cdot p\times_{s-t}\big).$$

*The coefficients satisfy $J_{0,s} = \langle\bar{\Omega}^*, \times_s\rangle$ and the diagonal recursion*

$$J_{t,s} \;=\; J_{t-1,s-1} \;+\; \Big\langle \Lambda^{t-1}\Big(-\nabla_{a'}\cdot\big(\mathsf{M}\cdot p\times_s\big)\Big),\; \bar{x}^*\Big\rangle \qquad (t\geq 1).$$

*Proof.* Differentiating the pairing gives one term per moving argument,

$$\left(\widehat{\int \bar{x}\,d\bar{\Omega}}\right)_t \;=\; \big\langle \hat{\Omega}_t,\; \bar{x}^*\big\rangle \;+\; \big\langle \bar{\Omega}^*,\; \hat{x}_t\big\rangle.$$

*The direct term.* The policy response to the whole path is the anti-causal convolution of Section 4.2; re-indexed by the calendar date of the aggregate movement, under the convention $\times_k = 0$ for $k<0$,

$$\hat{x}_t \;=\; \sum_{s\geq 0}\times_s \widehat{Y}_{t+s} \;=\; \sum_{s\geq 0}\times_{s-t}\widehat{Y}_s, \qquad \text{so} \qquad \big\langle\bar{\Omega}^*,\;\hat{x}_t\big\rangle \;=\; \sum_{s\geq 0}\big\langle\bar{\Omega}^*,\;\times_{s-t}\big\rangle\,\widehat{Y}_s.$$

*The distribution term.* The predetermined-state slices of the same convolution drive Proposition 4,

$$\hat{a}_\tau \;=\; p\,\hat{x}_\tau \;=\; \sum_{s\geq 0} p \times_{s-\tau} \widehat{Y}_s,$$

and the law of motion is linear in its driving direction, so the response splits date by date, by induction on $t$. Suppose $\hat{\Omega}_t = \sum_s \hat{\Omega}_{t,s}\,\widehat{Y}_s$, true at $t=0$ with every $\hat{\Omega}_{0,s}=0$ by predeterminedness; then

$$\hat{\Omega}_{t+1} \;=\; \Lambda\,\hat{\Omega}_t \;-\; \nabla_{a'}\cdot\big(\mathsf{M}\cdot\hat{a}_t\big) \;=\; \sum_{s\geq 0}\Big[\,\Lambda\,\hat{\Omega}_{t,s} \;-\; \nabla_{a'}\cdot\big(\mathsf{M}\cdot p\times_{s-t}\big)\Big]\,\widehat{Y}_s,$$

which is the displayed coefficient recursion.

*The convolution.* Collecting the two terms per unit of $\widehat{Y}_s$ gives $J_{t,s} = \langle \hat{\Omega}_{t,s}, \bar{x}^*\rangle + \langle \bar{\Omega}^*, \times_{s-t}\rangle$; at $t=0$ only the direct term survives, $J_{0,s} = \langle \bar{\Omega}^*, \times_s\rangle$.

*The diagonal recursion.* Write $S_k \equiv -\nabla_{a'}\cdot\big(\mathsf{M}\cdot p\times_k\big)$ for the source at horizon $k$ and unroll the coefficient recursion,

$$\hat{\Omega}_{t,s} \;=\; \sum_{\tau=0}^{t-1} \Lambda^{t-1-\tau}\,S_{s-\tau}.$$

Shifting both indices down by one and re-indexing the sum by $\tau \mapsto \tau+1$,

$$\hat{\Omega}_{t-1,s-1} \;=\; \sum_{\tau=0}^{t-2} \Lambda^{t-2-\tau}\,S_{s-1-\tau} \;=\; \sum_{\tau=1}^{t-1} \Lambda^{t-1-\tau}\,S_{s-\tau},$$

the same sum as for $\hat{\Omega}_{t,s}$ with the $\tau=0$ term removed, so

$$\hat{\Omega}_{t,s} \;-\; \hat{\Omega}_{t-1,s-1} \;=\; \Lambda^{t-1}\,S_s,$$

while the direct terms cancel outright, $\times_{s-t} = \times_{(s-1)-(t-1)}$. Pairing against $\bar{x}^*$ gives the recursion. □

The proposition wraps the whole first-order cascade: the direct term $\langle \bar{\Omega}^*, \times_{s-t}\rangle$ is the behavioral channel, households reacting today to aggregates $s-t$ periods ahead while the population has not yet moved, and the term $\langle \hat{\Omega}_{t,s}, \bar{x}^*\rangle$ is the compositional channel, the accumulated displacement of the population read against the unchanged steady-state behavior. The two passages that follow verify that the announcement algorithm of Auclert et al. [2021] and the operator recursion of Bhandari et al. [2023] compute exactly these coefficients.

**Identification with the fake-news algorithm of Auclert et al..** Auclert et al. [2021] reach the coefficients on the grid through a sequence of announcement experiments. Their environment is fully discretized, the cross-section a histogram vector $D_t$ with law of motion $D_{t+1} = \Lambda_t' D_t$ [Auclert et al., 2021, eq. (11)]; for a shock of size $dx$ announced at date 0 to land at date $s$ they write $dY_t^s$, $dD_t^s$, $d\Lambda_t^s$ for the date-$t$ responses [Auclert et al., 2021, §3.2]. The dictionary is immediate. The announced shock is the path $\widehat{Y}_\tau = \mathbf{1}\{\tau = s\}\,d\bar{Y}$, so the date-$t$ policy response is $\times_{s-t}\,d\bar{Y}$, their $dD_t^s$ is $\hat{\Omega}_{t,s}\,d\bar{Y}$, their transition-matrix response $d\Lambda_t^s$ is the weak-* derivative $D_{\bar{a}^*}\Lambda\cdot p\times_{s-t}$, and the predetermined start $dD_0^s = 0$, recorded in the proof of their Lemma 2, is $\hat{\Omega}_{0,s} = 0$. Their Lemma 1 [Auclert et al., 2021, eq. (16)] observes that the responses depend on the distance $s-t$ to the shock and not on the two dates separately; here that time invariance arrives pre-packaged, the coefficients $\times_{s-t}$ carrying a single index by construction, and it is what lets the increments of Proposition 5 slide along diagonals. Their fake-news matrix collects the diagonal differences, $\mathcal{F}_{t,s}\,dx = dY_t^s - dY_{t-1}^{s-1}$, and their Lemma 2 [Auclert et al., 2021, eq. (18)] evaluates it as $y_{ss}'\,(\Lambda_{ss}')^{t-1}\,dD_1^s$: in the notation here, with $\hat{\Omega}_{1,s} = -\nabla_{a'}\cdot(\mathsf{M}\cdot p\times_s)$,

$$\mathcal{F}_{t,s}\,d\bar{Y} \;=\; \big\langle \Lambda^{t-1}\,\hat{\Omega}_{1,s},\; \bar{x}^*\big\rangle\,d\bar{Y} \;=\; \big\langle \hat{\Omega}_{1,s},\; \mathcal{K}^{t-1}\bar{x}^*\big\rangle\,d\bar{Y},$$

the second equality by the adjunction of Lemma 2. That adjunction is the content of their transposes, and it is where their expectation vectors come from: $\mathcal{E}_t \equiv (\Lambda_{ss})^t\,y_{ss}$ [Auclert et al., 2021, Def. 1, eq. (23)], with the recursion $\mathcal{E}_t = \Lambda_{ss}\mathcal{E}_{t-1}$ of their Lemma 3, is the discrete Koopman iterate $\mathcal{K}^t\bar{x}^*$, and their own gloss says exactly this, the expected time path of the outcome for an agent starting at a given grid point. So where Bhandari et al. [2023] iterate the transfer operator on the distribution side of the

pairing, Auclert et al. [2021] iterate the expectation operator on the function side: two factorizations of one bilinear form, exchanged by $\Lambda = \mathcal{K}^*$. Their Proposition 1 [Auclert et al., 2021, eq. (24)] assembles the fake-news matrix, first row $dY_0^s = \langle \bar{\Omega}^*, \times_s \rangle\, dx$ and body $\mathcal{E}'_{t-1}\, dD_1^s$, and their accumulation rule $\mathcal{J}_{t,s} = \mathcal{J}_{t-1,s-1} + \mathcal{F}_{t,s}$ [Auclert et al., 2021, Prop. 1, whose eq. (25) is its summed form] is the diagonal recursion of Proposition 5, coefficient for coefficient. Two presentational differences remain. Auclert et al. derive the Jacobian column by column from the announcement experiments and do not display the convolution $\sum_s J_{t,s}\widehat{Y}_s$, which enters implicitly when the assembled matrix multiplies the input path inside their general-equilibrium system [Auclert et al., 2021, eqs. (27)–(28)]; and their objects live on the grid from the start, truncated at a horizon $T$, where Section 4.5 shows the same truncation arise as the last step of the construction here.

**Identification with Bhandari et al.'s Corollary 1.** On the state-space side the same coefficients are Bhandari et al. [2023, Corollary 1, eq. (35)]. Their remaining operator is defined here directly on fields,

$$\mathcal{I}^{(a)} \cdot \omega \;=\; \sum_j \big\langle \omega^j,\; \partial_{a^j} \bar{x}^* \big\rangle,$$

it weights a field by the slope of the aggregated policy $\bar{x}^*$ and integrates [Bhandari et al., 2023, §3.3]. It is the aggregate face of the weak divergence: one integration by parts, the boundary term dropped by no-flux, gives $\langle \nabla_a \cdot \omega,\; \bar{x}^* \rangle = -\,\mathcal{I}^{(a)} \cdot \omega$ for any field with a density, which is the step Bhandari et al. themselves take to make their recursion operational, $\int \bar{x}\, d\hat{\Omega}_t = -\,\mathcal{I}^{(a)} \cdot \frac{d}{d\theta}\hat{\Omega}_t$ in their notation [Bhandari et al., 2023, §3.3]. The increment of Proposition 5 then factors through their three operators.

**Lemma 5.** *Under Assumption 1, for every $t \geq 1$ and every bounded direction $y$ with $y^i = 0$ on $\{\bar{a}^{*i} = \underline{a}^i\}$ for each bound $i$,*

$$\Big\langle \Lambda^{t-1}\Big( -\nabla_{a'} \cdot \big(\mathsf{M} \cdot y\big)\Big),\; \bar{x}^* \Big\rangle \;=\; \mathcal{I}^{(a)} \cdot \big(\mathcal{L}^{(a)}\big)^{t-1}\, \mathsf{M} \cdot y.$$

*Proof.* Three steps, each an identity already established:

$$\begin{aligned}
\Big\langle \Lambda^{t-1}\Big( -\nabla_{a'} \cdot \big(\mathsf{M} \cdot y\big)\Big),\; \bar{x}^* \Big\rangle \;&=\; \Big\langle -\nabla_{a'} \cdot \Big(\big(\mathcal{L}^{(a)}\big)^{t-1}\, \mathsf{M} \cdot y\Big),\; \bar{x}^* \Big\rangle && \text{(Lemma 4)} \\
&=\; \sum_j \Big\langle \Big[\big(\mathcal{L}^{(a)}\big)^{t-1}\, \mathsf{M} \cdot y\Big]^j,\; \partial_{a^j} \bar{x}^* \Big\rangle && \text{(integration by parts, no-flux)} \\
&=\; \mathcal{I}^{(a)} \cdot \big(\mathcal{L}^{(a)}\big)^{t-1}\, \mathsf{M} \cdot y && \text{(definition of } \mathcal{I}^{(a)}\text{)},
\end{aligned}$$

each pairing an honest integral part by part: the interior density by Lemma 3, the layers through their traces, and the transported point masses against the per-regime continuously differentiable $\partial_{a^j} \bar{x}^*$, their positions kept off the binding boundary $\Gamma$ by clause 4 of Assumption 1. □

With $y = p \times_s$, Proposition 5 becomes, line for line, their Corollary 1: $J_{0,s} = \int \times_s d\bar{\Omega}^*$ and

$$J_{t,s} \;=\; J_{t-1,s-1} \;+\; \mathcal{I}^{(a)} \cdot \big(\mathcal{L}^{(a)}\big)^{t-1}\, \mathsf{M} \cdot p \times_s,$$

under the same convention $\times_k = 0$ for $k < 0$ [Bhandari et al., 2023, eq. (35)]. Even the proofs are parallel: their appendix establishes the corollary by running a two-index recursion on the velocity field, the field-level image of the measure-level recursion in the proof of Proposition 5, and Lemma 4 is the exchange rate between the two. The three identification passages together say one thing: the announcement bookkeeping of Auclert et al. [2021], the operator calculus of Bhandari et al. [2023], and the weak-derivative construction here compute the same array $J$, and they differ only in which side of the adjunction $\Lambda = \mathcal{K}^*$ each prefers to iterate.

**Closing the system.** It remains to feed the heterogeneous Jacobian back into the aggregate block. Along the equilibrium path the aggregate equations hold at every date; differentiating the date-$t$ equation of $G$ and expanding the relevant-aggregate slice into its constituents, $\widehat{Y}_t \subseteq [\mathcal{E}_t,\; P\widehat{X}_{t-1},\; \widehat{X}_t,\; \mathbb{E}_t\widehat{X}_{t+1}]$, the expectation dropping under perfect foresight,

$$D_{\widehat{Y}}G \cdot \widehat{Y}_t \;+\; D_x G \cdot \sum_{s \geq 0} J_{t,s}\, \widehat{Y}_s \;=\; 0 \qquad (t \geq 0).$$

Every slice $\widehat{Y}_s$ contains only entries of the unknown aggregate path $\{\widehat{X}_t\}$ and of the innovation path $\{\mathcal{E}_t\}$. Stacking the equations over the truncation horizon $T$ and sorting every appearance of the slices by what they contain splits the stacked system into two blocks,

$$\tilde{J}_X \cdot \left(\widehat{X}_t\right)_{t<T} \;+\; \tilde{J}_{\mathcal{E}} \cdot \left(\mathcal{E}_t\right)_{t<T} \;=\; 0, \qquad \text{so that} \qquad \left(\widehat{X}_t\right)_{t<T} \;=\; \underbrace{-\,\tilde{J}_X^{-1}\,\tilde{J}_{\mathcal{E}}}_{G}\; \left(\mathcal{E}_t\right)_{t<T},$$

where $\tilde{J}_X$ collects the derivatives of $G$ with respect to the lagged, current, and led aggregates, each composed with the heterogeneous Jacobian $J$ wherever the aggregates reach the households through $\widehat{Y}$, and $\tilde{J}_{\mathcal{E}}$ collects the same composition for the innovations. The matrix $G$ is the general-equilibrium response $\partial\widehat{X}/\partial\mathcal{E}$, which Dynare 7 returns as `oo_.heterogeneity.dr.G`; the two stacked blocks are the state and shock Jacobians of the pipeline in Section 4.5. The aggregate path is the keystone: once $\{\widehat{X}_t\}$ is known, so are the slices $\{\widehat{Y}_t\}$, and the heterogeneous responses follow with no further solving, the policy response by the anti-causal convolution $\hat{x}_t = \sum_{s\geq 0} \times_s \widehat{Y}_{t+s}$ of Section 4.2 and the distribution response by the law of motion of Proposition 4.

Whether the stacked operator remains invertible once the truncation is removed, which is what existence and uniqueness of a bounded first-order response amount to, is taken up in Section 5.1, where the time invariance built into every block above is put to work.

## 4.5 The finite-element projection Dynare 7 computes

Everything so far in Section 4 is exact and infinite-dimensional: the responses are functions on the state space, the distributions live in $E$, and the operators act between them. Dynare 7 computes with finite arrays. This subsection walks through the arrays the code actually forms, one for each continuous object of the previous subsections, and records, where each discrete object appears, the interpolation rate at which it converges to its continuous counterpart. The blueprint is the histogram analysis of Bhandari et al. [2023, App. C], recast so that every step is a named finite-element operation; their claims are stated for a single predetermined state, and the wording here keeps track of what carries to several.

**The discrete carriers.** Two meshes carry the computation, the two of Section 3.1, and they divide the labor by cost. The policy mesh $\boldsymbol{\theta} \times \boldsymbol{a}^p$ carries the policy functions and their responses: every node bears a nonlinear solve at the steady state and one linear solve per horizon in the dynamics, so this mesh is kept coarse. The distribution mesh $\boldsymbol{\theta} \times \boldsymbol{a}^d$ carries only mass, cheap per node, so it is refined independently, typically finer. Both use the same $P^1$ element, read two ways. On the policy mesh the policy arrays $\boldsymbol{x}$ are the nodal values of $P^1$ (multilinear) interpolants, the functions reconstructed off the grid in the time iteration of Section 3.3; on the nodal basis, whose hats take value one at their own node and zero at every other, the nodal values are at once the interpolant's coefficients (Definitions 15–16). On the distribution mesh the reading is the dual one: the distribution is carried by the projection

$$\mathcal{P}^{i,j}(a,\theta) \;=\; \mathcal{P}^i(a)\,\mathcal{Q}^j(\theta),$$

with $\mathcal{P}^i$ the piecewise-linear tent peaking at the asset node $\boldsymbol{a}^d_{[i]}$ and $\mathcal{Q}^j$ the indicator of the Markov node $\boldsymbol{\theta}_{[j]}$, since $\theta$ is discrete and carries no interpolation. With several continuous predetermined coordinates $a \in \mathcal{A}$ ($d_a > 1$) the asset factor is the tensor-product (multilinear) hat, the bilinear $Q^1$ element on a rectangle. This is the $P^1$ nodal basis on the distribution grid of Section 3.1: a partition of unity whose hats are dual to the nodal point-evaluations [Brenner and Scott, 2008, (3.1.3),(3.5.1)–(3.5.2)], and a point mass landing on a node, even a corner node where several bounds bind, is represented exactly. The histogram $\boldsymbol{\Omega}$ is the coefficient vector of the cross-section in this representation, and the mass-splitting lottery of Young [2010], which sends a point mass at $\bar{a}^*(\theta,a)$ to the bracketing nodes with the hat weights, is the projection dual to nodal interpolation against that basis (Lemma 11): not a convenient approximation but the transfer operator read on the nodal basis.

**The discrete operators.** Every discrete operator below is a continuous operator of the previous subsections read through the finite-element machinery of Appendix C: on the policy mesh through the interpolant, on the distribution mesh through pairings with the basis. How the toolkit evaluates that interpolant and its gradient in practice, by a single Gray-code corner walk shared across the operators below, is Appendix G. I take the policy side first, in the order of the computation.

The backward recursion of Lemma 1 needs two operators, and each discretizes by a named operation. The operator $J = F_x + F_{x^+}\,\mathbb{E}[D_a\bar{x}^* \mid \theta, a]\,p$ of Proposition 1 acts pointwise in the state: its coefficients

are derivatives of the pointwise residual $F$, evaluated state by state, so on the policy mesh collocation (Appendix C) turns it into one small matrix per node, inverted node by node. The slope inside it is read off the interpolant: $D_a \bar{x}^*$ at the landing point $\bar{a}^*(\theta, a)$ is the slope of a piecewise-linear function, the difference quotient of the two nodal values bracketing the landing point. The conditional expectation $\mathcal{K}$ itself compresses to the finite-element space as $\mathcal{I}_h \mathcal{K}$, the expectation followed by nodal interpolation: for a $P^1$ function the value at the landing point is computed exactly, so applying the compression to the coefficients of $\times_{s-1}$ returns the nodal values of $\mathbb{E}[\times_{s-1} \mid \theta, a]$ with no error beyond the one already made in carrying $\times_{s-1}$ by its interpolant; and since the landing points are the steady-state ones at every horizon, the compression is a single sparse matrix, tabulated once. It is the mirror image of the lottery: Lemma 11 compresses the transfer operator through the dual interpolant, $\Lambda^d = \mathcal{I}_h^* \Lambda$, the policy mesh compresses the Koopman operator through the interpolant itself, and the adjunction of Lemma 2 survives the discretization.

The approximation of the backward step is therefore located in one place: each $\times_s$ is replaced by its $P^1$ interpolant on the policy mesh. The response is continuously differentiable off finitely many surfaces, the regime boundaries of Proposition 2 and their pre-images under the policy, so the interpolant errs at first order in the policy mesh width away from the cells those surfaces cross, the estimate of Theorem 3 applied regime by regime; on a crossed cell the response jumps and the interpolant smears the jump over the cell. That the estimate holds regime by regime and not globally is worth dwelling on, because it is where the infinite-dimensional setting first refuses a finite-dimensional reflex. Across a regime boundary $\times_s$ jumps by an $O(1)$ amount, which at once expels it from $H^1$ and leaves the interpolant $O(1)$ wrong on the band of crossed cells, a set of measure $O(h)$; Theorem 3 therefore has no global purchase, and the honest global rate lives in the integral norms alone,

$$\left\| \times_s - \mathcal{I}_h \times_s \right\|_{L^p} \;=\; O\big(h^{1/p}\big) \quad (p < \infty), \qquad \left\| \times_s - \mathcal{I}_h \times_s \right\|_{L^\infty} \;=\; O(1),$$

the uniform norm registering no convergence at all. The computation never reads the response pointwise: it forms the aggregate increment $J_{t,s} = \int \times_s \, d\bar{\Omega}^*$, an integration against the stationary distribution, and that pairing converges at the full first-order rate $O(h)$, the rate the distribution side assembles below and Bhandari et al. [2023, App. C] establish for the histogram, the $O(h)$ living in the proofs of their Claims 12–13. The gap between the two rates is the inequivalence of norms of Appendix C in miniature: in finite dimension every norm returns the same verdict on convergence and one need not name one, whereas here the same sequence diverges in the uniform norm and converges in the integral norm the model pairs against, so naming the norm is not a formality but decides whether the object converges at all. Once solved, the responses cross to the distribution mesh by prolongation (Appendix C): the policy-mesh interpolant is evaluated at the distribution nodes, exactly, so the crossing adds no error of its own and the finer mesh inherits the coarse-mesh interpolation error unchanged. The steady-state landing points that drive the lottery below cross the same way: the lottery brackets the interpolated policy, not the exact one, and carries the same inherited error.

On the distribution mesh the histogram coefficients are themselves pairings with the basis: the mass stored at the node $(\boldsymbol{\theta}_{[j]}, \boldsymbol{a}^d_{[i]})$ is $\langle \mu, \; \mathcal{P}^i \mathcal{Q}^j \rangle$, the measure read against the hat, so a distribution enters the computer exactly as it enters the mathematics, through test functions. The discrete transfer operator acts on a node mass by the lottery. A unit mass at the node $(\boldsymbol{\theta}_{[j]}, \boldsymbol{a}^d_{[i]})$ lands at $\bar{a}^*(\boldsymbol{\theta}_{[j]}, \boldsymbol{a}^d_{[i]})$ and is split back onto the bracketing nodes with the hat weights,

$$\Lambda^d \, \delta_{(j,i)} \;=\; \sum_{j'} \boldsymbol{\pi}_{jj'} \sum_{i'} \mathcal{P}^{i'}\big(\bar{a}^*(\boldsymbol{\theta}_{[j]}, \boldsymbol{a}^d_{[i]})\big) \, \delta_{(j',i')},$$

which is Lemma 11 node by node: $\Lambda^d = \mathcal{I}_h^* \Lambda$, the push-forward followed by the interpolation-dual projection. The discrete Koopman operator is the same data read the other way: contracting an array of node values against the weights interpolates it at each landing point and averages over $\boldsymbol{\pi}$, the distribution-mesh instance of the compression $\mathcal{I}_h \mathcal{K}$ met above, and on the grid the two are one sparse matrix and its transpose. The discrete displacement is the same *stencil*, the fixed finite pattern of nodes and weights the operator reads at each node, one derivative up. By Proposition 3 the displacement pairs with a test function through its gradient, and testing against the hats replaces that gradient by hat derivatives, so the array the code assembles reads, node by node,

$$\big(\mathsf{M}^h \, \hat{a}\big)_{(j',i')} \;=\; \big\langle \mathsf{M} \cdot \hat{a}, \; \nabla\big(\mathcal{P}^{i'} \mathcal{Q}^{j'}\big) \big\rangle \;=\; \sum_{j,\,i} \boldsymbol{\pi}_{jj'} \; \boldsymbol{\Omega}_{(j,i)} \; \hat{a}(\boldsymbol{\theta}_{[j]}, \boldsymbol{a}^d_{[i]}) \cdot \nabla \mathcal{P}^{i'}\big(\bar{a}^*(\boldsymbol{\theta}_{[j]}, \boldsymbol{a}^d_{[i]})\big),$$

a sum over source nodes of mass times response times the hat derivative $\pm 1/h$ at the landing point, a bounded jump and never an impulse. Replacing the gradient of a test function by the derivative of

its interpolant is the one approximation of this side, and its rate is the interpolation estimate: first order in the mesh width, $O(h)$, in the $H^1$ norm (Theorem 3), the rate surviving for merely Lipschitz test functions through the averaging Scott–Zhang interpolant [Brenner and Scott, 2008, §4.8]; Bhandari et al. [2023, App. C, Claims 12–13] state the single-coordinate case and its $t$-step iterates, and Lemma 6 below carries the rate to several predetermined coordinates, uniformly in the horizon. Each assembled aggregate increment therefore converges at $O(h)$ to the weak-derivative increment of Proposition 5, which is in turn Bhandari et al.'s velocity composition (Lemma 5): the finite-element reading and the velocity-field reading meet in the limit. The hat derivative absorbs both the divergence and its sign into the basis, which is why the discrete law of motion,

$$\hat{\Omega}_{t+1} \;=\; \Lambda^d\, \hat{\Omega}_t \;+\; \mathsf{M}^h\, \hat{a}_t,$$

carries a plus sign where Proposition 4 carries a minus. This update is the $P^1$ specialization of the continuous law, and Bhandari et al.'s own code iterates precisely it, with a comment flagging that the iterated object is the change in the histogram, not in the distribution function.

The singular mass needs no special treatment on the grid. An atom sitting on a node is represented exactly by the partition of unity; the lottery displaces it with the same weights as any mass; and the hat-derivative stencil applied to a displaced node mass is the finite-dimensional shadow of the dipole $\dot{\delta}_{(\theta,u)}[v]$, the difference quotient of the hats standing in for the gradient the dipole would evaluate. The code computes the dipole terms of $D$ without ever naming them.

**The pipeline, array by array.** The pipeline is run by the command `heterogeneity_solve`, which requires a steady state in memory, loaded or computed by the commands of Section 3, and accepts a single option, `truncation_horizon`, the horizon $T$ of every array below, with default 300. The command populates `oo_.heterogeneity.dr`; the documented, user-facing subfield is `dr.G`, the remaining subfields being internal intermediates whose names and layout may change. Internally, the first-order solver `+heterogeneity/solve.m` forms, in order:

1. *The policy-response coefficients* $\times_s$. The per-node collocation blocks of $J$ are assembled as `f3` in `compute_impulse_responses`, the interpolant slopes at the landing points entering through the inverse mesh widths `mat.pol.inv_h` (`compute_expected_dx`); the compressed Koopman matrix is tabulated once at load time as `mat.pol.Phi_e`, from the landing-point brackets and weights `mat.pol.ind`, `mat.pol.w` and the shock kernel `mat.Mu` (`load_steady_state`). The backward recursion of Lemma 1 is then the horizon loop in `compute_impulse_responses`, one expectation product and one linear solve per node and horizon, nodal values converted to interpolant coefficients through the LU factors of the collocation matrix (`load_steady_state`; the identity for the nodal basis). The result is the array `impulse_responses`, the discrete $\{\times_s\}_{s<T}$.

2. *The displacements* `curlyDs`. The response coefficients are first prolonged onto the distribution mesh, one sparse product with the bridge matrix `mat.d.Phi` (assembled in `solve`; the matrix of policy-mesh hats at distribution nodes, built in `load_steady_state`). The stencil then contracts, node by node, the histogram `mat.d.hist`, the prolonged responses, the lottery indices and weights `mat.d.ind`, `mat.d.w`, tabulated from the prolonged steady-state policy (`load_steady_state`), the inverse mesh widths `mat.d.inv_h`, and the shock kernel `mat.Mu` (`solve`): the $P^1$ displacement $\mathsf{M}^h\, p \times_s$ for each horizon $s$, the discrete source terms of the law of motion.

3. *The impact row* `curlyYs` (in `solve`): the contraction of the prolonged $\times_s$ against the steady-state histogram, the discrete $J_{0,s} = \langle \bar{\Omega}^*, \times_s \rangle$ of Proposition 5, and the name Auclert et al. [2021] use for the same array.

4. *The expectation arrays* `expectations` (`compute_expected_pol`): the discrete Koopman iterates $\mathcal{K}^t \bar{x}^*$ for $t < T-1$, the expectation vectors $\mathcal{E}_t$ of Auclert et al. [2021, Def. 1, eq. (23)]. These arrays exist because of the adjunction of Section 4.4: filling the Jacobian with the contractions $\langle \hat{\Omega}_{1,s},\ \mathcal{K}^{t-1} \bar{x}^* \rangle$ costs cheap inner products where iterating the distribution column by column would cost a push-forward per entry, the computational content of Auclert et al.'s algorithm. The code centers each array to zero cross-sectional mean (in `compute_expected_pol`), which is harmless because every displacement it is later contracted with carries zero total mass, and numerically kinder.

5. *The matrices* $\mathcal{F}_{t,s}$, `transition_matrices` (computed in the Fortran kernel `compute_transition_matrices.f08`): first row the impact row `curlyYs`, body the contractions $\mathcal{F}_{t,s} = \langle \hat{\Omega}_{1,s},\ \mathcal{K}^{t-1} \bar{x}^* \rangle$ of `expectations` with `curlyDs`, the diagonal increments of Proposition 5.

6. *The heterogeneous Jacobian* `heterogeneous_jacobian` (`compute_heterogeneous_jacobian`): the accumulation $J_{t,s} = J_{t-1,s-1} + \mathcal{F}_{t,s}$, coded as one running sum along diagonals in the same routine, the fill rule of Auclert et al. [2021, Prop. 1, eq. (25)].

7. *The general-equilibrium response* `G` (in `solve`): the closure of Section 4.4. The code assembles the two stacked blocks, the state Jacobian $\tilde{J}_X$ and the shock Jacobian $\tilde{J}_\mathcal{E}$, from the aggregate-block derivatives composed with the heterogeneous Jacobian (`compute_state_jacobian` and `compute_shock_jacobian`), solves the stacked system, assembled sparse and densified for the solve, and returns $G = \partial\widehat{X}/\partial\mathcal{E}$ as `oo_.heterogeneity.dr.G`.

Truncation is the one departure from the continuous objects: every array stops at the horizon $T$, so the code's $J$ is the $T \times T$ upper-left corner of the operator of Proposition 5, the same corner Auclert et al. [2021] assemble. The returned `dr.G` is likewise the $T \times T$ corner of the untruncated response operator whose infinite inverse Proposition 9 constructs; that the corner converges to that inverse is not automatic but the finite-section fact of Section 5.1, holding under determinacy together with the reflected-symbol condition recorded there, not from invertibility of the infinite operator alone.

**What the $P^1$ basis cannot see.** Two consequences close the comparison. First, on the smooth part of the policy the discrete `curlyDs` of Dynare 7 and the finite-difference displacement of Auclert et al. [2021] converge to the same weak derivative, at the $O(h)$ rate recorded above, so the state-space construction of Bhandari et al. [2023] and the sequence-space Jacobian of Auclert et al. [2021] agree at first order, exactly as Bhandari et al. [2023] report numerically; the $P^1$ partition of unity is what keeps the discrete distribution a unit-mass object through the limit. This is the precise sense in which Dynare 7's mixed approach, the heterogeneous block of Bhandari et al. [2023] with the aggregate assembly of Auclert et al. [2021], is consistent. Second, the fidelity is a first-order fact. The $P^1$ histogram misses the second-order $C_{1,1}$ term that Bhandari et al. [2023, App. C.1] isolate, the contribution of the moving binding boundary $\Gamma$, because the $P^1$ hat has zero second derivative, $\mathcal{P}_{aa} = 0$, so the basis carries nothing for that term to land on; Bhandari et al. note that the histogram generically misses such second-order terms however fine the grid. At first order, the order Dynare 7 computes, the projection loses nothing but the $O(h)$ interpolation error recorded above; what the second order asks of the basis, and what the $P^1$ element cannot give it, is taken up in Section 5.2.
The $O(h)$ pairing rate above holds at any number of predetermined coordinates, not only the single one Bhandari et al. [2023] treat, and the reason is a stationarity count worth isolating, since the same mechanism reappears for the higher-order register in Appendix E.

**Lemma 6** (The pairing rate at several predetermined coordinates)**.** *Let Assumptions 1 and 2 hold. Then each aggregate increment $\langle \Lambda^{t-1} S_s, \bar{x}^* \rangle$ of Proposition 5, assembled on the policy mesh of width $h$ as in the pipeline above, converges to its continuous value with a horizon-summable rate: for a constant $C$ independent of $h, t, s$,*

$$\left| \langle \Lambda^{t-1} S_s, \bar{x}^* \rangle_h - \langle \Lambda^{t-1} S_s, \bar{x}^* \rangle \right| \;\le\; C\, \rho^s\, t\, \lambda^{t-1}\, h,$$

*at any number $d_a$ of predetermined coordinates; the aggregate pairing therefore converges at $O(h)$ uniformly in the horizon, and Bhandari et al. [2023, App. C, Claims 12–13]'s single-coordinate rate is its $d_a = 1$ specialization.*

*Proof.* By the adjunction of Lemma 2 and the centering of Step 2 of Proposition 6, the increment is $\langle S_s, \psi_t \rangle$ with $\psi_t = \mathcal{K}^{t-1}\bar{x}^* - \langle \bar{\Omega}^*, \bar{x}^* \rangle$ the centered iterate, and the pipeline replaces $\psi_t$ by its $P^1$ interpolant $\mathcal{I}_h\psi_t$, pairing the gradient error against the bounded direction $p \times_s$, $\|\times_s\|_\infty \le C\rho^s$, over the probability measure $\tilde{\mu}$ the source induces (Step 1 of Proposition 6, the kernel $\pi$ times $\bar{\Omega}^*$ read at the landing points):

$$\left| \langle S_s, \psi_t - \mathcal{I}_h\psi_t \rangle \right| \;\le\; C\, \rho^s \int \left| \nabla(\psi_t - \mathcal{I}_h\psi_t) \right| d\tilde{\mu}.$$

Split the integral at the seams, the iterated pre-images of $\Gamma$ under the per-regime policy across which $\psi_t$ is only piecewise smooth, $C^k$ and finitely many per horizon by the transversality clause of Section 5.2. Off the seams $\psi_t$ is $C^2$ per regime with $\|\nabla^2\psi_t\|_\infty \le C\lambda^{t-1}$ (clause 2 of Assumption 2, one grade up as in Appendix E), so the averaging Scott–Zhang interpolant returns the gradient error $O(h\,\|\nabla^2\psi_t\|_\infty)$ [Brenner and Scott, 2008, §4.8], at most $Ch\lambda^{t-1}$ against the unit measure. On the $h$-band around each

seam the gradient of $\psi_t$ jumps by at most $2\|\nabla\psi_t\|_\infty \le C\lambda^{t-1}$, so the gradient error there is $O(\lambda^{t-1})$ pointwise. The band around the depth-$j$ seam is the $j$-step pre-image of the band $B$ around $\Gamma$, and $\bar{\Omega}^*$ is $\Lambda$-invariant, so $\tilde{\mu}$ assigns it the mass $\tilde{\mu}(B) \le Ch$ that $B$ itself carries ($\Gamma$ codimension one, the density bounded); the band masses sum to an expected-visit count,

$$\sum_{j=0}^{t-1} \tilde{\mu}\big(\text{band}_j\big) \;=\; \sum_{j=0}^{t-1} \tilde{\mu}(B) \;=\; t\,\tilde{\mu}(B) \;\le\; C\,t\,h,$$

linear in the horizon, the geometric growth of the seam surface area not entering because each pre-image band carries the same invariant mass as $B$. The band contribution is thus $C\lambda^{t-1}\,t\,h$, and with the smooth part the increment error is $C\rho^s(h\lambda^{t-1} + t\lambda^{t-1}h) \le C\rho^s\,t\,\lambda^{t-1}h$, whose constant is summable in $s$ and $t$ since $\rho, \lambda < 1$. At $d_a = 1$ the band around the moving crossing point carries mass $O(h)$ and the sum is the same $O(th)$, recovering Claims 12–13. □

**Simulating.** The solved decision rule is consumed by the command `heterogeneity_simulate`, which requires `heterogeneity_solve` to have run and auto-detects one of two modes. In the stochastic mode, innovations are unanticipated: the command always returns impulse responses to each declared aggregate shock, over the horizon set by the option `irf` (default 40), restricted to the shocks listed in `irf_shocks`, and rescaled by `relative_irf` on request; when the option `periods` is positive it also simulates a stochastic path, drawing innovations each period from a normal law with the variance matrix `M_.Sigma_e`. In the news mode, triggered as soon as the `shocks` block carries `periods` and `values` clauses, households learn the whole announced sequence at date 0, the anticipation convention of Auclert et al. [2021], and the stochastic-mode options are unavailable. Outputs land in the standard Dynare slots, impulse responses in `oo_.irfs` and simulated paths in `oo_.endo_simul` and `oo_.exo_simul`; internally, both modes contract the relevant shock paths against `dr.G`, the convolution of Proposition 5 at work. One difference with the sequence-space practice deserves a flag: the toolkit of Auclert et al. [2021] consumes a user-supplied sequence of aggregate-shock values, while Dynare simulates innovations to an exogenous process whose equation must therefore be written in the `model` block; the response to any anticipated path is then read off `dr.G` by one matrix-vector product.

**Backing out the heterogeneous responses.** The command `heterogeneity_simulate` reports aggregate paths; the cross-sectional responses, of the policy functions and of the distribution, it leaves to be reconstructed, and each follows in a single pass once the aggregate path $\{\widehat{Y}_t\}$ is in hand, read off the first column of the relevant `dr.G` subfields. The policy response is the anti-causal convolution of Section 4.2,

$$\hat{x}_t(\theta, a) \;=\; \sum_{s \ge 0} \times_s(\theta, a)\,\widehat{Y}_{t+s},$$

the coefficients $\times_s$ being the array `impulse_responses` of step 1 of the pipeline, which the solver forms but does not persist; the helper recomputes them from the stored steady-state derivatives by the backward recursion of Lemma 1 and contracts them against the path, the timing of each relevant-aggregate slice $\widehat{Y}_t$, its lagged, current, and led aggregates, carried by the alignment of the convolution window. The distribution response is the forward law of motion of the discrete cross-section,

$$\hat{\Omega}_0 \;=\; 0, \qquad \hat{\Omega}_{t+1} \;=\; \Lambda^d\,\hat{\Omega}_t \;+\; \sum_{s \ge 0} \texttt{curlyDs}_s\,\widehat{Y}_{t+s},$$

the specialization to a known path of the update of Section 4.5: the per-period source $\mathsf{M}^h\hat{a}_t$ is, by linearity of $\hat{a}_t = p\,\hat{x}_t = \sum_{s\ge0} p \times_s \widehat{Y}_{t+s}$ and the identity $\texttt{curlyDs}_s = \mathsf{M}^h p \times_s$ of step 2, the stored displacement arrays `dr.curlyDs` convolved against the same path, while the steady-state push-forward $\Lambda^d$ is rebuilt from the lottery brackets and weights `mat.d.ind`, `mat.d.w` and the shock kernel `mat.Mu`, the discrete transfer operator of Section 4.5. Both constructions read subfields beneath `dr.G`, the displacements `dr.curlyDs` and the steady-state lottery data, which the warning above marks as internal intermediates whose names and layout may change; I supply them as the helpers `het_policy_responses.m`, `het_distribution_responses.m`, and `build_steady_state_pushforward.m`, listed in Appendix H.

# 5 Existence, determinacy, and the higher orders

The framework was built to compute the first-order response; this section records what it opens beyond that computation: the question of existence and uniqueness that the truncation hid, and the higher orders the same operators deliver. A one-equation illustration, the purely anticipatory economy $X_t = \beta\,\mathbb{E}_t[\varphi(X_{t+1})] + Z_t$, runs alongside every proof below, chosen so that each abstract object, the symbol and its winding number, the factorization, and the higher-order kernels, can be written out and checked by hand.

## 5.1 Existence and determinacy

Section 4.4 solved a truncated system, and the truncation hid a question. At the household level, Proposition 1 secured existence and uniqueness of the backward map through the isomorphism $J$; at the general-equilibrium level nothing so far guarantees that the untruncated stacked system determines a bounded aggregate response, or determines only one. No finite computation can settle this on its own: a $T \times T$ corner may be invertible while the infinite system is not, and conversely. The question is determinacy, and it is live in the literature this document builds on: neither Bhandari et al. [2023] nor Auclert et al. [2021] state a criterion, and the criterion now available, due to Auclert et al. [2025], is proved for the discretized heterogeneous-agent block. This subsection establishes that criterion for the continuous objects of Section 4, on which Auclert et al. [2025] do not work, and later certifies (Proposition 10) that the winding number a discretized solve reads off is the continuous model's; the framework is what makes both short: the backward operator $\mathfrak{r}$ and the transfer operator $\Lambda$ carry no date index, so the stacked Jacobian is a linear-time-invariant object, and the invertibility of such operators is classical mathematics, settled between the 1950s and the 1970s and collected in Böttcher and Silbermann [1999]; the economics application is the recent step, Onatski [2006] for linear rational-expectations models and Auclert et al. [2025] for the sequence space.

The right home for the unknown path is $\ell^2$, the space of square-summable sequences: the aggregate response $(\widehat{X}_t)_{t\geq 0}$ is required to obey $\sum_t \|\widehat{X}_t\|^2 < \infty$. Square-summability of the impulse response is finite variance for the stochastic economy the perturbation approximates, and the requirement has bite: it excludes the random-walk component of a linearized representative-agent economy [Auclert et al., 2025, §2.1, fn. 5], while the stationary heterogeneous-agent models of this document sit inside comfortably. On $\ell^2$ the operators of Section 4.4 fall into a classical class, defined next.

**Definition 4** (Block Toeplitz and quasi-Toeplitz operators, symbol, winding number)**.** Fix a block size $n$ and let $\ell^2$ carry the square-summable sequences $x = (x_t)_{t\geq 0}$ of vectors $x_t \in \mathbb{R}^n$, one block per date; the matrix of a bounded operator on $\ell^2$ is then a doubly indexed array of $n \times n$ blocks. The operator is *block Toeplitz* when its blocks are constant along each diagonal,

$$\big[\mathsf{T}(a)\big]_{t,s} \;=\; a_{t-s} \qquad (t, s \geq 0),$$

for a two-sided family of matrices $(a_k)_{k\in\mathbb{Z}}$, $a_k \in \mathbb{R}^{n\times n}$, with $\sum_k \|a_k\| < \infty$; the matrix-valued function on the unit circle

$$a(z) \;=\; \sum_{k\in\mathbb{Z}} a_k\, z^k, \qquad |z| = 1,$$

is its *symbol*, and the scalar case $n = 1$ is a Toeplitz operator proper. The operator is *block quasi-Toeplitz* when it is the sum of a block Toeplitz operator and a compact correction, $\mathsf{T}(a) + \mathsf{C}$; the classical theory of both classes is collected in Böttcher and Silbermann [1999, Chs. 1 and 6]. For a continuous scalar function $f$ that does not vanish on the unit circle, the *winding number* $\mathrm{wind}(f)$ is the signed number of turns the closed curve $\omega \mapsto f(e^{i\omega})$ makes around the origin.

The time convention matches Section 4.4: the entry $a_{t-s}$ carries the date-$t$ response to a date-$s$ movement, so the coefficient on $z$ acts with one period of delay and the coefficient on $z^{-1}$ with one period of anticipation; the lag operator has symbol $z$, the lead operator $z^{-1}$, and a forward-looking block loads the negative powers.

Two geometric-decay hypotheses power the whole subsection, and one word of operator theory phrases the first. The *spectral radius* of a bounded operator $A$ is the geometric growth rate of its powers,

$$r(A) \;=\; \lim_{m\to\infty} \big\|A^m\big\|^{1/m},$$

the spectral-radius formula, equivalently the largest modulus in the spectrum of $A$; a spectral radius strictly below one says exactly that the powers decay geometrically, $\|A^m\| \le C\,\rho^m$ for every $\rho$ strictly between $r(A)$ and one. The two hypotheses are the continuous statement of the pair that defines a *stationary* heterogeneous-agent block in Auclert et al. [2025, Def. 6], where the backward iteration has its Jacobian eigenvalues inside the unit circle and the transition matrix has all eigenvalues except the unit one inside.

**Assumption 2** (Geometric stability)**.** At the steady state,

1. the backward map contracts: the spectral radius of $D_{\bar{x}+\mathfrak{x}}$ on $L^\infty(\Theta \times \mathcal{A})$ is strictly smaller than one;

2. the heterogeneous dynamics mix through the policy $\bar{x}^*$: there are $C > 0$ and $\lambda \in (0,1)$ with

$$\big\|\mathcal{K}^m \bar{x}^* - \langle \bar{\Omega}^*, \bar{x}^* \rangle\big\|_\infty \;+\; \big\|\nabla_a\big(\mathcal{K}^m \bar{x}^*\big)\big\|_\infty \;\le\; C\,\lambda^m \qquad (m \ge 0).$$

   The suprema are essential, over the seam-free set where $\mathcal{K}^m \bar{x}^*$ is classically differentiable, the complement of the iterated pre-images of $\Gamma$ under the per-regime policy (Section 5.2); across those measure-zero seams the gradient jumps but stays bounded, so the clause bounds the per-regime gradients and not a global smoothness the iterates do not have.

Clause 1 is discounting at work: it makes the powers of the backward map decay geometrically, so with Lemma 1, which writes each coefficient as such a power, the policy responses decay too, $\|\times_s\|_\infty \le C\rho^s$ for any $\rho$ strictly between the spectral radius and one, and a shock announced $s$ periods ahead moves today's policy like $\rho^s$. Clause 2 says that the expected policy path of a household forgets its starting state at a geometric rate, uniformly with its slope: $\mathcal{K}^m \bar{x}^*$ is the expectation-vector object $\mathcal{E}_m$ of Auclert et al. [2021] identified in Section 4.4, so the clause asserts for the continuous operator exactly the decay their arrays display; the gradient term is included because the dipoles of $D$ pair through slopes at the states they sit at, $\langle \dot{\delta}_{(\theta,u)}[v],\ \mathcal{K}^m \bar{x}^* \rangle = v \cdot \nabla_a\big(\mathcal{K}^m \bar{x}^*\big)(\theta, u)$ (Section 4.3), and clause 2 is what controls them.

**Proposition 6** (The heterogeneous Jacobian is block quasi-Toeplitz)**.** *Let Assumptions 1 and 2 and the setting of Proposition 2 be in force, write $S_k \equiv -\nabla_{a'} \cdot (\mathsf{M} \cdot p \times_k)$ for the sources of Proposition 5, and fix $\rho$ strictly between the spectral radius of clause 1 of Assumption 2 and one. Then, for a constant $C$ that may change from line to line:*

1. *the increments decay geometrically in both indices,*

$$\big\|\big\langle \Lambda^{t-1} S_s,\ \bar{x}^* \big\rangle\big\| \;\le\; C\,\lambda^{t-1}\rho^s \qquad (t \ge 1,\ s \ge 0);$$

2. *along each diagonal the blocks of the heterogeneous Jacobian converge, $j_k \equiv \lim_{t\to\infty} J_{t,\,t-k}$, with $\|j_k\| \le C \max(\rho,\lambda)^{|k|}$, and the symbol they define is*

$$j(z) \;=\; \big\langle \bar{\Omega}^*,\ \times(1/z) \big\rangle \;+\; \Big\langle S(1/z),\ z\,(I - z\,\mathcal{K})^{-1}\,\big(\bar{x}^* - \langle \bar{\Omega}^*, \bar{x}^* \rangle\big) \Big\rangle,$$

   *where*

$$\times(w) \;\equiv\; \sum_{s\ge 0} \times_s w^s \;=\; \big(I - w\,D_{\bar{x}+\mathfrak{x}}\big)^{-1} D_{\bar{Y}\mathfrak{x}}, \qquad S(w) \;\equiv\; -\nabla_{a'} \cdot \big(\mathsf{M} \cdot p \times(w)\big),$$

   *all three series converging on an annulus $\rho < |z| < 1/\lambda$ containing the unit circle, on which $j$ is analytic;*

3. *the heterogeneous Jacobian is block quasi-Toeplitz, $J = \mathsf{T}(j) + \mathsf{C}$, the correction carrying the missing anticipation of shocks that would have had to be announced before date 0,*

$$\mathsf{C}_{t,s} \;=\; -\sum_{m>t} \big\langle \Lambda^{m-1} S_{m-t+s},\ \bar{x}^* \big\rangle, \qquad \|\mathsf{C}_{t,s}\| \;\le\; C\,\lambda^t \rho^s,$$

   *and it is a Hilbert–Schmidt, hence compact, operator: the bound concentrates it near the entry $t = s = 0$, fading geometrically in the row and in the column.*

*Proof. Step 1: how a source pairs, and its two consequences.* Fix a bounded direction $y$ with $y^i = 0$ on $\{\bar{a}^{*i} = \underline{a}^i\}$ for each bound $i$, and a test function $\psi$, Lipschitz and continuously differentiable near the forward orbits $\mathcal{O}(\theta_n, u_n)$ of the atoms. The source $-\nabla_{a'} \cdot (\mathsf{M} \cdot y)$ is a weak divergence: by definition it pairs with $\psi$ by moving the derivative onto $\psi$, the boundary term dropped by no-flux (Assumption 1), and the displacement pairing of Proposition 3 then evaluates the result,

$$\big\langle -\nabla_{a'} \cdot (\mathsf{M} \cdot y),\ \psi \big\rangle \;=\; \sum_j \big\langle (\mathsf{M} \cdot y)^j,\ \partial_{a'^j} \psi \big\rangle \;=\; \sum_j \iint \partial_{a'^j} \psi\big(\theta', \bar{a}^*(\theta, a)\big)\ y^j(\theta, a)\, d\pi(\theta' \mid \theta)\, d\bar{\Omega}^*(\theta, a).$$

Two consequences follow. The first is a bound: the kernel and the stationary distribution both carry unit mass, so the double integral is at most the supremum of its integrand,

$$\big\| \big\langle -\nabla_{a'} \cdot (\mathsf{M} \cdot y),\ \psi \big\rangle \big\| \;\le\; C\, \|y\|_\infty\, \|\nabla_{a'} \psi\|_\infty,$$

the gradient of $\psi$ read at the landing points $\bar{a}^*(\theta, a)$; the atoms of $\bar{\Omega}^*$ land inside regimes where $\psi$ is continuously differentiable, by clause 4 of Assumption 1, so the essential supremum controls their landing points too. The second consequence is about mass: only the *gradient* of $\psi$ enters the pairing, so a constant test function pairs to zero,

$$\big\langle -\nabla_{a'} \cdot (\mathsf{M} \cdot y),\ c \big\rangle \;=\; 0 \qquad \text{for every constant } c.$$

In words: a divergence only moves mass around, it carries none of its own.

*Step 2: the increments.* The increment $\langle \Lambda^{t-1} S_s, \bar{x}^* \rangle$ pushes the source forward $t-1$ periods and reads the result against the steady-state policy. Move the transfer operators onto the function side by the adjunction of Lemma 2, applied $t-1$ times:

$$\big\langle \Lambda^{t-1} S_s,\ \bar{x}^* \big\rangle \;=\; \big\langle S_s,\ \mathcal{K}^{t-1} \bar{x}^* \big\rangle.$$

The right side is not yet usable, because $\mathcal{K}^{t-1} \bar{x}^*$ does not decay as $t$ grows: it converges to the constant $\langle \bar{\Omega}^*, \bar{x}^* \rangle$, the unconditional mean of the policy, not to zero. What decays, by clause 2 of Assumption 2, is the *centered* iterate, and the zero mass of Step 1 is exactly the licence to center: subtracting a constant from the test function does not change its pairing with a source, so

$$\big\langle S_s,\ \mathcal{K}^{t-1} \bar{x}^* \big\rangle \;=\; \big\langle S_s,\ \mathcal{K}^{t-1} \bar{x}^* - \langle \bar{\Omega}^*, \bar{x}^* \rangle \big\rangle.$$

Now apply the bound of Step 1 with $y = p \times_s$ and $\psi = \mathcal{K}^{t-1} \bar{x}^* - \langle \bar{\Omega}^*, \bar{x}^* \rangle$. The direction is small because the shock is far: $\|\times_s\|_\infty \le C\rho^s$, clause 1 of Assumption 2 bounding the powers of the backward map through the spectral-radius formula and Lemma 1 writing $\times_s$ as such a power. The test function is small because the dynamics mix: clause 2 bounds the centered iterate, gradient included, by $C\lambda^{t-1}$. The product of the two bounds is item 1.

*Step 3: the diagonal limits.* Fix a diagonal $k = t - s$ and unroll the coefficient recursion of Proposition 5 along it,

$$J_{t,\,t-k} \;=\; \big\langle \bar{\Omega}^*,\ \times_{-k} \big\rangle \;+\; \sum_{m=1}^{t} \big\langle \Lambda^{m-1} S_{m-k},\ \bar{x}^* \big\rangle,$$

under the conventions $\times_n = 0$ and $S_n = 0$ for $n < 0$: moving down the diagonal appends new terms to the sum and never alters a term already appended. Item 1 bounds the term at $m$ by $C\lambda^{m-1}\rho^{m-k}$, a summable family, so the series converges absolutely and the diagonal limit exists,

$$j_k \;=\; \big\langle \bar{\Omega}^*,\ \times_{-k} \big\rangle \;+\; \sum_{m=1}^{\infty} \big\langle \Lambda^{m-1} S_{m-k},\ \bar{x}^* \big\rangle.$$

Its size follows from the same bounds, one estimate per side of the main diagonal. On the anticipation side $k \le 0$, both the direct term and every source carry the announcement horizon $-k$,

$$\|j_k\| \;\le\; C\,\rho^{-k} \;+\; C\,\rho^{-k} \sum_{m \ge 1} (\lambda\rho)^{m-1} \;\le\; C\,\rho^{-k};$$

on the propagation side $k > 0$ the direct term vanishes and the first surviving source needs $m \ge k$ pushes,

$$\|j_k\| \;\le\; C \sum_{m \ge k} \lambda^{m-1}\rho^{m-k} \;\le\; C\,\lambda^{k-1}.$$

Together, $\|j_k\| \le C \max(\rho, \lambda)^{|k|}$.

*Step 4: the symbol.* By Lemma 1 the coefficients are the powers of one operator, $\times_s = (D_{\bar{x}^+}\mathfrak{x})^s\, D_{\bar{Y}}\mathfrak{x}$, so they are generated by the single Neumann series

$$\times(w) \;=\; \sum_{s\ge 0} \times_s w^s \;=\; \big(I - w\, D_{\bar{x}^+}\mathfrak{x}\big)^{-1} D_{\bar{Y}}\mathfrak{x},$$

absolutely convergent for $|w| < 1/\rho$ by clause 1, and the sources inherit the same generating function, $S(w) = -\nabla_{a'} \cdot (\mathsf{M} \cdot p \times(w))$, by linearity of the displacement in its direction. Now sum $j(z) = \sum_k j_k z^k$ from the series of Step 3, each increment first written in the centered form of Step 2, and reorganize the double sum by the index $u = m - k$ of the source and the power $m$ of the Koopman operator; the geometric bounds make the double sum absolutely convergent, so the reordering is legitimate, and with $z^k = z^m z^{-u}$,

$$\begin{aligned}\sum_{k\in\mathbb{Z}} z^k \sum_{m\ge 1} \big\langle \Lambda^{m-1} S_{m-k},\ \bar{x}^* \big\rangle \;&=\; \sum_{u\ge 0} z^{-u} \Big\langle S_u,\ \sum_{m\ge 1} z^m\, \mathcal{K}^{m-1}\big(\bar{x}^* - \langle \bar{\Omega}^*, \bar{x}^*\rangle\big)\Big\rangle \\ &=\; \Big\langle S(1/z),\ z\,\big(I - z\,\mathcal{K}\big)^{-1}\big(\bar{x}^* - \langle \bar{\Omega}^*, \bar{x}^*\rangle\big)\Big\rangle,\end{aligned}$$

while the direct terms sum to $\langle \bar{\Omega}^*, \times(1/z)\rangle$. The two resolvents delimit the annulus: the anticipation side, a series in $1/z$, converges for $|z| > \rho$; the propagation side, a series in $z$ whose terms clause 2 bounds by $C(\lambda|z|)^m$, converges for $|z| < 1/\lambda$. On the annulus $\rho < |z| < 1/\lambda$, which contains the unit circle, $j$ is therefore a convergent Laurent series, hence analytic.

*Step 5: the correction, and the two operators.* The correction is the entry-by-entry gap between the operator and its diagonal limits,

$$\mathsf{C}_{t,s} \;=\; J_{t,s} - j_{t-s} \;=\; -\sum_{m>t} \big\langle \Lambda^{m-1} S_{m-t+s},\ \bar{x}^* \big\rangle,$$

minus the terms the diagonal would still append below row $t$: the anticipation a shock announced before date 0 would have supplied. Item 1 bounds the tail,

$$\|\mathsf{C}_{t,s}\| \;\le\; C \sum_{m>t} \lambda^{m-1}\, \rho^{\,m-t+s} \;\le\; C\, \lambda^t \rho^s,$$

geometric in the row and in the column at once. Squaring and summing over all $(t, s)$ therefore gives a finite number: the correction has square-summable entries, so it is a Hilbert–Schmidt operator, and Hilbert–Schmidt operators are compact [Gohberg et al., 1990, Ch. VIII, §2]. The Toeplitz part is bounded on $\ell^2$ with no citation needed: the $k$-th diagonal acts as a shift composed with multiplication by the fixed matrix $j_k$, an operator of norm $\|j_k\|$, and the triangle inequality sums the diagonals, $\|\mathsf{T}(j)\| \le \sum_k \|j_k\| < \infty$ by Step 3. So $J = \mathsf{T}(j) + \mathsf{C}$ with $\mathsf{T}(j)$ bounded and $\mathsf{C}$ compact, which is item 3. □

**Corollary 2** (The stacked system is block quasi-Toeplitz)**.** *Remove the truncation: stack the differentiated aggregate equations of Section 4.4 over all $t \ge 0$ and keep the names $\tilde{J}_X$ and $\tilde{J}_\mathcal{E}$ for the resulting operators on $\ell^2$. Write the relevance slice as a selection,*

$$\widehat{Y}_t \;=\; Q_\mathcal{E}\, \mathcal{E}_t \;+\; Q_-\, P\widehat{X}_{t-1} \;+\; Q_0\, \widehat{X}_t \;+\; Q_+\, \widehat{X}_{t+1},$$

*with constant matrices $Q_\mathcal{E}, Q_-, Q_0, Q_+$ picking each constituent out of $\widehat{Y}_t \subseteq [\mathcal{E}_t, P\widehat{X}_{t-1}, \widehat{X}_t, \mathbb{E}_t\widehat{X}_{t+1}]$. Then both stacked operators are block quasi-Toeplitz,*

$$\begin{aligned}\tilde{J}_X \;&=\; \mathsf{T}(j_X) + \mathsf{C}_X, \qquad j_X(z) \;=\; \big(D_{\bar{Y}}G \;+\; D_x G\; j(z)\big)\,\big(Q_-\, P\, z \;+\; Q_0 \;+\; Q_+\, z^{-1}\big), \\ \tilde{J}_\mathcal{E} \;&=\; \mathsf{T}(j_\mathcal{E}) + \mathsf{C}_\mathcal{E}, \qquad j_\mathcal{E}(z) \;=\; \big(D_{\bar{Y}}G \;+\; D_x G\; j(z)\big)\, Q_\mathcal{E},\end{aligned}$$

*with $\mathsf{C}_X$ and $\mathsf{C}_\mathcal{E}$ compact. The symbol $j_X$ is square, one row per aggregate equation and one column per aggregate variable, and analytic on an annulus containing the unit circle.*

*Proof.* Name the three kinds of operators the stacked equation composes, then multiply them out one product at a time.

*The factors.* The *selection* maps an aggregate path to the path of relevance slices it induces; by the display in the statement it is the block Toeplitz operator $\mathsf{T}(q)$ with the banded symbol

$$q(z) \;=\; Q_-\, P\, z \;+\; Q_0 \;+\; Q_+\, z^{-1},$$

and it needs no correction at the edge: at $t = 0$ the matrix of $\mathsf{T}(q)$ simply has no column $s = -1$ to read, which is exactly the predetermined start $\widehat{X}_{-1} = 0$, so the one-sided time axis implements predeterminedness by itself. The *date-wise derivatives* apply $D_{\bar{Y}}G$ and $D_xG$ date by date: block diagonal operators, hence block Toeplitz with the constant symbols $D_{\bar{Y}}G$ and $D_xG$. The *heterogeneous Jacobian* is $J = \mathsf{T}(j) + \mathsf{C}$ by Proposition 6. Differentiating the aggregate equations date by date and stacking, as in Section 4.4 but without the truncation, the system reads

$$\Big[\mathsf{T}\big(D_{\bar{Y}}G\big) \;+\; \mathsf{T}\big(D_xG\big)\,J\Big]\,\big(\widehat{Y}_t\big)_{t\geq 0} \;=\; 0, \qquad \big(\widehat{Y}_t\big)_{t\geq 0} \;=\; \mathsf{T}(q)\,\big(\widehat{X}_t\big)_{t\geq 0} \;+\; \mathsf{T}\big(Q_{\mathcal{E}}\big)\,\big(\mathcal{E}_t\big)_{t\geq 0},$$

so that $\tilde{J}_X$ and $\tilde{J}_{\mathcal{E}}$ are the two compositions obtained by substituting the second display into the first.

*The product rule.* Two facts multiply the factors out. First, the product of two block Toeplitz operators is block Toeplitz plus compact, with the symbols multiplying: the product identity of Böttcher and Silbermann [1999, Prop. 1.12 and eq. (6.2)] writes $\mathsf{T}(a)\mathsf{T}(b) - \mathsf{T}(ab)$ as a product of two Hankel operators, and Hankel operators with continuous symbols are compact [Böttcher and Silbermann, 1999, Thm. 1.16]; when one factor has a *constant* symbol the Hankel term vanishes outright and the identity is exact. Second, a product with a compact factor on either side is compact, so every term involving a correction stays a correction.

One split fixes every exactness claim that follows, here and in Proposition 9. The factors with a *constant* symbol,

$$D_{\bar{Y}}G, \qquad D_xG, \qquad Q_{\mathcal{E}},$$

multiply exactly and carry no Hankel remnant; the factors merely *continuous* on the circle,

$$j, \qquad q, \qquad j_X, \qquad j_-^{-1},\; j_+^{-1},$$

do carry one, so a Toeplitz product is exact precisely when no two continuous factors meet in it. The selection $q$ is a Laurent polynomial, continuous rather than constant, even though its blocks $Q_-, Q_0, Q_+$ are each constant matrices.

*Multiplying out.* Abbreviate $A \equiv D_{\bar{Y}}G + D_xG\,j$ and apply the two facts to the inner bracket first,

$$\mathsf{T}\big(D_{\bar{Y}}G\big) \;+\; \mathsf{T}\big(D_xG\big)\big(\mathsf{T}(j) + \mathsf{C}\big) \;=\; \mathsf{T}(A) \;+\; \mathsf{C}_1, \qquad \mathsf{C}_1 \;\equiv\; \mathsf{T}\big(D_xG\big)\,\mathsf{C},$$

the Toeplitz product exact because $D_xG$ is a constant symbol, and $\mathsf{C}_1$ compact as a bounded operator times a compact one; then compose with the selection,

$$\big[\mathsf{T}(A) + \mathsf{C}_1\big]\,\mathsf{T}(q) \;=\; \mathsf{T}\big(A\,q\big) \;+\; \mathsf{C}_2, \qquad \mathsf{C}_2 \;\equiv\; \big(\mathsf{T}(A)\,\mathsf{T}(q) - \mathsf{T}(A\,q)\big) \;+\; \mathsf{C}_1\,\mathsf{T}(q),$$

the first summand of $\mathsf{C}_2$ the Hankel product of the cited identity, compact for continuous symbols, the second a compact times a bounded operator; the product symbol $A\,q$ is the displayed $j_X$. The same composition with $\mathsf{T}(Q_{\mathcal{E}})$ in place of $\mathsf{T}(q)$ is exact in its last step, $Q_{\mathcal{E}}$ being constant, and gives $j_{\mathcal{E}}$. The symbol $j_X$ is analytic on an annulus because $j$ is (Proposition 6) and $q$ is a Laurent polynomial, and it is square because the stacked system carries one equation row per aggregate variable, as the counting of Section 1 requires. □

Three readings are worth recording. First, the symbol is the transfer function of the heterogeneous block, in the sense given to infinite-dimensional linear systems by Curtain and Zwart [1995, Def. 4.3.5]: the backward resolvent $\times(1/z) = (I - z^{-1}\,D_{\bar{x}^+}\mathfrak{x})^{-1}D_{\bar{Y}}\mathfrak{x}$ carries anticipation; the Koopman resolvent $z\,(I - z\,\mathcal{K})^{-1}$ carries propagation; and the correction $\mathsf{C}$, with $\|\mathsf{C}_{t,s}\| \leq C\,\lambda^t\rho^s$ (Proposition 6, item 3), is the missing anticipation of shocks that would have needed announcing before date 0, the object Auclert et al. [2025, §2] isolate on the grid. Second, item 1 of Proposition 6 is the continuous mirror of the discrete bound of Auclert et al. [2025, Lem. 2] and item 3 of their Proposition 4, proved there for the discretized block; and through Lemma 5 the increments are also Bhandari et al. [2023]'s velocity compositions, so the quasi-Toeplitz structure is a property of the object, not of the method that computes it. Third, the complementarity constraints leave the structure untouched: whatever bounds the $\perp$-clauses of Section 1 impose, Proposition 2 keeps every response coefficient a bounded function, the decay stays geometric, and the symbol stays continuous on the circle, so the constraints reach determinacy only through the values of $j$.

With the structure in hand, invertibility is classical.

**Proposition 7** (The winding-number criterion)**.** *In the setting of Corollary 2, suppose* $\det j_X(z) \neq 0$ *for every* $z$ *on the unit circle, and set* $w \equiv \operatorname{wind}(\det j_X)$*. Then* $\tilde{J}_X$ *is a* Fredholm *operator, meaning that its kernel is finite-dimensional and its range is closed with finite codimension; the difference of the two counts, the surplus of free solution directions over unreachable right-hand sides, is its* index*, and here*

$$\dim\ker \tilde{J}_X \;-\; \operatorname{codim}\operatorname{ran} \tilde{J}_X \;=\; -w.$$

*In particular:*

1. *if* $w < 0$*, the homogeneous stacked system has at least* $-w$ *independent square-summable solutions: the bounded response is not unique, the model is indeterminate at first order;*
2. *if* $w > 0$*, the range has codimension at least* $w$*: for some right-hand sides no square-summable response exists;*
3. *if* $w = 0$*, the operator is invertible if and only if its kernel is trivial, and this holds generically, on an open dense set of block quasi-Toeplitz operators, a refinement due to Auclert et al. [2025, Props. 8 and 11]; the genericity is in operator space, and its translation to the model's primitives is open, as Auclert et al. note.*

*If instead* $\det j_X$ *vanishes somewhere on the unit circle,* $\tilde{J}_X$ *is not Fredholm.*

*Proof. Step 1: the Toeplitz part.* For a block Toeplitz operator with a continuous matrix symbol the statement is the classical criterion of Gohberg: $\mathsf{T}(j_X)$ is Fredholm if and only if $\det j_X$ has no zeros on the unit circle, in which case

$$\dim\ker \mathsf{T}(j_X) \;-\; \operatorname{codim}\operatorname{ran} \mathsf{T}(j_X) \;=\; -\operatorname{wind}(\det j_X)$$

[Böttcher and Silbermann, 1999, Thm. 6.5].

*Step 2: from the Toeplitz part to the full operator.* A compact perturbation changes neither Fredholmness nor the index [Gohberg et al., 1990, Ch. XI, Thm. 4.2]. Applied to $\tilde{J}_X = \mathsf{T}(j_X) + \mathsf{C}_X$, this carries Step 1's conclusion to $\tilde{J}_X$ whenever $\det j_X$ has no zeros on the circle. Applied in the other direction, to $\mathsf{T}(j_X) = \tilde{J}_X - \mathsf{C}_X$, it gives the final claim: if $\tilde{J}_X$ were Fredholm while $\det j_X$ vanished somewhere, the Toeplitz part would be Fredholm too, against Step 1.

*Step 3: reading the index.* The index is a difference of two nonnegative integers,

$$\dim\ker \tilde{J}_X \;=\; \operatorname{codim}\operatorname{ran} \tilde{J}_X \;-\; w,$$

so a negative $w$ forces $\dim\ker \tilde{J}_X \geq -w$, which is item 1, and a positive $w$ forces $\operatorname{codim}\operatorname{ran} \tilde{J}_X \geq w$, which is item 2.

*Step 4: index zero.* At $w = 0$ the kernel and the cokernel have the same dimension, so the operator is injective exactly when it is surjective, and invertibility reduces to the triviality of the kernel; that this triviality holds on an open dense set of block quasi-Toeplitz operators is the genericity refinement of Auclert et al. [2025, Prop. 8]. □

Proposition 7 counts unreachable right-hand sides but leaves open whether the model ever presents one, and it never presents an arbitrary one: along the equilibrium path the innovation enters only through the forcing $\tilde{J}_{\mathcal{E}}\left(\mathcal{E}_t\right)_{t\geq 0}$, the right-hand side of the stacked system of Corollary 2. The unreachable right-hand sides form a subspace, the *cokernel* $\ker \tilde{J}_X^*$ of the adjoint (the operator with $\langle \tilde{J}_X^*\phi,\, x\rangle = \langle \phi,\, \tilde{J}_X x\rangle$), of dimension $\operatorname{codim}\operatorname{ran} \tilde{J}_X$; a right-hand side is reachable exactly when it is orthogonal to every cokernel direction. Whether the model has a bounded response is therefore whether its own forcing avoids them.

**Proposition 8** (Existence of the response to an innovation)**.** *In the setting of Proposition 7 with* $w > 0$*, a bounded response* $\left(\widehat{X}_t\right)_{t\geq 0} \in \ell^2$ *to the innovation path* $\left(\mathcal{E}_t\right)_{t\geq 0}$ *exists if and only if the forcing is orthogonal to the cokernel,*

$$\left\langle \phi,\; \tilde{J}_{\mathcal{E}}\left(\mathcal{E}_t\right)_{t\geq 0} \right\rangle \;=\; 0 \qquad \text{for every } \phi \in \ker \tilde{J}_X^*.$$

*These are* $\dim\ker \tilde{J}_X + w$ *conditions, at least* $w$ *of them. A positive winding number therefore opens the possibility of nonexistence without settling it: the response fails exactly when the forcing meets an unreachable direction, a matter of the shock and not of* $w$ *alone.*

*Proof.* $\tilde{J}_X$ is Fredholm (Proposition 7), so its range is closed and, on $\ell^2$, equals the orthogonal complement of the cokernel, $\operatorname{ran}\tilde{J}_X = \left(\ker \tilde{J}_X^*\right)^\perp$ [Gohberg et al., 1990, Ch. XIV, Prop. 2.5]; a bounded response solves $\tilde{J}_X\left(\widehat{X}_t\right)_{t\geq 0} = -\tilde{J}_{\mathcal{E}}\left(\mathcal{E}_t\right)_{t\geq 0}$ exactly when the forcing lies there, the displayed orthogonality. The cokernel has dimension $\operatorname{codim}\operatorname{ran}\tilde{J}_X = \dim\ker\tilde{J}_X + w$ by the index $-w$ of Proposition 7, at least $w$ since the kernel dimension is nonnegative. □

The zero-winding case deserves the honesty it rarely gets. For a scalar Toeplitz operator with no compact correction, index zero does imply invertibility, by Coburn's lemma: a nonzero symbol leaves the kernel of the operator or of its adjoint trivial [Böttcher and Silbermann, 1999, Thm. 1.10]. Both hypotheses matter: with matrix blocks the lemma fails, an explicit two-by-two counterexample being Böttcher and Silbermann [1999, Ex. 6.1], and with a compact correction it fails already for scalars: advancing a path one period and then delaying it one period has symbol 1, yet it erases the date-0 entry, a one-dimensional kernel [Auclert et al., 2025, §3]. What saves practice is the pairing of failures: Auclert et al. [2025, Props. 8 and 9] show the non-generic, index-zero, non-invertible case always couples nonexistence with indeterminacy, so checking existence directly rules both out; and the safe direction carries no caveat, since a unique bounded response forces $w = 0$. A nonzero winding number is thus a proof of trouble, a zero one a generic certificate.

**Proposition 9** (Wiener–Hopf factorization and the solution operator)**.** *Suppose* $\det j_X(z) \neq 0$ *on the unit circle and let* $n$ *be the size of* $j_X$*. Then:*

1. *the symbol factors as*

$$j_X(z) \;=\; j_-(z)\;\operatorname{diag}\left(z^{\kappa_1},\dots,z^{\kappa_n}\right)\,j_+(z),$$

   *where* $j_+$ *and its inverse are absolutely convergent series in nonnegative powers of* $z$*,* $j_-$ *and its inverse in nonpositive powers, and the integers* $\kappa_1 \leq \cdots \leq \kappa_n$ *are unique [Böttcher and Silbermann, 1999, Thm. 6.6]; they are the* partial indices*: each diagonal factor* $z^{\kappa_i}$ *is a shift of Fredholm index* $-\kappa_i$*, so the factorization splits the index* $-w$ *of Proposition 7 into one integer per coordinate, and the* $\kappa_i$ *sum to* $w$*;*

2. $\mathsf{T}(j_X)$ *is invertible if and only if every partial index vanishes, in which case* $j_X = j_-\,j_+$ *and*

$$\mathsf{T}(j_X)^{-1} \;=\; \mathsf{T}\left(j_+^{-1}\right)\,\mathsf{T}\left(j_-^{-1}\right);$$

3. *if the full operator* $\tilde{J}_X$ *is invertible, its inverse is block quasi-Toeplitz with symbol* $j_X^{-1}$*, and the untruncated general-equilibrium response* $G = -\tilde{J}_X^{-1}\tilde{J}_{\mathcal{E}}$ *is block quasi-Toeplitz with symbol* $-\,j_X(z)^{-1}\,j_{\mathcal{E}}(z)$*; this is the* infinite *inverse, and the* $T \times T$ *corner the truncated solve returns as* `dr.G` *converges to it only under the finite-section condition of the discussion below, not from invertibility of* $\tilde{J}_X$ *alone.*

*Proof. Step 1: the factorization.* Item 1 is the factorization theorem of Gohberg and Krein for symbols in the Wiener class, absolutely summable coefficients, which the annulus analyticity of Corollary 2 more than grants; the same theorem computes the two defects of the Toeplitz operator, its kernel dimension and its range codimension, from the partial indices,

$$\dim\ker\mathsf{T}(j_X) \;=\; \sum_{\kappa_i<0}|\kappa_i|, \qquad \operatorname{codim}\operatorname{ran}\mathsf{T}(j_X) \;=\; \sum_{\kappa_i>0}\kappa_i$$

[Böttcher and Silbermann, 1999, Thm. 6.6].

*Step 2: the equivalence.* Both defects are sums of nonnegative contributions, one per nonzero partial index, so they vanish together exactly when every $\kappa_i = 0$; and since $\det j_X$ has no zeros on the circle the operator is Fredholm with closed range (Step 1 of Proposition 7), so it is invertible exactly when both defects vanish. This is the equivalence of item 2, and at $\kappa_1 = \cdots = \kappa_n = 0$ the middle factor of item 1 is the identity, $j_X = j_-\,j_+$.

*Step 3: the inverse of the Toeplitz part.* One classical identity powers the computation: a Toeplitz product with only nonpositive powers in its left factor, or only nonnegative powers in its right factor, carries no correction, $\mathsf{T}(ab) = \mathsf{T}(a)\mathsf{T}(b)$ exactly [Böttcher and Silbermann, 1999, Prop. 1.13]. Applied to the canonical factorization it gives $\mathsf{T}(j_X) = \mathsf{T}(j_-)\mathsf{T}(j_+)$. Now multiply the candidate inverse against

$\mathsf{T}(j_X)$ and contract one adjacent pair at a time, each contraction another exact instance of the same identity, the inverses $j_\pm^{-1}$ staying one-sided by item 1:

$$\mathsf{T}(j_+^{-1})\,\mathsf{T}(j_-^{-1})\,\mathsf{T}(j_-)\,\mathsf{T}(j_+) \;=\; \mathsf{T}(j_+^{-1})\,\mathsf{T}(j_-^{-1}\,j_-)\,\mathsf{T}(j_+) \;=\; \mathsf{T}(j_+^{-1})\,\mathsf{T}(j_+) \;=\; \mathsf{T}(j_+^{-1}\,j_+) \;=\; I,$$

and symmetrically in the other order, $\mathsf{T}(j_-)\mathsf{T}(j_+)\mathsf{T}(j_+^{-1})\mathsf{T}(j_-^{-1}) = I$. This proves the displayed inverse formula.

*Step 4: the inverse of the full operator, and the response.* First, the symbol $j_X^{-1}$ is itself admissible: $\det j_X$ is continuous and nonvanishing on the circle, hence nonvanishing on a slightly thinner annulus, on which $j_X^{-1}$ is analytic; its coefficients therefore decay geometrically, and $\mathsf{T}(j_X^{-1})$ is a bounded block Toeplitz operator. Second, $\mathsf{T}(j_X^{-1})$ inverts $\tilde{J}_X$ up to a compact operator, by the product rule of Corollary 2's proof:

$$\tilde{J}_X\,\mathsf{T}(j_X^{-1}) \;=\; \big(\mathsf{T}(j_X) + \mathsf{C}_X\big)\,\mathsf{T}(j_X^{-1}) \;=\; I \;+\; \mathsf{C}_3, \qquad \mathsf{C}_3 \;\equiv\; \big(\mathsf{T}(j_X)\,\mathsf{T}(j_X^{-1}) - I\big) \;+\; \mathsf{C}_X\,\mathsf{T}(j_X^{-1}),$$

the first summand of $\mathsf{C}_3$ compact by the product identity, $j_X\,j_X^{-1} = I$ making it the Hankel product, and the second compact times bounded. If $\tilde{J}_X$ is invertible, multiply this display by $\tilde{J}_X^{-1}$ on the left and rearrange,

$$\tilde{J}_X^{-1} \;=\; \mathsf{T}(j_X^{-1}) \;-\; \tilde{J}_X^{-1}\,\mathsf{C}_3,$$

and the last term is again compact: the inverse is block quasi-Toeplitz with symbol $j_X^{-1}$, with no need for the Toeplitz part $\mathsf{T}(j_X)$ itself to be invertible. The response operator follows by substituting this inverse and expanding,

$$G \;=\; -\,\tilde{J}_X^{-1}\,\tilde{J}_{\mathcal{E}} \;=\; -\Big(\mathsf{T}(j_X^{-1}) - \tilde{J}_X^{-1}\,\mathsf{C}_3\Big)\Big(\mathsf{T}(j_{\mathcal{E}}) + \mathsf{C}_{\mathcal{E}}\Big) \;=\; \mathsf{T}\big(-\,j_X^{-1}\,j_{\mathcal{E}}\big) \;+\; \mathsf{C}_G,$$

with

$$\mathsf{C}_G \;\equiv\; -\big(\mathsf{T}(j_X^{-1})\,\mathsf{T}(j_{\mathcal{E}}) - \mathsf{T}(j_X^{-1}\,j_{\mathcal{E}})\big) \;-\; \mathsf{T}(j_X^{-1})\,\mathsf{C}_{\mathcal{E}} \;+\; \tilde{J}_X^{-1}\,\mathsf{C}_3\,\tilde{J}_{\mathcal{E}},$$

each summand compact, the first a Hankel product, the second bounded times compact, the third bounded times compact times bounded, which is item 3. □

The factorization is the payoff of the whole linear-time-invariant reading, because its two factors are the two halves of this document. $\mathsf{T}(j_+^{-1})$ is lower triangular, a causal convolution, letting the past act on the present; $\mathsf{T}(j_-^{-1})$ is upper triangular, an anti-causal convolution, sweeping the announced future back to today. The solved response therefore acts in two passes: an anti-causal pass that aggregates anticipation, the general-equilibrium image of the anti-causal policy convolution $\hat{x}_t = \sum_{s\geq 0} \times_s \widehat{Y}_{t+s}$ of Section 4.2, and a causal pass that accumulates propagation, the image of the forward law of motion of Proposition 4; in the representative-agent world the two passes are the familiar pair, expectational roots solved forward and predetermined ones backward. The same factorization that certifies determinacy constructs the solution operator, and the winding number is the obstruction it can meet: read informally, each nonzero partial index is a power of $z$ that no pair of invertible one-sided factors can absorb, a boundary condition missing or in excess. Determinacy pins only their sum $w = \sum_i \kappa_i = 0$; the two-pass inverse of Proposition 9 and the cheap cascade of Proposition 12 need every $\kappa_i = 0$, a *canonical* factorization, which is strictly stronger, since the block case admits $w = 0$ with, say, $\kappa = (1, -1)$, the middle factor $\mathrm{diag}(z, z^{-1})$ then not the identity and only the general, harder factorization serving. Computing the partial indices themselves is a genuinely hard problem [Böttcher and Silbermann, 1999, §6.1]; their sum, the winding number, is the computable part, and by the genericity above it is also the decisive part for determinacy.

Two classical bridges close the theory. The first is to the root-counting tradition. For a representative-agent model the heterogeneous term drops from Corollary 2 and $j_X$ is a matrix Laurent polynomial, so $\det j_X$ is a rational function of $z$ whose only pole inside the disc sits at the origin; the argument principle [Brown and Churchill, 2014, §93] converts the winding number into a zero count net of the pole order, and the zero-winding condition becomes the equality of a stable-root count with a predetermined-variable count, the condition of Blanchard and Kahn [1980]. The equivalence is a theorem: Onatski [2006] proves it for linear rational-expectations models, introducing the winding-number test to economics, and Auclert et al. [2025, App. B.8] restate it in the sequence space; the tradition continues in Meyer-Gohde [2026], who no longer only tests but solves and estimates linear models wholly in the frequency domain, seeking

the solution as an analytic function on the disc. The gain of the winding form is that it consumes only the impulse-response coefficients, never the eigenvalues of a transition matrix, which is what lets it survive the passage to heterogeneous agents, where the state carries a distribution; a discretized state-space route can still count eigenvalues, Reiter-style, but at the cost of the very high-dimensional eigenproblem the sequence space was designed to avoid.

The second bridge is to the truncation Dynare 7 actually performs, and the mathematics of that step also has a name, the finite-section method. Convergence of the $T \times T$ corner solutions to the true response is not automatic: it is equivalent to the corners being invertible, for $T$ large, with inverses that converge [Böttcher and Silbermann, 1999, Prop. 2.2], and the classical theory answers exactly for the operators here. The corner of $\tilde{J}_X = \mathsf{T}(j_X) + \mathsf{C}_X$ is the truncated Toeplitz matrix plus the truncated correction, and for that class the criterion is complete: the corner sequence is stable if and only if both $\tilde{J}_X$ itself and the *associated* pure Toeplitz operator, built on the reflected symbol $z \mapsto j_X(1/z)$, are invertible [Böttcher and Silbermann, 1999, Thm. 6.10, with Thms. 2.11 and 6.9 as the pure-Toeplitz antecedents]. The reflected condition is the one addition, and it is mild: reflecting the symbol runs the curve backwards, so the winding number flips sign and the associated operator is Fredholm of index $-w$; at $w = 0$ it is invertible outright in the scalar case, by Coburn's lemma again, and generically in the block case, the generic reading being the one Auclert et al. [2025, Prop. 12] record, once more a genericity in operator space rather than in model parameters (Proposition 7, item 3). Determinacy is therefore what licenses the truncation of Section 4.5: when $w = 0$ and the model is determinate, the corner solve converges as $T$ grows, its residual error the *spurious missing anticipation* of Auclert et al. [2025, §4.1], the corner forcing every response to vanish beyond $T$ and thereby erasing anticipation the infinite system carries; when $w \neq 0$, no horizon repairs the corner, and the failure shows up not at any single $T$, where the corner may look perfectly well conditioned, but across horizons, as inverses that grow or solutions that drift. The winding number, not the conditioning of one corner, is the sound diagnostic.

The diagnostic is cheap on the arrays Dynare 7 already forms. The arrays $\mathcal{F}_{t,s}$, `transition_matrices` of Section 4.5, hold the impact responses $\langle \bar{\Omega}^*, \times_s \rangle$ in their first row and the increments $\langle \Lambda^{t-1} S_s, \bar{x}^* \rangle$ in their body, so summing each array along its diagonals is exactly the diagonal-limit construction in Step 3 of the proof of Proposition 6, truncated at the horizon: it recovers the symbol blocks $j_k$ up to a geometrically small tail. The aggregate-block derivatives and the selections then assemble $j_X$ pointwise, a fast Fourier transform evaluates $\det j_X(e^{i\omega})$ on a frequency grid, and the winding number is read off the signed crossings of the positive real axis by the determinant's path around the origin, as Auclert et al. [2025, App. A.1] implement it. A natural extension of the toolkit is therefore a determinacy check reported by `heterogeneity_solve` alongside `dr.G`, at a cost negligible next to the solve it would certify; as argued above, the truncated solve alone cannot provide one. Appendix I makes this check executable on the operators the solver already forms and records its verdict on the two shipped models.

The check reads the discretized symbol, so it owes one guarantee: that its verdict is the continuous model's, not an artifact of the mesh or the horizon. This is a question the continuous reading raises and Auclert et al. [2025], working on the discretized block from the outset, do not; it has a finite-sample answer.

**Proposition 10** (The determinacy verdict is the continuous model's)**.** *In the setting of Corollary 2 and Proposition 7, suppose* $\det j_X \neq 0$ *on the unit circle, and write*

$$m_0 \;\equiv\; \min_{|z|=1} \left| \det j_X(z) \right| \;>\; 0.$$

*Let* $j_X^{(h,T)}$ *be the symbol the diagnostic assembles: the blocks* $j_k$ *recovered by summing the arrays* $\mathcal{F}_{t,s}$ *along their diagonals, truncated at horizon* $T$*, on the policy mesh of width* $h$*. Then there are* $h_0 > 0$ *and* $T_0$ *such that, for every* $h < h_0$ *and* $T > T_0$*,* $\det j_X^{(h,T)}$ *has no zero on the unit circle and*

$$\operatorname{wind}\left(\det j_X^{(h,T)}\right) \;=\; \operatorname{wind}\left(\det j_X\right) \;=\; w.$$

*The verdict is exact, not asymptotic, as soon as the checkable inequality*

$$L \left\| j_X^{(h,T)} - j_X \right\|_\infty \;<\; m_0$$

*holds,* $L$ *a Lipschitz constant of* $z \mapsto \det j_X(z)$ *over the symbols' shared range and* $\|\cdot\|_\infty$ *the supremum over the unit circle; its left side is controlled by the mesh width* $h$ *and the truncation tail* $\max(\rho, \lambda)^T$*, its right side read off the same frequency grid as* $w$*. The result holds at any number* $d_a$ *of predetermined coordinates.*

*Proof. Step 1: the block error splits.* Only the heterogeneous symbol $j(z) = \sum_k j_k z^k$ is discretized; the aggregate-block derivatives $D_{\bar{Y}}G, D_xG$ and the selection $q$ that Corollary 2 composes with it are exact. The error therefore lives in the blocks and splits into truncation and discretization,

$$\|j_k^{(h,T)} - j_k\| \;\le\; \|j_k^{(h,T)} - j_k^{(T)}\| \;+\; \|j_k^{(T)} - j_k\|,$$

$j_k^{(T)}$ the diagonal sum stopped at $m \le T$.

*Step 2: the truncation tail is uniform.* By Step 3 of Proposition 6, $j_k = \langle \bar{\Omega}^*, \times_{-k}\rangle + \sum_{m\ge 1}\langle \Lambda^{m-1} S_{m-k}, \bar{x}^*\rangle$, and $j_k^{(T)}$ drops the terms past $m = T$; item 1 of that proposition bounds the tail,

$$\|j_k^{(T)} - j_k\| \;\le\; C \sum_{m>T} \lambda^{m-1}\rho^{m-k} \;\le\; C'\,\eta^T \sigma^{|k|}, \qquad \eta = \max(\rho,\lambda) < 1,\ \sigma < 1,$$

geometrically small, uniformly in $k$.

*Step 3: the discretization error vanishes.* Each diagonal sum is a finite combination of the aggregate increments $\langle \Lambda^{t-1} S_s, \bar{x}^*\rangle$, which the pipeline of Section 4.5 computes at the rate $O(h)$ uniformly in the horizon (Lemma 6), so $\|j_k^{(h,T)} - j_k^{(T)}\| \to 0$ as $h \to 0$; through the gradient control of clause 2 of Assumption 2 the discrete blocks inherit the geometric bound $\|j_k^{(h,T)}\| \le C\sigma^{|k|}$ of item 2 of Proposition 6, uniformly in $h$ and $T$.

*Step 4: the symbol converges uniformly on the circle.* Fix $\varepsilon > 0$ and split the Laurent series at a radius $K$,

$$\|j_X^{(h,T)} - j_X\|_\infty \;\le\; \|D_xG\|\,\|q\|_\infty \Big( \sum_{|k|\le K} \|j_k^{(h,T)} - j_k\| \;+\; \sum_{|k|>K} \big(\|j_k^{(h,T)}\| + \|j_k\|\big)\Big).$$

The geometric bounds of Steps 2–3 push the $|k| > K$ tail below $\varepsilon/3$ once $K = K(\varepsilon)$; Step 2 pushes the truncation part of the finite band below $\varepsilon/3$ for $T$ large; Step 3 pushes the discretization part of the finitely many remaining blocks below $\varepsilon/3$ for $h$ small. Hence $\|j_X^{(h,T)} - j_X\|_\infty \to 0$.

*Step 5: the winding number is stable.* The determinant is a polynomial in the entries of its matrix, hence Lipschitz on the bounded set the symbols occupy, so $\sup_{|z|=1} |\det j_X^{(h,T)} - \det j_X| \le L\,\|j_X^{(h,T)} - j_X\|_\infty$. Choose $h, T$ past the thresholds of Step 4 so this supremum falls below $m_0$. Two continuous closed curves on the circle that stay everywhere closer than the distance of the smaller to the origin wind the origin equally, by the argument principle [Brown and Churchill, 2014, §93]: the perturbation cannot carry a crossing across the origin. So $\det j_X^{(h,T)}$ is itself nonvanishing and $\mathrm{wind}(\det j_X^{(h,T)}) = \mathrm{wind}(\det j_X) = w$. The threshold is the displayed inequality, and the winding number being an integer, once the gap sits below the floor $m_0$ there is no residual error. □

The finite frequency grid is a resolution, not a source of error. Because $\det j_X$ is analytic and nonvanishing on an annulus containing the unit circle (Corollary 2), the phase $\omega \mapsto \arg\det j_X(e^{i\omega})$ is real-analytic with a bounded derivative,

$$M \;\equiv\; \Big\| \tfrac{(\det j_X)'}{\det j_X} \Big\|_\infty \;\le\; \frac{\|\det j_X\|_\infty}{m_0\,\delta},$$

the last inequality Cauchy's estimate with $\delta$ the distance from the circle to the annulus boundary, fixed by $\rho$ and $\lambda$. A grid of $N > 2M$ nodes then turns the phase by less than $\pi$ per cell, so accumulating the signed increments recovers the winding exactly, and $\min_n |\det j_X(e^{i\omega_n})|$ reads the floor $m_0$ to $O(1/N)$. The determinacy grid is set by the same decay rates the truncation tail already estimates, and refining it removes the last discretization of the check.

The same complex analysis that decides determinacy also accelerates the algebra, because every large array of the pipeline has a small frequency-domain twin. Take the Toeplitz part $\mathsf{T}(j_X)$ first. Truncated at the horizon, the $T \times T$ array is carried entirely by the $2T-1$ symbol blocks $j_{-(T-1)}, \dots, j_{T-1}$, and its action on an aggregate path is a convolution against them,

$$\Big(\mathsf{T}(j_X)\,(\widehat{X}_s)_{s\ge 0}\Big)_t \;=\; \sum_s j_{t-s}\,\widehat{X}_s,$$

which one fast Fourier transform pass evaluates in $O(T \log T)$ operations where the dense matrix-vector product costs $O(T^2)$. Composition works the same way: by the product rule of Corollary 2, composing

two Toeplitz parts $\mathsf{T}(j_1)\,\mathsf{T}(j_2)$ amounts to multiplying the two small matrices $j_1(e^{i\omega})$ and $j_2(e^{i\omega})$ at each frequency $\omega$ of the grid, the Hankel leftovers joining the compact correction. Even inversion, ordinarily the expensive step, turns pointwise: by Proposition 9 the Toeplitz part of $\tilde{J}_X^{-1}$ has symbol $j_X(z)^{-1}$, so it is assembled by inverting, frequency by frequency, the same small matrices $j_X(e^{i\omega})$ whose determinant the winding count above already evaluated: the determinacy test and the fast solver consume one object.

The corrections are the part the frequency domain cannot see, and they are exactly the part Proposition 6 made small. Its tail bound $\|\mathsf{C}_{t,s}\| \le C\,\lambda^t\rho^s$ says the entries of $\mathsf{C}_X$ fade geometrically away from the upper-left corner: keeping only the rows $t \le t_\varepsilon$ and the columns $s \le s_\varepsilon$ at which $\lambda^{t_\varepsilon}$ and $\rho^{s_\varepsilon}$ reach a tolerance $\varepsilon$ discards an error of order $\varepsilon$, and a matrix supported on a corner of side $t_\varepsilon \sim \log(1/\varepsilon)$ has rank $r$ at most that side, so the numerical rank of the correction grows only logarithmically in the accuracy demanded. A factorization $\mathsf{C}_X \approx UV^\top$ into two thin $T \times r$ matrices then stores the correction in $2Tr$ numbers and applies it in $O(Tr)$ operations. Auclert et al. [2025, §§4–5 and App. A.3] assemble these pieces into practice, on the quasi-Toeplitz matrix arithmetic of Bini et al. [2019]: they precondition the truncated system with $\mathsf{T}(j_X^{-1})$, the pointwise-inverted symbol, optionally augmented by a small corner of $\mathsf{C}_X$, and a Krylov iteration, a solver that only ever multiplies by the matrix, then converges in a handful of steps. The factorization of Proposition 9 suggests one refinement: at a canonical factorization $j_X = j_-\,j_+$ the product $\mathsf{T}(j_+^{-1})\,\mathsf{T}(j_-^{-1})$ inverts the Toeplitz part exactly, one anti-causal and one causal convolution in place of an approximate preconditioner. For a scalar symbol the two factors are cheap to build, split the Fourier series of $\log j_X$ into its nonnegative and nonpositive halves and exponentiate each, the construction inside the factorization theorem [Böttcher and Silbermann, 1999, Thm. 1.14]; for blocks the partial indices return. None of this asks for new computations: the symbol $j_X$ and the correction $\mathsf{C}_X$ were already assembled for the determinacy check, and the diagnostic and the fast algebra are the same complex analysis read twice. The factorization is also consumed once more, unchanged, by every higher order of perturbation (Proposition 12).

## 5.2 Higher-order perturbation

The dynamics of Section 4 computed the first-order term of the moving-average representation of Section 4.1 and stopped at the $O(\sigma^2)$ remainder; the analysis of Section 5.1 then certified and factorized the operator behind that computation. This subsection records what the same objects yield at every order of perturbation, and the pattern is uniform: the order-$k$ term solves a linear system in the operators the first order already factored, the heterogeneous $J$ of Proposition 1 and the stacked $\tilde{J}_X$ of Corollary 2, with a right-hand side assembled from the orders below. Nothing new is ever inverted, the determinacy certificate is paid once and covers every order, and the marginal cost of an order is the assembly of its right-hand side. At the heterogeneous level this carries the second-order construction of Bhandari et al. [2023] to arbitrary order, on the representative-agent higher-order perturbation of Schmitt-Grohé and Uribe [2004] and Lan and Meyer-Gohde [2013]; the order-$k$ statement, the symbol form, and the constraint-layer corrections are what the sequence-space reading adds. Dynare 7 ships the first order, but the preprocessor already emits the second derivatives of the heterogeneous residual, the file `dynamic_het1_g2.m` of Section 1, so the ingredients of the order-two step are on disk. Throughout, $k$ is the target order; orders two and three appear as the worked specializations that reveal the pattern.

**The expansion in the scale of risk.** Make the scaling of Section 4.1 explicit, $\mathcal{E}_t = \sigma\,\bar{\mathcal{E}}_t$, with unit innovations $\bar{\mathcal{E}}_t$ independent and identically distributed, mean zero, unit variance, and with *vanishing odd moments* up to the order of the expansion, $\mathbb{E}\,\bar{\mathcal{E}}^{\,r} = 0$ for odd $r \le k$, the Gaussian being the leading case; the risk scale is a standard deviation, so its domain is $\sigma \in [0, 1)$ and every $\sigma$-derivative below is one-sided at zero. Where the preceding sections solved the perfect-foresight economy, which receives a known innovation path, the stochastic economy anticipates the law of future innovations, and its date-$t$ solution is a function of the risk scale and of the innovations realized up to $t$,

$$X_t \;=\; \mathfrak{X}\big(\sigma;\; \mathcal{E}_t,\, \mathcal{E}_{t-1},\, \dots\big),$$

the heterogeneous-agent image of the nonlinear infinite moving average policy function of Lan and Meyer-Gohde [2013, eq. (2)]. The map $\mathfrak{X}$ is the object this subsection expands, the gothic capital echoing the backward operator $\mathfrak{x}$; I write $\partial\mathfrak{X}/\partial\mathcal{E}_{-j}$ for the derivative in the innovation slot lagged $j$ periods, always evaluated at the base point $(\sigma, \mathcal{E}) = (0, 0)$, the deterministic steady state.

A running illustration accompanies every proof of this subsection, chosen so that each abstract object can be written out by hand: the one-equation economy

$$X_t \;=\; \beta\,\mathbb{E}_t\big[\varphi(X_{t+1})\big] \;+\; Z_t, \qquad Z_t \;=\; Z^* + \mathcal{E}_t,$$

with no heterogeneous block, a discount factor $\beta \in (0,1)$, and a smooth aggregate feedback $\varphi$ with $|\beta\varphi'| < 1$ at the steady state, the derivative $\varphi'$ always taken there. Differentiating once and stacking as in Corollary 2, its first-order rows read $\widehat{X}_t - \beta\varphi'\,\widehat{X}_{t+1} = \widehat{Z}_t$, so its symbol, winding number, and factorization are computable at sight:

$$j_X(z) \;=\; 1 - \beta\varphi'\, z^{-1} \;=\; \frac{z - \beta\varphi'}{z}, \qquad \operatorname{wind}(\det j_X) \;=\; 1 - 1 \;=\; 0, \qquad j_- = j_X, \quad j_+ = 1,$$

one zero and one pole of the determinant inside the disc, and a factorization whose causal factor is trivial: the economy is purely anticipatory, everything in it is expectation and nothing is a predetermined state. Solving the rows forward, the first-order response to a unit innovation at date $s$ is $(\beta\varphi')^{s-t}$ at every date $t \le s$ and zero afterwards, a sequence worth keeping in sight below.

**Proposition 11** (The expansion in the scale of risk)**.** *Let the unit innovations be independent, identically distributed, mean zero, unit variance, with vanishing odd moments up to order $k$; let the perfect-foresight stacked system be determinate (Proposition 7 with $w = 0$ and trivial kernel); and let the stochastic solution $\mathfrak{X}$ exist on a neighbourhood of the base point and be $k$ times continuously differentiable there. Then*

$$X_t \;=\; \bar{X}^* \;+\; \sum_{m=1}^{k} \frac{1}{m!} \sum_{\substack{q+n=m \\ q \text{ even}}} \binom{m}{q} \sigma^q \sum_{j_1,\dots,j_n \ge 0} \widehat{X}^{(q)}_{j_1\cdots j_n}\, \mathcal{E}_{t-j_1} \cdots \mathcal{E}_{t-j_n} \;+\; O\big(\sigma^{k+1}\big),$$

*with kernels*

$$\widehat{X}^{(q)}_{j_1\cdots j_n} \;\equiv\; \frac{\partial^{\,q+n}\mathfrak{X}}{\partial\sigma^q\, \partial\mathcal{E}_{-j_1} \cdots \partial\mathcal{E}_{-j_n}}\,(0;0),$$

*symmetric in the lags, and:*

1. *the pure shock kernels, $q = 0$, are perfect-foresight objects: $\widehat{X}^{(0)}_{j_1\cdots j_n}$ is the derivative of the deterministic transition path with respect to unanticipated innovations landing $j_1, \dots, j_n$ periods before the reading date, computable by the apparatus of Section 4 alone; in particular $\widehat{X}^{(0)}_j = \widehat{X}_j$;*
2. *every kernel with $q$ odd vanishes (proved below at $q = 1$, and for all odd $q$ in Proposition 13, item 1);*
3. *the kernels with $q \ge 2$ even, the risk corrections, are determined by the same stacked operator as the first order, with the sources of Proposition 13.*

*Proof.* At the representative-agent level the proposition and its proof are those of Lan and Meyer-Gohde [2013]; the steps below record what the heterogeneous sequence space changes, which operator closes each equation and why it is invertible. Throughout, $\mathfrak{X}$ is a function of the scalar $\sigma$ and of countably many finite-dimensional innovation slots, so every derivative is an ordinary partial derivative, and the hypothesis of $k$-fold continuous differentiability lets mixed partials commute and legitimates the Taylor polynomial.

*Step 1: at $\sigma = 0$, $\mathfrak{X}$ is the deterministic transition map.* At $\sigma = 0$ every future innovation is zero almost surely, so each conditional expectation collapses to an evaluation at the zero future path, and the stochastic equilibrium conditions coincide, date by date, with those of a perfect-foresight economy hit by a sequence of unannounced innovations. Determinacy turns this equality of conditions into an equality of solutions: by Proposition 7, item 3, with the trivial kernel assumed, the bounded perfect-foresight path from a given cross-section and innovation is locally unique, so $\mathfrak{X}(0;\cdot)$ *is* the deterministic transition map. Every derivative taken purely in the innovation slots at the base point is then a derivative of that map, which is item 1; at one derivative, $\widehat{X}^{(0)}_j = \widehat{X}_j$. In the running illustration, $\mathfrak{X}(0;\mathcal{E}_t, \mathcal{E}_{t-1}, \dots)$ is the value $X_t$ of the path solving $X_\tau = \beta\varphi(X_{\tau+1}) + Z^* + \mathcal{E}_\tau$ forward from every date $\tau$, given the realized innovations and zero future ones.

*Step 2: the first risk derivative is zero, by moment-homogeneity.* Once the realized innovations are given as arguments, $\sigma$ lives in one place only, the integrated one-period-ahead innovation $\sigma\bar{\mathcal{E}}_{t+1}$ inside each date-$t$ conditional expectation. Write $\mathcal{G}_t$ for the date-$t$ row of the stacked system of Corollary 2; each row holds identically in $\sigma$, so its $\sigma$-derivative at the base point vanishes and collects the two channels,

$$0 \;=\; \frac{d}{d\sigma}\,\mathcal{G}_t\Big|_0 \;=\; \big(\tilde{J}_X\,\widehat{X}^{(1)}\big)_t \;+\; \big(D_{\mathcal{E}_{t+1}}\mathcal{G}_t\big)\Big|_{\mathcal{Z}^*} \mathbb{E}\,\bar{\mathcal{E}}_{t+1} \qquad\Longrightarrow\qquad \tilde{J}_X\,\widehat{X}^{(1)} \;=\; 0\,:$$

the first term is the slope against the unknown $\sigma$-derivative of the path, by construction the stacked operator acting on the kernel; the second differentiates through the integrated innovation, the impact loading of the row, a fixed steady-state vector, times the mean of what was integrated, and that mean is zero by hypothesis. The surviving system is linear and homogeneous, and the assumed invertibility forces $\widehat{X}^{(1)} = 0$. This is eqs. (20)–(22) of Lan and Meyer-Gohde [2013] run in the sequence space: their no-unit-roots matrix $f_y + f_{y^-} + f_{y^+}$, the sum of a row's coefficients, is the symbol at the point $z = 1$, $j_X(1)$, and the stacked $\tilde{J}_X$ stands in its place here; the technique is that of eq. (6) of Schmitt-Grohé and Uribe [2004]. The integrated moment vanishes: in their words, opening the expansion to a moment of the future distribution of shocks changes nothing if that moment is exactly zero [Lan and Meyer-Gohde, 2013, following eq. (22)]. The general odd-$q$ case, the higher odd moments in the role of $\mathbb{E}\,\bar{\mathcal{E}}$, is item 1 of Proposition 13.

*Step 3: the display.* The display is the $k$-th Taylor polynomial of $\mathfrak{X}$ jointly in all its arguments: each choice of $q$ slots for $\sigma$ out of $m$ carries the binomial weight $\binom{m}{q}$, the innovations fill the remaining $n = m - q$ slots with their lags summed over, and the odd-$q$ groups vanish by Step 2 and item 1 of Proposition 13. Every coefficient is a one-sided derivative at $\sigma = 0$, the only ones the domain $[0, 1)$ defines, and the kernels may be taken symmetric in the lags because mixed partials commute. At the representative-agent level the display is the $M$th-order approximation of Lan and Meyer-Gohde [2013, Cor. 2.4]. □

At $k = 1$ the proposition is certainty equivalence: item 2 kills the linear-in-$\sigma$ constant, item 1 identifies the surviving coefficients with the perfect-foresight impulse responses $\widehat{X}_j$, and the display collapses to the moving average of Section 4.1; at the representative-agent level the zero derivatives behind this are Theorem 1 of Schmitt-Grohé and Uribe [2004, p. 763]. At $k = 2$ three new objects appear: the quadratic kernels $\widehat{X}^{(0)}_{j_1 j_2}$, whose diagonal entries bend the response in the innovation's own size and whose off-diagonal entries let the response depend on the history the innovation lands on; and one constant, $\frac{1}{2}\sigma^2\,\widehat{X}^{(2)}$, the risk constant, the first mark of anticipated variance on the solution. For the heterogeneous framework this quadratic expansion, with its risk block, is Lemma 2 of Bhandari et al. [2023, eq. (12)]. The general vanishing of item 2 was conjectured from numerical experiments in Schmitt-Grohé and Uribe [2004, fn. 10, p. 764]; the moment argument of Proposition 13 proves it and delimits it at once, for the odd *rungs*, the kernel families $\widehat{X}^{(q)}_{j_1 \cdots j_n}$ read at fixed $(q, n)$ with $q$ odd, are not dead letters: they revive exactly when the odd moments do, third moments entering at order three [Lan and Meyer-Gohde, 2013, eq. (46) and fn. 23].

The proposition names the targets, the kernels $\widehat{X}^{(q)}_{j_1 \cdots j_n}$; the rest of the subsection computes them. Lemma 8 differentiates the heterogeneous block once more, order after order, always solving the same operator $J$. Lemma 7 pins the geometry of the binding boundary $\Gamma$, Lemma 10 and Corollary 3 supply the calculus for what those derivatives do there, and the law of motion of the cross-section propagates the order-$m$ policy kernels through $\Lambda$ and the displacements $\mathsf{M}, \mathsf{M}_2, \ldots$ Proposition 12 closes general equilibrium in the same stacked operator $\tilde{J}_X$, symbol included; Proposition 13 adds the even risk rungs; and the algorithm paragraph collects the pieces in execution order.

**The heterogeneous block at order $k$.** The pure shock kernels being perfect-foresight objects, the apparatus of Section 4 computes them by differentiating once more, and then again; the general statement costs no more than the second order once the derivative of a composition is organized by partitions. The vocabulary of the constraint, on top of Definition 1, needs stating before the recursions use it: first the geometry of the binding boundary $\Gamma$, then the layers it carries.

**The geometry of the binding boundary.** Definition 1 produced each $\Gamma_b$ as the zero set of the Lipschitz function $\phi_b = \mu^* - g^*$; the calculus below differentiates across $\Gamma_b$ while a perturbation moves it, and for that it needs each $\Gamma_b$ to be, per Markov row, a genuinely smooth hypersurface carrying a unit normal $\mathbf{n}$ and a surface measure $d\Gamma$. That regularity is not free: it comes out of the model, through one new object per bound, the policy the household would choose if the bound were not there.

**Definition 5** (Virtual policy and virtual gap)**.** Fix a bound $b$ on the coordinate $a^i$ and a point of its boundary $\Gamma_b$; away from the other bounds' boundaries their statuses are locally constant, and the constructions below carry them as frozen. Release bound $b$: drop its multiplier and its slackness row, keep every other row of the unfolded steady-state system, and solve the resulting smooth square system on a ball around the point. The solution, delivered by the implicit function theorem under the strong regularity of Proposition 2 (Step 1 of the proof of Lemma 7 supplies the Jacobian), is the *virtual policy*

$\tilde{a}^*(\theta, a)$, a $C^k$ function of the state: what the household would choose if bound $b$ were ignored. It agrees with the true policy $\bar{a}^*$ on the interior regime and continues it smoothly past the frontier. The *virtual gap* is its slack in the released bound,

$$\tilde{g}(\theta, a) \;=\; \tilde{a}^{*\,i}(\theta, a) \;-\; \underline{a}^i,$$

$C^k$ on a neighbourhood of $\Gamma_b$ and equal to the true gap $g^*$ on the interior regime.

The object is not exotic; sequence-space codes already compute it. The two-asset household routine of Auclert et al. [2021] solves the unconstrained liquid-saving problem first, which is exactly the virtual policy, then the constrained branch, and glues the two by the sign of the virtual violation: its binding region is exactly $\{\tilde{g} \le 0\}$ (replication package `ecta200326-sup-0002`, `models/hank_2a.py`, `backward_iterate`, the mask at lines 88–90).

**Assumption 3** (Boundary regularity)**.** For each bound $b$, on a neighbourhood of its boundary $\Gamma_b$:

1. *(non-degenerate virtual saving)* clause 2 of Assumption 1, read on the virtual policy: $|\det D_a \tilde{a}^*| \ge \epsilon > 0$. On the interior regime this *is* clause 2, every coordinate being free there; the virtual policy carries the bound across $\Gamma_b$ by continuity, and the clause adds only the uniform neighbourhood.

2. *(shadow-price monotonicity)* for $s$ near 0, pin the constrained coordinate at the shifted bound $\underline{a}^i + s$, solve the remaining rows as Step 2 of the proof of Lemma 7 makes precise, and write $R(s; \theta, a)$ for the resulting Euler-row residual, the multiplier of the problem whose bound sits at $\underline{a}^i + s$. The *marginal shadow price of tightening*,

   $$\lambda_b(\theta, a) \;\equiv\; \partial_s R(0; \theta, a),$$

   satisfies $\lambda_b \ge \underline{\lambda} > 0$: tightening a binding bound raises its shadow price at first order. (This household-level $\lambda_b$ is unrelated to the propagation rate $\lambda$ of Assumption 2.) In concave consumption–saving models this is the standard envelope fact, driven by $u'' < 0$: forcing an extra unit of saving on a constrained household moves its Euler residual at first order. The abstract framework needs no concavity, and unlike clauses 1 and 3 this is not an independent hypothesis: $\lambda_b > 0$ follows from the standing strong regularity, derived in Step 4 of Lemma 7 as a nonzero Schur complement of the binding-branch Jacobian. I keep it here for the definition of $\lambda_b$ that the lemma consumes.

3. *(trace of the interior density)* the interior density $\rho^*$ of Assumption 1 admits a trace on $\Gamma_b$ in $L^{p'}(\Gamma_b)$, and is moreover continuous across $\Gamma_b$, its two one-sided traces agreeing, so the single trace $\rho^*|_{\Gamma_b}$ is well-defined. Only the aggregate readings below consume this clause, which is why it is stated here and not in Assumption 1: the first order never pairs anything on $\Gamma$. Continuity across $\Gamma_b$ is what makes those single-trace readings unambiguous; it is stronger than the bare one-sided $L^{p'}$ trace, which a jumping density would still satisfy, yet weaker than the global continuity Bhandari et al. [2023, Assumption 2(b)] assume of the interior density at $d_a = 1$, and the scalar aggregate display already presumed it when it evaluated $\rho^*(\theta, a_\Gamma(\theta))$ at a point.

**Lemma 7** (The binding boundary is a $C^k$ hypersurface)**.** *Let the steady-state household problem be strongly regular (Proposition 2), let clause 1 of Assumption 3 hold, and let $\lambda_b$ be the shadow price its clause 2 defines, $k$ being the per-regime smoothness of the primitives. Then, near each point of $\Gamma_b$:*

1. (multiplier-violation identity) *the binding regime's multiplier extends to a $C^k$ function $\tilde{\mu}$ past the frontier, and it prices the virtual violation,*

   $$\tilde{\mu} \;=\; -\lambda_b\, \tilde{g} \;+\; O\big(\tilde{g}^2\big), \qquad \nabla_a \tilde{\mu} \;=\; -\lambda_b\, \nabla_a \tilde{g} \quad \text{on } \Gamma_b;$$

   *in particular $\tilde{g} < 0$ on the binding regime, where $\tilde{\mu} = \mu^* > 0$, and the shadow price is strictly positive, $\lambda_b \ge \underline{\lambda} > 0$;*

2. (regular level set) *the virtual gap has a non-vanishing state gradient, $|\nabla_a \tilde{g}| \ge c \equiv \epsilon / L^{d_a - 1} > 0$ with $L$ the Lipschitz bound of the virtual policy, and near the point*

   $$\Gamma_b \;=\; \{\tilde{g} = 0\} :$$

*per Markov row $\theta$, $\Gamma_b$ is a regular level set of the $C^k$ function $\tilde{g}$, hence a $C^k$ hypersurface of codimension one, $C^{1,1}$ in particular, with unit normal*

$$\mathbf{n} \;=\; \frac{\nabla_a \tilde{g}}{|\nabla_a \tilde{g}|},$$

*pointing out of the binding regime and into the interior regime, surface measure $d\Gamma$, and bounded tangential divergence $\operatorname{div}_\Gamma \mathbf{n}$ of the normal;*

3. (transversality) *along $e = \mathbf{n}$ the one-sided slopes of the complementarity pair are bounded away from zero, $D_e\, g^* \geq c$ from the interior side and $D_e\, \mu^* \leq -\underline{\lambda} c$ from the binding side: neither the gap nor the multiplier can die out tangentially flat at the frontier, which is the general-$d_a$ form of the knife-edge exclusion of Bhandari et al. [2023, App. B.1].*

*At $d_a = 1$ the slice of $\Gamma_b$ in a row $\theta$ is a finite set of crossing points; at each, $\mathbf{n} = \pm 1$, and the running example's borrowing-constraint picture, binding at low $a$, has $\mathbf{n} = +1$ with the interior regime on the side $a > a_\Gamma(\theta)$.*

*Proof. Step 1: the two one-sided systems extend past the frontier.* Off $\Gamma_b$ each regime solves a smooth square system and the implicit function theorem applies regime by regime (proof of Proposition 2). At a point $p \in \Gamma_b$ strict complementarity fails, but each *branch* system remains smooth and square: the interior branch drops the multiplier and the slackness row, the binding branch pins the coordinate at $\underline{a}^i$ and carries the multiplier as an unknown. Strong regularity at a point where strict complementarity fails forces every branch Jacobian to be nonsingular there [Facchinei and Pang, 2003, Cor. 9.1.24], so the implicit function theorem solves each branch on a *full* ball around $p$: the interior branch returns the virtual policy $\tilde{a}^*$ of Definition 5, the binding branch a $C^k$ continuation $\tilde{\mu}$ of the multiplier past the frontier, each agreeing with the true solution on its own regime.

*Step 2: the family of shifted problems, and its uniqueness ball.* For $s$ near 0, pin the constrained coordinate at $\underline{a}^i + s$ and keep the remaining rows: every member of this family has the binding branch's structure, the pin row differentiating to a unit for every $s$ and $s$ entering only the residual, so the family's Jacobian in the free unknowns is the binding branch's own sub-block, nonsingular at $p$ by Step 1 and uniformly invertible on a neighbourhood by continuity. The quantitative implicit function theorem then yields radii such that, for every small $|s|$ and every state $a$ near $p$, the shifted system has *exactly one* solution inside a fixed ball around the frontier solution, $C^k$ in $(s, a)$ jointly; I take the family to be that in-ball solution, and $R(s; \theta, a)$ its Euler-row residual, as in Assumption 3. No claim is made that this Jacobian equals the interior branch's: the two branch systems differ, the jump between their solution maps' derivatives across $\Gamma_b$ is real, and it is exactly the leading jump $[\partial_{\mathbf{n}}\, \bar{a}^*]$ the interface calculus below runs on.

*Step 3: three readings of $R$.* At $s = 0$ the family is the binding branch itself, so $R(0; a) = \tilde{\mu}(a)$. At $s = \tilde{g}(a)$ the interior branch's solution satisfies the shifted system, the pin then demanding exactly what the household chooses on its own, and it lies in the uniqueness ball by continuity, so it *is* the family member there and its Euler residual vanishes: $R(\tilde{g}(a); a) = 0$. And $\partial_s R(0; a) = \lambda_b(a)$ is the definition of the shadow price. On the binding regime the relevant range is $s \in [\tilde{g}(a), 0]$ with $\tilde{g}(a) < 0$, the pin sliding below the bound, a mathematical member of the family and nothing more.

*Step 4: mean value in the scalar $s$.* The pin is one scalar whatever $d_a$ is, the state $a$ a spectator parameter, so the one-variable mean value theorem applies to $s \mapsto R(s; a)$:

$$\tilde{\mu}(a) \;=\; R(0; a) - R\big(\tilde{g}(a); a\big) \;=\; -\,\partial_s R(\xi_a; a)\; \tilde{g}(a), \qquad \xi_a \text{ between } \tilde{g}(a) \text{ and } 0,$$

and $\partial_s R(\xi_a; a) \to \lambda_b(a)$ as $a$ approaches $\Gamma_b$, which is item 1. For the gradient form, both $\tilde{\mu}$ and $\tilde{g}$ are $C^k$ and vanish on $\Gamma_b$, so their gradients there are purely normal, a $C^1$ function vanishing on a hypersurface having its gradient along the normal, and dividing the identity by the normal distance gives $\partial_{\mathbf{n}}\, \tilde{\mu} = -\,\lambda_b\; \partial_{\mathbf{n}}\, \tilde{g}$ pointwise on $\Gamma_b$, hence $\nabla_a \tilde{\mu} = -\,\lambda_b\, \nabla_a \tilde{g}$ there. On the binding regime $\tilde{g} < 0$, the virtual policy violating the bound being what makes the household constrained (Definitions 1 and 5); on the interior regime $\tilde{g} = g^* > 0$; so near $\Gamma_b$ the zero set $\{\tilde{g} = 0\}$ is exactly the common frontier. The same readings fix the sign of the shadow price, so clause 2 of Assumption 3 need not assume it. Since $R(\tilde{g}; a) = 0$ and $R(0; a) = \mu^*(a) > 0$ on the binding regime (Step 3), the mean-value slope is positive, $\partial_s R(\xi_a; a) = \mu^*(a)/|\tilde{g}(a)| > 0$, so $\lambda_b = \lim_{a \to \Gamma_b} \partial_s R(0; a) \geq 0$. It is moreover nonzero: $\lambda_b = \partial_s R(0)$ is the Schur complement of the free block, nonsingular by Step 2, in the Jacobian of the pinned binding-branch system, which strong regularity makes nonsingular (Step 1, Facchinei and Pang [2003, Cor. 9.1.24]), and

a Schur complement of a nonsingular matrix along a nonsingular block cannot vanish. Hence $\lambda_b > 0$, uniformly $\geq \underline{\lambda} > 0$ on compacta by continuity, which is clause 2, now derived rather than assumed.

*Step 5: the gradient bound, by Hadamard's inequality.* Clause 1 bounds the determinant below, $|\det D_a \tilde{a}^*| \geq \epsilon$, while every row of $D_a \tilde{a}^*$ is bounded by the Lipschitz constant $L$; Hadamard's determinant inequality, $|\det M| \leq \prod_j \|\mathrm{row}_j M\|$, then traps the constrained coordinate's own row from below,

$$\left|\nabla_a \tilde{a}^{*\,i}\right| \;\geq\; \frac{\epsilon}{L^{\,d_a - 1}} \;=\; c \;>\; 0,$$

and $\nabla_a \tilde{g} = \nabla_a \tilde{a}^{*\,i}$. The classical implicit function theorem applied to $\tilde{g}$ at a regular level set writes $\Gamma_b$, near $p$, as the graph of a $C^k$ function transverse to $\mathbf{n} = \nabla_a \tilde{g} / |\nabla_a \tilde{g}|$; the normal points where $\tilde{g}$ grows, out of the binding regime ($\tilde{g} < 0$) and into the interior one, $\mathrm{div}_\Gamma \mathbf{n}$ is $C^{k-2}$, bounded on compacta, and the one-sided slopes of item 3 read $D_{\mathbf{n}}\, g^* = |\nabla_a \tilde{g}| \geq c$ from the interior side and, through item 1, $D_{\mathbf{n}}\, \mu^* = -\,\lambda_b\, D_{\mathbf{n}}\, \tilde{g} \leq -\,\underline{\lambda}\, c$ from the binding side.

*Reduction at $d_a = 1$.* The ball of Step 1 is an interval, $\nabla_a \tilde{g}$ a nonzero scalar slope, and the level set $\{\tilde{g} = 0\}$ a single crossing point per chart: the slice of $\Gamma_b$ is a finite set of points $a_\Gamma(\theta)$, and in the borrowing-constraint picture, binding at low $a$, the virtual gap rises through zero along $+a$, so $\mathbf{n} = +1$ and the interior regime sits at $a > a_\Gamma(\theta)$, the sign convention every scalar reduction below is checked against. □

Two remarks bound the lemma's reach. First, the regularity does not rest on any free-boundary subtlety, and the comparison with the obstacle problem is instructive: there the gap detaches *quadratically* from the free boundary, solution and gradient vanishing together, no one-sided branch with a non-vanishing gradient exists, and smoothness holds only near regular points, singular points being genuinely possible [Figalli, 2018, Sect. 5 and Thm. 7.3]; the positive theory needs blow-up machinery [Petrosyan et al., 2012, Ch. 4, Problem A, and Ch. 7]. Here clauses 1–2 make the complementarity pair transversal, first-order one-sided slopes on both sides, so $\Gamma_b$ is a regular level set of a smooth branch object and the elementary implicit-function route suffices: the obstacle literature is an analogy for what can go wrong without the clauses, not a transfer. Second, at $d_a = 1$ the published antecedent again assumes what the lemma derives: Bhandari et al. [2023, Assumption 2(a)] take the points of non-differentiability to be described by finitely many sufficiently differentiable functions $\kappa_j$, the scalar shadow of item 2.

*Single crossing as monotone virtual saving.* One labeled strengthening of clause 1 recovers a global picture. If the virtual policy is monotone in the bound's own coordinate near $\Gamma_b$, $\partial \tilde{a}^{*\,i} / \partial a^i \geq c > 0$, then along each $a^i$-line the virtual gap is strictly increasing and crosses zero at most once, so the boundary is globally a graph over the remaining coordinates,

$$\Gamma_b \;=\; \left\{\, a \;:\; a^i = \gamma_b\big(\theta,\, a^{-i}\big) \,\right\}.$$

At $d_a = 1$ this is single crossing, one point $a_\Gamma(\theta)$ per row. In the two-asset model it reads: the virtual liquid saving $\tilde{b}'$ is what the household would save ignoring its liquid bound; monotonicity in own liquid wealth, $\partial \tilde{b}' / \partial b > 0$, makes the liquidity frontier the curve $b = b^*(a)$ of Figure 5, with slope

$$b^{*\prime}(a) \;=\; -\,\frac{\partial \tilde{b}' / \partial a}{\partial \tilde{b}' / \partial b}$$

read off virtual-policy derivatives; clause 1 asks that marginal liquid saving not flatten in every state direction at the frontier, clause 2 that forcing an extra unit of liquid saving on a constrained household move its liquid Euler residual at first order.

**Definition 6** (Jump, layer, ledger)**.** Fix the Markov state $\theta$, suppressed throughout, and a bound whose boundary $\Gamma$ is the $C^k$ hypersurface of Lemma 7, now moving along a perturbation with parameters $q \in \mathbb{R}^m$, generic here, in the model the sizes of the announced deviations; the unit normal $\mathbf{n}$ points out of the binding regime, and $g(x^+)$, $g(x^-)$ are the one-sided limits at $x \in \Gamma$ from the side $\mathbf{n}$ points to and from the other. For a function $g$ of the state smooth on each regime:

- the *jump* of $g$ across $\Gamma$ is the function on $\Gamma$ collecting the differences of its one-sided limits, $[g](x) \equiv g(x^+) - g(x^-)$;
- a *layer* of degree $i$ and weight $w$, a function on $\Gamma$, is the distribution $w\, \partial^i_{\mathbf{n}} \delta_\Gamma$, pairing with test functions through the traces of their normal derivatives,

$$\left\langle w\, \partial^i_{\mathbf{n}} \delta_\Gamma,\; \varphi \right\rangle \;=\; (-1)^i \int_\Gamma w\, \left(\partial^i_{\mathbf{n}} \varphi\right)\big|_\Gamma \, d\Gamma,$$

the degree-zero layer $w\,\delta_\Gamma$ being the measure $w\,d\Gamma$ on the hypersurface; a weight in $L^{p'}(\Gamma)$ serves a plain trace pairing, and a weight in the fractional class $W^{1/p,\,p'}(\Gamma)$ of Appendix C serves the pairings that consume tangential gradients of traces, the same two classes that clause 1 of Assumption 1 runs on its codimension-one carriers;

- on the algebra of two-regime functions plus finite layer sums, closed under differentiation by Lemma 10 below, every element decomposes as $f = \{f\} + \sum_{i\le I} w_i\,\partial^i_{\mathbf{n}}\delta_\Gamma$, and the *ledger* of $f$ is its singular part,

$$\mathsf{L}_\Gamma f \;\equiv\; f - \{f\} \;=\; \sum_{i\le I} w_i(f)\,\partial^i_{\mathbf{n}}\delta_\Gamma.$$

Two layer families now coexist, and one sentence keeps them apart: the layers $\sigma^*_f$ of Assumption 1 are singular parts of the *distribution*, carried by the sets $\{a : a^i = \underline{a}^i\}$ on the boundary of $\mathcal{A}$ and present from the first order on; the layers above are singular parts of *policy derivatives*, carried by the interior hypersurface $\Gamma$ and present from the second order on. The same trace machinery pairs both (the trace remark of Appendix C); the carriers and the owners differ. Derivatives of regime-smooth functions produce jumps at the first order in the state and layers from the second on, with weights the interface calculus of Lemma 10 computes. At $d_a = 1$ the slice of $\Gamma$ is the crossing point $a_\Gamma(\theta; q)$, the surface integral is an evaluation, and a layer of degree $i$ collapses to the functional pairing by $(-1)^i\,w\,\varphi^{(i)}(a_\Gamma)$, the derivatives sitting on the test function; in the Dirac notation of Bhandari et al. [2023] this functional is the $i$-th derivative $\delta^{(i)}$ of the scalar Dirac, extending to general degree a notation whose printed instances stop at the first derivative, and it appears below only where their displays are being matched.

**Lemma 8** (The policy recursion at order $k$)**.** *In the setting of Proposition 1, let the heterogeneous residual $\mathcal{F}$ be $k$ times continuously Fréchet differentiable on each regime at the steady state. Then the $k$-th derivative of the backward operator solves*

$$J\,D^k\mathfrak{x}\,(h_1,\dots,h_k) \;=\; -\sum_{\substack{\pi\vdash\{1,\dots,k\}\\ |\pi|\ge 2}} D^{|\pi|}\mathcal{F}\Big(D^{|B_1|}(\mathfrak{x},\mathrm{id})\,(h_{B_1}),\;\dots,\;D^{|B_{|\pi|}|}(\mathfrak{x},\mathrm{id})\,(h_{B_{|\pi|}})\Big),$$

*where the notation reads as follows. A* partition $\pi = \{B_1,\dots,B_{|\pi|}\}$ *of the index set $\{1,\dots,k\}$ is a collection of disjoint nonempty* blocks $B_i \subseteq \{1,\dots,k\}$ *whose union is the whole set, $|\pi|$ counting the blocks; the symbol $\pi \vdash \{1,\dots,k\}$ declares that $\pi$ is a partition of $\{1,\dots,k\}$, and the sum runs over those with $|\pi| \ge 2$ blocks. For a block $B$, $h_B$ collects the directions with labels in $B$, and the inner derivative feeding one slot of the $|\pi|$-linear form $D^{|\pi|}\mathcal{F}$ is*

$$D^{|B|}(\mathfrak{x},\mathrm{id})\,(h_B) \;=\; \begin{cases} \big(D\mathfrak{x}\,h_B,\;h_B\big) & \text{if } |B| = 1,\\ \big(D^{|B|}\mathfrak{x}\,(h_B),\;0\big) & \text{if } |B| \ge 2, \end{cases}$$

*its first component the derivative of the backward operator in the block's directions, its second the derivative of the identity map on $(\bar{x}^+, \bar{Y})$, which is the direction itself for one derivative and zero beyond, since $D\,\mathrm{id} = \mathrm{id}$ at every base point. Every term on the right involves derivatives of $\mathfrak{x}$ of order strictly below $k$, and the operator on the left is the $J$ that Proposition 1 inverted: one factorization serves every order. On each regime the identity is classical; across the binding boundary $\Gamma$ the state-direction derivatives on the right carry the layers of Lemma 10, with the weights computed there.*

*Proof. Step 0: derivatives of maps between spaces of functions, and multilinear forms.* The residual $\mathcal{F}$ and the backward operator $\mathfrak{x}$ act on functions, so their derivatives are the Fréchet derivatives of Appendix C: $Dg(y)\,h$ is the best linear approximation to $g(y+h) - g(y)$, a linear map applied to the direction $h$. Differentiating once more produces an object taking *two* directions: differentiate $y \mapsto Dg(y)\,h_1$ along $h_2$ and call the result $D^2g(y)(h_1,h_2)$. It is linear in $h_1$ and in $h_2$ separately, and symmetric in the pair, for the same reason mixed partials commute; the finite-dimensional anchor is the Hessian, $D^2g(h_1,h_2) = h_1^\top(\nabla^2 g)\,h_2$, and "bilinear form" names exactly that object with function-valued directions. Iterating $k$ times gives the $k$-linear symmetric form $D^kg(h_1,\dots,h_k)$, the multilinear form of the statement: a machine that eats $k$ directions and returns a residual-sized object, linearly in each.

*Step 1: the second derivative of a composition, by hand.* Before any general formula, differentiate $G(y) = g(f(y))$ twice. Once, the chain rule:

$$DG(y)\,h_1 \;=\; Dg\big(f(y)\big)\,\big(Df(y)\,h_1\big).$$

Differentiate this along $h_2$. Two factors depend on $y$: the outer derivative $Dg(f(y))$, whose differentiation produces the bilinear form of $g$ fed with the transported direction $Df\,h_2$; and the inner factor $Df(y)\,h_1$, whose differentiation produces the bilinear form of $f$. The product rule collects both,

$$D^2G\,(h_1,h_2) \;=\; D^2g\big(Df\,h_1,\; Df\,h_2\big) \;+\; Dg\;D^2f\,(h_1,h_2)\,:$$

the curvature of the outer map against two transported directions, plus the slope of the outer map against the curvature of the inner one. Differentiating once more would produce, by the same two rules, the third derivative of $g$ against three transported directions, three mixed terms in which one pair of directions has passed through $D^2f$, and the slope of $g$ against $D^3f$.

*Step 2: the pattern is indexed by partitions.* The terms just produced are enumerated by a combinatorial object. A *partition* $\pi$ of the direction labels $\{1,\dots,k\}$ is a way of grouping them into disjoint nonempty *blocks* $B_1,\dots,B_{|\pi|}$ covering all labels, $|\pi|$ counting the blocks. The set $\{1,2\}$ has exactly two partitions, and they are exactly the two terms of Step 1:

$$\begin{aligned}\pi=\big\{\{1\},\{2\}\big),\;\; B_1=\{1\},\;\; B_2=\{2\},\;\; |\pi|=2:\qquad & D^2g\big(Df\,h_1,\; Df\,h_2\big),\\ \pi=\big\{\{1,2\}\big),\;\; B_1=\{1,2\},\;\; |\pi|=1:\qquad & Dg\;D^2f\,(h_1,h_2),\end{aligned}$$

one slot of the outer derivative per block, each block passing through an inner derivative of its own size. The set $\{1,2,3\}$ has five partitions, and they are the five terms of the third derivative:

$$\begin{aligned}\big\{\{1\},\{2\},\{3\}\big),\;\; |\pi|=3:\qquad & D^3g\big(Df\,h_1,\; Df\,h_2,\; Df\,h_3\big),\\ \big\{\{1,2\},\{3\}\big),\;\; |\pi|=2:\qquad & D^2g\big(D^2f(h_1,h_2),\; Df\,h_3\big),\\ \big\{\{1,3\},\{2\}\big),\;\; |\pi|=2:\qquad & D^2g\big(D^2f(h_1,h_3),\; Df\,h_2\big),\\ \big\{\{2,3\},\{1\}\big),\;\; |\pi|=2:\qquad & D^2g\big(D^2f(h_2,h_3),\; Df\,h_1\big),\\ \big\{\{1,2,3\}\big),\;\; |\pi|=1:\qquad & Dg\;D^3f\,(h_1,h_2,h_3),\end{aligned}$$

the outer trilinear form fed by three transported directions, the three mixed terms, one per way of pairing two directions inside the inner map, and the inner third derivative under the outer slope. The rule visible in both cases is general: each block of directions enters through one derivative of the inner map, of order the block's size, and the blocks fill the slots of a derivative of the outer map, of order the number of blocks,

$$D^k(g\circ f)\,(h_1,\dots,h_k) \;=\; \sum_{\pi\vdash\{1,\dots,k\}} D^{|\pi|}g\Big(D^{|B_1|}f\,(h_{B_1}),\;\dots,\;D^{|B_{|\pi|}|}f\,(h_{B_{|\pi|}})\Big),$$

where $h_B$ collects the directions with labels in $B$; the Banach-space statement and proof are Fraenkel [1978, p. 161], and the induction is the one Step 1 began.

*Step 3: apply the formula to the defining identity.* Take $g=\mathcal{F}$ and, as the inner map, $f(y)=(\mathfrak{x}(y),\,y)$ with $y=(\bar{x}^+,\bar{Y})$, so that $g\circ f$ is the map $y\mapsto\mathcal{F}\big(\mathfrak{x}(y),\bar{x}^+,\bar{Y}\big)$, identically zero on a neighbourhood of the steady state by Proposition 1; all its derivatives vanish. Two simplifications shape the partition sum. First, the identity component of $f$ has derivative $h\mapsto h$ at every base point; since this linear map does not vary with the base point, its second and higher derivatives are zero, so blocks of size two or more feed *only* the $\mathfrak{x}$ component. Second, run the $k=2$ case in full before the general one: with $Df\,h=(D\mathfrak{x}\,h,\;h)$ and $D^2f(h_1,h_2)=(D^2\mathfrak{x}(h_1,h_2),\;0)$, Step 1's two terms read

$$0 \;=\; D^2\mathcal{F}\Big(\big(D\mathfrak{x}\,h_1,\;h_1\big),\;\big(D\mathfrak{x}\,h_2,\;h_2\big)\Big) \;+\; D\mathcal{F}\cdot\big(D^2\mathfrak{x}(h_1,h_2),\;0\big),$$

and the last term, the slope of $\mathcal{F}$ applied to a direction moving only the first argument group, is by definition $D_{\bar{x}}\mathcal{F}\,D^2\mathfrak{x}(h_1,h_2)=J\,D^2\mathfrak{x}(h_1,h_2)$: rearranged, the $k=2$ display of the statement. At general $k$ the same two observations apply verbatim: the only partition whose block reaches size $k$ is the single-block one, contributing $J\,D^k\mathfrak{x}$, and moving every other partition to the right-hand side gives the display, every block there being of size below $k$ and hence carrying an already-known derivative of $\mathfrak{x}$.

*Step 4: regimes and layers.* Off the binding boundary $\Gamma$, strict complementarity fixes each agent's regime on a neighbourhood (Proposition 2), so on each regime $\mathcal{F}$ is $k$ times differentiable by hypothesis and the implicit function theorem delivers $D^k\mathfrak{x}$ classically, regime by regime, with the display holding pointwise. The right-hand side, however, evaluates state derivatives of objects that jump across the binding boundary $\Gamma$: the conditional-expectation argument composes next-period policies with the

landing point $\bar{a}(\theta, a)$, and differentiating that composition twice or more in the state differentiates the jumps that Proposition 2 left bounded, the slope $D_a\bar{x}^*$ of the steady-state policy and the coefficients $\times_s$ themselves. Those state derivatives exist as distributions with layers on the binding boundary $\Gamma$, computed by Lemma 10; the recursion is then read as a per-regime classical identity plus its ledger $\mathsf{L}_\Gamma$, both constructive.

*Where the derivatives of $J$ went.* One may wonder why the final formula never differentiates $J$, although $J(y) = D_{\bar{x}}\mathcal{F}(\mathfrak{x}(y), y)$ certainly varies along the path. It does, and the formula has not lost it; the point is worth one explicit reconciliation at $k = 2$. The first-order identity is $J(y)\, D\mathfrak{x}(y)\, h_1 + D\mathcal{F}\big(\mathfrak{x}(y), y\big) \cdot (0, h_1) = 0$, the second term the slope of $\mathcal{F}$ against the direction moving the $(\bar{x}^+, \bar{Y})$ group alone; differentiate it directly along $h_2$, with the product rule:

$$0 \;=\; \big(D_{h_2}J\big)\, D\mathfrak{x}\, h_1 \;+\; J\, D^2\mathfrak{x}(h_1, h_2) \;+\; D^2\mathcal{F}\Big(\big(0,\; h_1\big),\; \big(D\mathfrak{x}\, h_2,\; h_2\big)\Big),$$

and compute the first term: $J(y)$ is itself a composition, so its derivative along $h_2$ is the bilinear form of $\mathcal{F}$ with one slot held at the $\bar{x}$-direction it acts on,

$$\big(D_{h_2}J\big)\, v \;=\; D^2\mathcal{F}\Big(\big(v,\; 0\big),\; \big(D\mathfrak{x}\, h_2,\; h_2\big)\Big),$$

the $\bar{x}$-slice of $D^2\mathcal{F}$ against the transported second direction, while the last term of the display carries the remaining slices. At $v = D\mathfrak{x}\, h_1$, linearity of the bilinear form in its first slot sums the two pieces,

$$D^2\mathcal{F}\Big(\big(D\mathfrak{x}\, h_1,\; 0\big),\; \big(D\mathfrak{x}\, h_2,\; h_2\big)\Big) \;+\; D^2\mathcal{F}\Big(\big(0,\; h_1\big),\; \big(D\mathfrak{x}\, h_2,\; h_2\big)\Big) \;=\; D^2\mathcal{F}\Big(\big(D\mathfrak{x}\, h_1,\; h_1\big),\; \big(D\mathfrak{x}\, h_2,\; h_2\big)\Big),$$

reassembling exactly $D^2\mathcal{F}\big((D\mathfrak{x}\, h_1, h_1), (D\mathfrak{x}\, h_2, h_2)\big) + J\, D^2\mathfrak{x}(h_1, h_2) = 0$, the $k = 2$ display of Step 3: the variation of $J$ along the path is not dropped, it is the two-block partition terms, and the $J$ of the final formula is frozen at the steady state because every object in it is evaluated at the base point. □

The horizon reading mirrors Lemma 1, and it is worth deriving rather than asserting, because it shows where each piece of the recursion comes from. Fix the date-$t$ identity along a perturbed path,

$$\mathcal{F}\Big(\bar{x}_t,\; \mathbb{E}\big[\,\bar{x}_{t+1}\big(\theta', \bar{a}_t(\theta, a)\big)\,\big|\,\theta\big],\; \bar{Y}_t\Big) \;=\; 0,$$

and rerun the announcement experiment of Section 4.2 that isolated the coefficients $\times_s$: households learn today that their exogenous input, the path of relevant aggregates $\bar{Y}$, will deviate at a single date, $s$ periods ahead, and at no other date. The one-date deviation is the experimenter's design, not a claim about equilibrium aggregates, whose full path response is rebuilt afterwards by superposing these experiments. The first argument moves by today's response, $\times_s$ per unit of perturbation (Lemma 1). The middle argument is a composition and its derivative has two terms, one per moving piece:

$$d\,\mathbb{E}\big[\,\bar{x}_{t+1}\big(\theta', \bar{a}_t(\theta, a)\big)\,\big|\,\theta\big] \;=\; \mathbb{E}\Big[\;\underbrace{\times_{s-1}\big(\theta', \bar{a}^*(\theta, a)\big)}_{\text{next period's policy moves}} \;+\; \underbrace{D_a\bar{x}^*\big(\theta', \bar{a}^*(\theta, a)\big)\; p \times_s (\theta, a)}_{\text{the landing point moves}} \;\Big|\; \theta\Big],$$

next period seeing the shock $s-1$ periods ahead, and the landing point displaced by today's predetermined response $p \times_s$. The third argument $\bar{Y}_t$ moves directly only when the announced deviation lands, a loading I write $e_s$, nonzero only at $s = 0$: the anticipation the announcement stirs before landing travels entirely through the first two arguments, while the date-$t$ residual's own aggregate argument sits at the steady state until the deviation date arrives. Collecting the three movements into one composite first-order slice,

$$u_s \;\equiv\; \Big(\times_s,\;\; \mathbb{E}\big[\times_{s-1} + D_a\bar{x}^*\; p \times_s \,\big|\, \theta, a\big],\;\; e_s\Big),$$

one object per horizon, feeding the three argument groups of $\mathcal{F}$. Now differentiate a second time, along a second horizon $s_2$. Differentiating $\mathcal{F}$'s slopes produces its bilinear form against the two slices, $D^2\mathcal{F}(u_{s_1}, u_{s_2})$; differentiating the *inside* of the middle argument, the product and chain rules applied to the display above, produces, writing $\times_{s_1 s_2}{}^{(2)}$ for the order-two policy kernel at the pair of horizons and evaluating every function at the steady-state landing point $\bar{a}^*(\theta, a)$,

$$\begin{aligned} d^2\,\mathbb{E}\big[\,\bar{x}_{t+1}\big(\theta', \bar{a}_t(\theta, a)\big)\,\big|\,\theta\big] \;=\; \mathbb{E}\Big[\; & \times_{s_1-1,\, s_2-1}{}^{(2)} \;+\; D_a\times_{s_1-1}\big(p \times_{s_2}\big) \;+\; D_a\times_{s_2-1}\big(p \times_{s_1}\big) \\ & +\; D_a^2\bar{x}^*\;\big(p \times_{s_1}\big)\big(p \times_{s_2}\big) \;+\; D_a\bar{x}^*\; p \times_{s_1 s_2}{}^{(2)} \;\Big|\; \theta\Big]. \end{aligned}$$

Each term is one moving piece caught by the other perturbation: next period's policy responds at second order; the next-period coefficient $\times_{s_1-1}$ is read at a landing point the second perturbation displaces, and symmetrically with the horizons exchanged; the slope $D_a\bar{x}^*$ of the landing-point displacement is itself read at a displaced point, once for the pair; and the landing point responds at second order. The coefficient of that last term once the outer slope is applied, $\mathcal{F}_{x^+}\,\mathbb{E}[D_a\bar{x}^* \mid \theta, a]\,p$, is precisely the second summand of $J$ as Proposition 2 displays it, so the term joins the left-hand side inside $J$. What remains on the right is the backward recursion

$$J\times_{s_1 s_2}{}^{(2)} \;=\; -\,D^2\mathcal{F}\big(u_{s_1}, u_{s_2}\big) \;-\; \mathcal{F}_{x^+}\,\mathbb{E}\Big[\, D_a\times_{s_1-1}\big(p\times_{s_2}\big) \;+\; D_a\times_{s_2-1}\big(p\times_{s_1}\big) \\ +\; D_a^2\bar{x}^*\,\big(p\times_{s_1}\big)\big(p\times_{s_2}\big) \;+\; \times_{s_1-1,\,s_2-1}{}^{(2)} \,\Big|\, \theta, a\Big],$$

solved backward in the pair of horizons from zero at large horizons, one solve in the factored $J$ per node and pair. In the bracket, the state derivatives $D_a\times_s$ of the first-order coefficients and the second state derivative $D_a^2\bar{x}^*$ of the steady-state policy are the objects that jump across the binding boundary $\Gamma$; their ledgers, $\mathsf{L}_\Gamma(D_a^2\,\bar{x}^*) = [\partial_{\mathbf{n}}\,\bar{x}^*]\,(\mathbf{n}\otimes\mathbf{n})\,\delta_\Gamma$ for the latter, are computed by Lemma 10. At order three the same two moves repeat. The middle argument differentiates a third time by the identical landing-point chain rule, every function evaluated at the steady-state landing point $\bar{a}^*(\theta, a)$ and each sum running over the three splits of $\{1,2,3\}$ into a pair and a singleton,

$$d^3\,\mathbb{E}\big[\,\bar{x}_{t+1}\big(\theta', \bar{a}_t(\theta,a)\big)\,\big|\,\theta\big] \;=\; \mathbb{E}\Big[\; \times_{s_1-1,\,s_2-1,\,s_3-1}{}^{(3)} \;+\; \sum_{\{i,j\},\,k} D_a\times_{s_i-1,\,s_j-1}{}^{(2)}\,\big(p\times_{s_k}\big) \;+\; \sum_{\{j,k\},\,i} D_a\times_{s_i-1}\big(p\times_{s_j s_k}{}^{(2)}\big) \\ +\; \sum_{\{j,k\},\,i} D_a^2\times_{s_i-1}\big(p\times_{s_j}\big)\big(p\times_{s_k}\big) \;+\; D_a^2\bar{x}^*\sum_{\{i,j\},\,k}\big(p\times_{s_i s_j}{}^{(2)}\big)\big(p\times_{s_k}\big) \\ +\; D_a^3\bar{x}^*\,\big(p\times_{s_1}\big)\big(p\times_{s_2}\big)\big(p\times_{s_3}\big) \;+\; D_a\bar{x}^*\;p\times_{s_1 s_2 s_3}{}^{(3)} \;\Big|\; \theta\Big],$$

fifteen terms in seven families. The last family joins $J$ exactly as at order two; write $R_{s_1 s_2 s_3}$ for the sum of the first six, the order-three source of the middle argument. The partitions of the three directions then sort the outer derivatives into the trilinear form and three bilinear-times-linear terms, the order-two composite slice $w_{s_i s_j} \equiv \big(\times_{s_i s_j}{}^{(2)},\; d^2\,\mathbb{E}\big[\bar{x}_{t+1}(\theta', \bar{a}_t) \mid \theta\big],\; 0\big)$ collecting the order-two response of the three argument groups,

$$J\times_{s_1 s_2 s_3}{}^{(3)} \;=\; -\,D^3\mathcal{F}\big(u_{s_1}, u_{s_2}, u_{s_3}\big) \;-\; \sum_{\{i,j\},\,k} D^2\mathcal{F}\big(w_{s_i s_j},\; u_{s_k}\big) \;-\; \mathcal{F}_{x^+}\,\mathbb{E}\big[\,R_{s_1 s_2 s_3}\,\big|\,\theta, a\big].$$

The middle term is the point: it consumes the order-two object exactly as the order-two solve returned it, contracted once against a first-order slice. Order three needs nothing from order two beyond its output; the only genuinely new inputs are the third derivatives of $\mathcal{F}$ and one more pass through the factored $J$. Inside $R_{s_1 s_2 s_3}$ the state derivatives jump across the binding boundary $\Gamma$ and their ledgers reach one degree deeper than at order two (Lemma 10); displayed along the normal, every slot of the multilinear ledger contracted against $\mathbf{n}$, the tangential slots adding tangential gradients of the same jumps,

$$\mathsf{L}_\Gamma\big(D_a^3\,\bar{x}^*\big)(\mathbf{n},\mathbf{n},\mathbf{n}) \;=\; \Big(\big[\partial_{\mathbf{n}}^2\,\bar{x}^*\big] \;-\; \big(\operatorname{div}_\Gamma \mathbf{n}\big)\big[\partial_{\mathbf{n}}\,\bar{x}^*\big]\Big)\,\delta_\Gamma \;+\; \big[\partial_{\mathbf{n}}\,\bar{x}^*\big]\;\partial_{\mathbf{n}}\delta_\Gamma,$$
$$\mathsf{L}_\Gamma\big(D_a^2\times_s\big)(\mathbf{n},\mathbf{n}) \;=\; \Big(\big[\partial_{\mathbf{n}}\times_s\big] \;-\; \big(\operatorname{div}_\Gamma \mathbf{n}\big)\big[\times_s\big]\Big)\,\delta_\Gamma \;+\; \big[\times_s\big]\;\partial_{\mathbf{n}}\delta_\Gamma,$$

the order-three ledger; at $d_a = 1$ the curvature terms vanish and the two ledgers read against a test function $\varphi$, per Markov row, as

$$\big\langle \mathsf{L}_\Gamma\big(D_a^3\,\bar{x}^*\big),\;\varphi\big\rangle \;=\; \big[D_a^2\,\bar{x}^*\big]\,\varphi(a_\Gamma) \;-\; \big[D_a\,\bar{x}^*\big]\,\varphi'(a_\Gamma), \qquad \big\langle \mathsf{L}_\Gamma\big(D_a^2\times_s\big),\;\varphi\big\rangle \;=\; \big[D_a\times_s\big]\,\varphi(a_\Gamma) \;-\; \big[\times_s\big]\,\varphi'(a_\Gamma),$$

the derivative sitting on the test function; in the generalized-function notation of Bhandari et al. [2023] these are the combinations $[\cdot]\,\delta + [\cdot]\,\delta'$. The general order collects the same two moves into one display.

For a block $B \subseteq \{1,\dots,k\}$ of announcement horizons, the order-$|B|$ composite slice feeding one slot of the outer forms is

$$w_{s_B} \;\equiv\; \Big(\times_{s_B}{}^{(|B|)},\;\; d^{|B|}\,\mathbb{E}\big[\,\bar{x}_{t+1}\big(\theta', \bar{a}_t\big)\,\big|\,\theta\big],\;\; e_{s_B}\Big), \qquad w_{s_{\{i\}}} = u_{s_i}, \qquad e_{s_B} = 0 \;\text{ for } |B| \geq 2,$$

the third component vanishing beyond one derivative because the announced path is linear in the sizes. The middle argument expands by the same two moving pieces as at orders two and three, one term per

choice of the directions $T \subseteq \{1,\dots,k\}$ routed through the landing point and per partition $\pi$ of $T$ into displacement blocks,

$$d^k\,\mathbb{E}\big[\,\bar{x}_{t+1}\big(\theta',\bar{a}_t(\theta,a)\big)\,\big|\,\theta\big] \;=\; \mathbb{E}\Big[\sum_{T\subseteq\{1,\dots,k\}}\;\sum_{\pi\vdash T} D_a^{|\pi|}\times_{s_{T^c}-1}{}^{(k-|T|)}\;\prod_{B\in\pi}\big(p\times_{s_B}{}^{(|B|)}\big)\;\Big|\;\theta\Big],$$

every function at the steady-state landing point $\bar{a}^*(\theta,a)$, with the conventions $\times_{s_{T^c}-1}{}^{(0)}=\bar{x}^*$ at $T^c=\emptyset$, $s_{T^c}-1$ shifting every date in $T^c$ by one, and $D_a^0$ the identity; at $k=2$ the pairs $(T,\pi)$ number five and at $k=3$ fifteen, reproducing the two displays above term for term. Writing $R_{s_1\cdots s_k}$ for this sum less its top term $(T,\pi)=(\{1,\dots,k\},\{T\})$, which is $D_a\bar{x}^*\;p\times_{s_1\cdots s_k}{}^{(k)}$ and joins $J$ through the outer slope $\mathcal{F}_{x^+}$, the partition formula of Lemma 8 reads, horizon by horizon,

$$J\times_{s_1\cdots s_k}{}^{(k)} \;=\; -\sum_{\substack{\pi\vdash\{1,\dots,k\}\\ |\pi|\geq 2}} D^{|\pi|}\mathcal{F}\big(w_{s_{B_1}},\dots,w_{s_{B_{|\pi|}}}\big)\;-\;\mathcal{F}_{x^+}\,\mathbb{E}\big[\,R_{s_1\cdots s_k}\,\big|\,\theta,a\big],$$

solved backward in the horizon tuple from zero at large horizons, one solve in the factored $J$ per node and tuple. Inductively, the order-$k$ right-hand side consumes the outputs of orders $2,\dots,k-1$ and the derivatives of $\mathcal{F}$ up to order $k$, and nothing else; at the representative-agent level this sequential structure is Schmitt-Grohé and Uribe [2004, §3.3]. The arrays the order-two step consumes are already emitted, `dynamic_het1_g2.m` tabulating the second derivatives of the heterogeneous residual; higher orders extend the same differentiation.

**The interface calculus.** Proposition 2 kept the first derivative clean: the coefficients $\times_s$ are bounded, with jumps across the binding boundary $\Gamma$ and no impulse. The higher orders differentiate those jumps, and the derivative of a jump is a layer: this is where the constraint re-enters the expansion, and it does so lawfully. What the recursion needs is not a case-by-case accounting but a calculus that stays closed at every order. The lemma is therefore stated for a generic two-regime function with a moving discontinuity hypersurface, the *interface* $\Gamma(q)$; in the model the interface is a bound's binding boundary, the hypersurface of Lemma 7, its position moved by the perturbation parameters $q$, the sizes of the announced deviations at the horizons under study, with the steady state sitting at $q=0$. With several bounds the calculus runs pairwise: at a generic point of $\Gamma$ exactly two regimes meet, points where three or more meet lying in the intersection of two per-bound boundaries, of codimension at least two and $\bar{\Omega}^*$-null, the lemma's one-sided pieces $f^{\pm}$ belonging one to each. One clause disciplines those meeting points: *the per-bound clauses of Assumption 3 hold with uniform constants up to the closure of each regime, and distinct boundaries intersect transversally with a uniform angle,* $|\mathbf{n}_b\cdot\mathbf{n}_{b'}|\leq 1-c_0$ *on* $\Gamma_b\cap\Gamma_{b'}$. The clause pays for itself: near a transversal intersection four regimes meet, each per-regime extension is $C^k$ up to its closed regime (Lemma 7 run per bound with the uniform constants), so every layer weight is a difference of two such extensions, bounded with its derivatives, the floors on $1/|\nabla_a\tilde{g}|$ and on the transversal slopes survive to the closure, and the aggregate surface readings extend from the $d\Gamma$-full subset away from intersections to all of $\Gamma$ by dominated convergence, the bounded weights integrating against the $L^{p'}$ trace. In the two-asset model the angle condition is the non-tangency of the liquidity frontier to the illiquid face, read off the slope of the single-crossing remark.

Its output is consumed three times below: the ledger displays of the recursions read off which layers a source carries and with what weights; the normal-speed hierarchy $v_\Gamma^{(j)}$ of Corollary 3, $v_\Gamma=-[\partial_{\mathbf{n}}\bar{x}^*]^{-1}[\times_s]$ at first order, is its constraint chain solved for the motion of $\Gamma$ order by order; and the aggregate pairing after the Bhandari et al. [2023] comparison is where each layer finally becomes a number.

One standing fact underlies all three, and the calculus below assumes it: that the moving boundary $\Gamma(q)$ is a $C^k$ hypersurface not just at the steady state $q=0$, where Lemma 7 proves it, but jointly in $q$ along the whole announced path. The date-$t$ problem along a path is infinite-horizon, its continuation depending on all of $q$, so the regularity must be earned uniformly.

**Lemma 9** (Uniform strong regularity along the path)**.** *Let the standing strong-regularity hypothesis of Section 4.2 hold at the steady state, and let the boundary states range over a compact set. Then there is $r>0$ such that the unfolded steady-state system is strongly regular uniformly for every perturbation $|q|<r$: the perturbed virtual gap $\tilde{g}(\cdot\,;q)$ is jointly $C^k$ in the state and $q$, reducing to $\tilde{g}$ of Lemma 7 at $q=0$, and $\Gamma(q)=\{\tilde{g}(\cdot\,;q)=0\}$ is a $C^k$ family of hypersurfaces, with unit normal, surface measure, and* $\operatorname{div}_\Gamma\mathbf{n}$ *depending $C^{k-2}$-ly on $q$. This is the standing hypothesis Lemma 10 makes on each $\Gamma(q)$.*

*Proof.* At $q = 0$ every branch Jacobian is nonsingular, by strong regularity where strict complementarity fails [Facchinei and Pang, 2003, Cor. 9.1.24], as in Step 1 of Lemma 7. Along an announced path of size $q$, the date-$t$ unfolded problem differs from the steady-state one only through the aggregate deviations, which enter through the backward recursion of Lemma 1; clause 1 of Assumption 2 discounts that recursion geometrically, so the map from $q$ to the date-$t$ problem's data is $C^k$ with derivatives bounded by a convergent geometric series, uniformly in $t$, and the perturbation is $O(|q|)$ uniformly. Strong regularity is stable under such a perturbation [Robinson, 1980, Thm. 2.1]: it persists with a Lipschitz solution map on a neighbourhood whose radius depends only on the perturbation size, so compactness of the boundary states yields a single radius $r > 0$ for all of them. The data being $C^k$ in $(\text{state}, q)$ and the branch Jacobian nonsingular, the parametric implicit function theorem [Abraham et al., 1988, Thm. 2.5.7] makes $\tilde{g}(\cdot\,; q)$ jointly $C^k$. Its state gradient stays bounded below, $|\nabla_a \tilde{g}| \geq c$ uniformly on $|q| < r$ by continuity (Step 5 of Lemma 7), so $\Gamma(q)$ is a regular level set for each $q$ and the regular-level-set construction runs uniformly, giving the $C^k$ family and the stated $q$-regularity of its normal, measure, and curvature. □

**Lemma 10** (Interface calculus)**.** *Fix the Markov state, a suppressed spectator, and let $\Gamma(q)$ be the moving interface, for each parameter vector $q$ in an open neighbourhood of the origin of $\mathbb{R}^m$ a $C^k$ hypersurface as in Lemma 7, with unit normal $\mathbf{n}$, surface measure $d\Gamma$, and layers $w\,\partial_{\mathbf{n}}^i \delta_\Gamma$ as in Definition 6. Let $f(a; q)$ be smooth in the state on each regime, up to the interface, with jump $[f]$, a function on $\Gamma(q)$. Write $v_\Gamma(x; q)$ for the* normal speed *of the interface at $x \in \Gamma(q)$, the normal component of the velocity of its points, the only part of that velocity any rule below sees; and, for a weight $w$ on the moving interface, write $\dot{w}$ for its* material derivative*, the derivative of $w$ along the flow carrying each point of $\Gamma(q)$ at the normal speed. Then, as distributions in the state,*

$$
\begin{aligned}
&\text{(JF)} && \nabla f \;=\; \{\nabla f\} \;+\; [f]\;\mathbf{n}\,\delta_\Gamma,\\
&\text{(PT)} && D_q f \;=\; \{D_q f\} \;-\; [f]\;v_\Gamma\,\delta_\Gamma,\\
&\text{(LT}_a\text{)} && D_i\big(w\,\delta_\Gamma\big) \;=\; w\,\mathbf{n}_i\;\partial_{\mathbf{n}}\delta_\Gamma \;+\; \Big[\big(\nabla_\Gamma w\big)_i \;-\; \big(\operatorname{div}_\Gamma \mathbf{n}\big)\,w\,\mathbf{n}_i\Big]\,\delta_\Gamma,\\
&\text{(LT}_q\text{)} && D_q\big(w\,\delta_\Gamma\big) \;=\; \dot{w}\,\delta_\Gamma \;-\; w\,v_\Gamma\;\partial_{\mathbf{n}}\delta_\Gamma \;+\; \big(\operatorname{div}_\Gamma \mathbf{n}\big)\,w\,v_\Gamma\,\delta_\Gamma.
\end{aligned}
$$

*Consequently the algebra of two-regime functions plus finite layer sums is closed under state and parameter derivatives; every mixed derivative of total order $K$ of a function continuous across the interface is a two-regime classical function plus layers of degree at most $K - 2$, with weights built from jumps of one-sided derivatives, the normal speeds $v_\Gamma^{(j)}$ of Corollary 3, tangential gradients of lower weights, and $\operatorname{div}_\Gamma \mathbf{n}$ with its tangential derivatives. If moreover $f$ stays continuous across the interface as $q$ varies, $[f](\cdot\,; q) \equiv 0$, the jumps of its first derivatives are purely normal and chained,*

$$
[\nabla f] \;=\; \big[\partial_{\mathbf{n}} f\big]\;\mathbf{n}, \qquad\qquad [D_q f] \;=\; -\,\big[\partial_{\mathbf{n}} f\big]\;v_\Gamma.
$$

*Proof. Step 0: what an identity between distributions says.* An identity between distributions asserts that both sides return the same number against every smooth, compactly supported test function $\varphi$; nothing more is claimed, and every rule below is proved by computing both pairings. Derivatives move onto the test function, $\langle D_i f, \varphi\rangle = -\langle f, D_i\varphi\rangle$, the weak derivative of Appendix C [Brenner and Scott, 2008, (1.2.4)], and the layers pair through the traces of normal derivatives, as Definition 6 records. The printed source for the calculus in this generality is Kanwal [1998, Ch. 5], pinned rule by rule below; the shape-calculus transport formulas are those of Delfour and Zolésio [2011, Ch. 9].

*Step 1: jumps of gradients of continuous functions are purely normal.* Let $[f] \equiv 0$ and take a local graph chart at a base point of $\Gamma$: rotate the coordinates so the tangent plane at the point is horizontal, write $a = (a', s)$ with $\Gamma = \{s = \gamma(a')\}$ and $\nabla'\gamma = 0$ at the base point. The trace identity $f^+\big(a', \gamma(a')\big) = f^-\big(a', \gamma(a')\big)$ holds identically in $a'$; differentiating along $a'_j$,

$$
\big[D_{a'_j} f\big] \;+\; \big[D_s f\big]\,\gamma_j \;=\; 0, \qquad j = 1, \ldots, d_a - 1,
$$

so the jump of the gradient annihilates every tangent vector $(e_j, \gamma_j)$ and only its normal component survives, $[\nabla f] = [\partial_{\mathbf{n}} f]\,\mathbf{n}$ [Kanwal, 1998, Sect. 5.5, eqs. (12)–(13)]. Differentiating the trace identity twice constrains the Hessian jump as well: at the base point $[D_{a'_j} D_{a'_k} f] = -\,[D_s f]\,\gamma_{jk}$, the tangential-tangential block of the Hessian jump being the first normal jump against the curvature of the interface, not a free quantity [Kanwal, 1998, Sect. 5.6, Thm. 2, eqs. (9)–(13)]. At $d_a = 1$ there is no tangential direction and both statements are empty: this is exactly the content a scalar calculus never sees.

*Step 2: (JF), by the per-regime divergence theorem.* Pair $D_i f$ against $\varphi$ and integrate by parts separately on the two regimes; along $\Gamma$ the binding regime's outer normal is $+\mathbf{n}$ and the interior regime's is $-\mathbf{n}$, and the terms on the boundary of $\mathcal{A}$ drop by the no-flux conventions of Assumption 1:

$$\langle D_i f,\ \varphi\rangle \ =\ -\int f\ D_i\varphi\ da \ =\ \int \{D_i f\}\,\varphi\ da \ +\ \int_\Gamma \left(f^+ - f^-\right)\mathbf{n}_i\ \varphi\big|_\Gamma\ d\Gamma,$$

the pairing of $\{D_i f\} + [f]\,\mathbf{n}_i\,\delta_\Gamma$: a layer is born on the interface with the vector weight $[f]\,\mathbf{n}$ [Kanwal, 1998, Sect. 5.5, Thm. 1, eq. (4)].

*Step 3: ($\mathrm{LT}_a$), by the tangential Green formula.* For a layer, $\langle D_i(w\,\delta_\Gamma),\varphi\rangle = -\int_\Gamma w\,(D_i\varphi)|_\Gamma\,d\Gamma$, and the derivative of the test function splits along the interface into its normal and tangential parts, $(D_i\varphi)|_\Gamma = \mathbf{n}_i\,(\partial_{\mathbf{n}}\varphi)|_\Gamma + (\nabla_\Gamma(\varphi|_\Gamma))_i$. The normal part is the degree-one pairing, $w\,\mathbf{n}_i\,\partial_{\mathbf{n}}\delta_\Gamma$. The tangential part integrates by parts *on* $\Gamma$, by the tangential Green formula [Delfour and Zolésio, 2011, Ch. 9, Sect. 5.5, eq. (5.27)], whose curvature term is exactly the divergence of the normal field: it returns $(\nabla_\Gamma w)_i\,\delta_\Gamma$ and the correction $-(\mathrm{div}_\Gamma\,\mathbf{n})\,w\,\mathbf{n}_i\,\delta_\Gamma$. Contracting the rule with $\mathbf{n}$ reproduces the printed one-line form $\partial_{\mathbf{n}}$-of-a-Dirac-layer, curvature term included [Kanwal, 1998, Sect. 5.3, eq. (9), with his mean-curvature convention translated into $\mathrm{div}_\Gamma\,\mathbf{n}$]. A sanity check to keep: on a circle of radius $R$ in the plane, $\mathrm{div}_\Gamma\,\mathbf{n} = 1/R$ when $\mathbf{n}$ points out of the disc, and the rule is the classical $\nabla\delta_S = \mathbf{n}\,\partial_{\mathbf{n}}\delta_S - (\mathrm{div}_\Gamma\,\mathbf{n})\,\mathbf{n}\,\delta_S$ picture.

*Step 4: (PT), by domain transport.* Pair $f$ against a fixed $\varphi$ and differentiate the two regime integrals along $q$ by the transport formula for a moving domain, the boundary moving at normal speed $v_\Gamma$ [Delfour and Zolésio, 2011, Ch. 9, Thm. 4.2, eq. (4.7)], the outer normals as in Step 2:

$$\begin{aligned}
\frac{d}{dq}\int_{\mathrm{binding}(q)} f^-\ \varphi\ da \ &=\ \int \{D_q f^-\}\,\varphi\ da \ +\ \int_\Gamma f^-\ v_\Gamma\ \varphi\big|_\Gamma\ d\Gamma,\\
\frac{d}{dq}\int_{\mathrm{interior}(q)} f^+\ \varphi\ da \ &=\ \int \{D_q f^+\}\,\varphi\ da \ -\ \int_\Gamma f^+\ v_\Gamma\ \varphi\big|_\Gamma\ d\Gamma,
\end{aligned}$$

the binding regime gaining the integrand where its boundary advances and the interior regime losing it; adding the two lines assembles $\{D_q f\} - [f]\,v_\Gamma\,\delta_\Gamma$ [Kanwal, 1998, Sect. 5.5, Thm. 1, eq. (5)]. Only the normal speed appears: a tangential reparameterization of the moving interface changes neither line, the structure D-Z record by reducing their velocity field to its normal component [Delfour and Zolésio, 2011, Ch. 9, eq. (4.14)].

*Step 5: ($\mathrm{LT}_q$), by boundary transport.* The pairing $h(q) = \int_{\Gamma(q)} w\,\varphi|_\Gamma\ d\Gamma$ is an integral over a moving hypersurface; the boundary transport formula, written in material form, reads

$$\frac{d}{dq}\int_{\Gamma(q)} \psi\ d\Gamma \ =\ \int_\Gamma \Big[\ \dot\psi\ +\ \left(\mathrm{div}_\Gamma\,\mathbf{n}\right)v_\Gamma\ \psi\ \Big]\ d\Gamma$$

[Delfour and Zolésio, 2011, Ch. 9, Thm. 4.3, eq. (4.17)]: the integrand moves by its material derivative, and the surface measure itself dilates at the rate $(\mathrm{div}_\Gamma\,\mathbf{n})\,v_\Gamma$ under the normal flow, the density derivative boxed at their eq. (4.12), read through eqs. (5.19) and (5.22) of the same chapter. Apply it to $\psi = w\,\varphi|_\Gamma$: the test function is fixed on $\mathcal{A}$ and only its evaluation point rides the interface, so the trace $\varphi|_\Gamma$ moves along the flow at the rate $v_\Gamma\,(\partial_{\mathbf{n}}\varphi)|_\Gamma$, whose pairing is the degree-one term $-\,w\,v_\Gamma\,\partial_{\mathbf{n}}\delta_\Gamma$ with the sign of Definition 6; collecting the three pieces is ($\mathrm{LT}_q$) [Kanwal, 1998, Sect. 5.5, Thm. 2, eq. (9), the unit-weight case, and Sect. 5.6, eqs. (7b)–(7c) for the weighted structure]. The material derivative is the trap the scalar case hides: the weight moves *with* the interface, not at a frozen position, and for the weights in actual use, jumps and traces of model objects on $\Gamma(q)$, $\dot w$ is computed by the chain rule along the moving trace point, which is exactly what the hierarchy of Corollary 3 differentiates.

*Step 6: closure, the layer count, and a familiar instance.* The four rules return objects of the same two kinds, two-regime functions and layers, so the algebra is closed and every mixed derivative is reached by iterating them: normal parts raise the degree, tangential parts differentiate the weights along $\Gamma$ and add curvature terms. For the count: a layer is born only when a derivative hits a nonzero jump, and a function continuous across the interface has $[f] = 0$, so its first derivatives carry jumps but no layer; the first layer appears at total order two, of degree zero, and each further derivative raises the degree by at most one, so total order $K$ reaches degree $K-2$. The general-degree formulas are printed in Kanwal [1998, Sect. 5.6, Thm. 3, eqs. (16)–(23), and Sect. 5.7] and are consumed here only through the displays above and this count. The constraint chain closes the same way: follow any point $x_q$ of the moving interface, differentiate the trace identity $f^+(x_q;q) = f^-(x_q;q)$ along $q$, and decompose the point's velocity into $v_\Gamma\,\mathbf{n}$ plus tangential slip; Step 1 kills the slip against $[\nabla f]$, so $0 = [\nabla f]\cdot\dot x_q + [D_q f] = [\partial_{\mathbf{n}} f]\,v_\Gamma + [D_q f]$,

independently of the flow chosen. The instance to keep in mind is the saving policy itself: flat at the bound on the binding regime, rising on the interior one, so its state gradient is a bounded two-regime function, no layer yet, and its second state derivative carries the first layer, the surface Dirac on $\Gamma$ weighted by the slope jump, the example $1-|x|$ of Appendix C pushed one derivative further.

*Step 7: the scalar reduction, $d_a = 1$, in Bhandari et al.'s notation.* In a Markov row the interface is the moving crossing point $a_\Gamma(q)$ of Lemma 7's reduction, with $\mathbf{n} = +1$ in the borrowing-constraint picture; there is no tangential direction, $\nabla_\Gamma w = 0$, a point does not bend, $\operatorname{div}_\Gamma \mathbf{n} = 0$, the material and partial derivatives coincide, $\dot{w} = D_q w$, and the normal speed is the point's speed, $v_\Gamma = D_q a_\Gamma$. This rung exists to match the scalar generalized-function displays of Bhandari et al. [2023], so it is written in their Dirac notation, the one place outside the comparison paragraphs where it appears: the four rules collapse to

$$\begin{aligned} D_a f &= \{D_a f\} + [f]\,\delta, & D_q f &= \{D_q f\} - [f]\,\big(D_q a_\Gamma\big)\,\delta, \\ D_a\big(w\,\delta^{(i)}\big) &= w\,\delta^{(i+1)}, & D_q\big(w\,\delta^{(i)}\big) &= \big(D_q w\big)\,\delta^{(i)} - w\,\big(D_q a_\Gamma\big)\,\delta^{(i+1)}, \end{aligned}$$

and the chain to $[D_q f] = -[D_a f]\,(D_q a_\Gamma)$. At this rung the whole proof is elementary: split $f = f^- + H\big(a - a_\Gamma(q)\big)(f^+ - f^-)$ with $H$ the Heaviside step and run the Leibniz rule for integrals with moving limits, which is the one-dimensional case of Steps 2 and 4 line for line. □

*Level-set normalization.* When the interface is presented as a level set $\Gamma = \{g = 0\}$ with $\nabla g \neq 0$ on $\Gamma$, the virtual gap $\tilde{g}$ of Definition 5 being the canonical instance, the composite Dirac normalizes by the coarea formula,

$$\delta(g) \;=\; \frac{\delta_\Gamma}{|\nabla g|}$$

(Evans and Gariepy, 2025, Thm. 3.10; Kanwal, 1998, Sect. 5.4, eqs. (1)–(3), who credit Gel'fand and Shilov, 1964): the factor $|\nabla g|$ converts level-set thickness into surface measure. The crease computation of Appendix C, the mixed second derivative of $\max(0, g)$ with its Dirac of strength $(\partial_i g)(\partial_j g)\,\delta(g)$ on the crease, is (JF) applied to $\nabla \max(0, g) = \mathbf{1}\{g > 0\}\,\nabla g$ and read through this normalization: the same object by two routes.

*Where the interface ends.* A boundary $\Gamma_b$ need not be closed: per Markov row it can end, on the boundary of $\mathcal{A}$ or on an intersection with another bound's boundary; at $d_a = 1$ its slice is a finite point set and there is no rim at all, so everything in this remark concerns $d_a \geq 2$. The rim terms are explicit, and they are carried, not killed: the layer weights need not vanish at a rim, the no-flux property governing the distribution's mass at the boundary of $\mathcal{A}$ and not the policy-derivative jumps, and in the two-asset model $[\partial_{\mathbf{n}}\,\bar{x}^*]$ stays away from zero where the liquidity frontier meets the illiquid face. Differentiating a layer along a direction with a tangential component $\tau$ adds, by the Gauss–Green theorem on surfaces [Maggi, 2012, Thm. 11.8, eq. (11.14)], the conormal line term

$$w\;\big(\tau \cdot \nu_{\partial\Gamma}\big)\;\delta_{\partial\Gamma},$$

a codimension-two layer on the rim $\partial\Gamma$ with conormal $\nu_{\partial\Gamma}$, the outward unit tangent of $\Gamma$ normal to its rim; and the parameter rules (PT), ($\mathrm{LT}_q$) add the analogous rim-flux term when the rim itself slides under the perturbation, by the transport theorem on a surface with boundary. Both are legitimate functionals on $V$, a line integral of a continuous test function. What keeps them out of every computation the document runs is the structure of the consumers: rim terms arise only from tangential differentiation, and the recursion ledgers are contracted against $\mathbf{n}$ in every slot, the normal-speed hierarchy is pointwise on $\Gamma$, and the aggregate readings are plain surface integrals; only the uncontracted tangential clauses of the general ledger carry the line term, and they now carry it explicitly. The rim does find a consumer one order up: the order-three aggregate reading at $d_a \geq 2$ differentiates a second-order layer tangentially and meets it, so the bookkeeping is not idle; every display in this document stops short of that, being normal-contracted throughout, which is why none invokes the line term here.

**Corollary 3** (Boundary sensitivities at every order)**.** *Let the perturbation parameters $q$ be the sizes of the announced deviations, as in Lemma 10, and let $f$ be the constrained coordinate of the perturbed policy, continuous across the moving interface along the whole path, $[f](\cdot\,; q) \equiv 0$. Then, pointwise on $\Gamma$:*

1. *at first order, along a perturbation of the relevant aggregates at horizon $s$, the interface moves at the normal speed*

$$v_\Gamma \;=\; -\big[\partial_{\mathbf{n}}\,\bar{x}^*\big]^{-1}\,\big[\times_s\big],$$

*the jump* $[\times_s]$ *read on the constrained coordinate's row, rescaled by the jump of the steady-state normal slope, both bounded functions on* $\Gamma$ *by Proposition 2, the inverse taken on that coordinate; the constrained coordinate does not respond on the binding side, so by Proposition 2 the jump* $[\times_s]$ *is the one-sided interior value, which is what makes* $v_\Gamma$ *computable from the stored arrays; the speed is well defined without choosing a flow on the moving interface, tangential slip dropping out by the purely-normal jump structure of Lemma 10;*

2. *at second order, writing* $v_\Gamma^{(2)}$ *for the material derivative of the speed along the normal flow,*

$$v_\Gamma^{(2)} \;=\; -\,\big[\partial_{\mathbf{n}} f\big]^{-1}\Big(\big[D_q^2 f\big] \;+\; 2\,\big[\partial_{\mathbf{n}} D_q f\big]\,v_\Gamma \;+\; \big[\mathbf{n}\cdot\big(D^2 f\big)\,\mathbf{n}\big]\,v_\Gamma^2\Big),$$

*every jump on the right an already-computed function on* $\Gamma$*; and at every order* $j$ *the speed* $v_\Gamma^{(j)}$ *solves against the* same *leading normal jump* $[\partial_{\mathbf{n}}\,\bar{x}^*]$*, the right-hand side collecting jumps of mixed derivatives of total order at most* $j$*, material where they involve the flow, contracted against speeds of orders below* $j$*, plus, from order three on, terms in* $\nabla_\Gamma v_\Gamma$ *from the rotation of the normal;*

3. *the speed has a closed level-set form in the virtual gap: the moving interface is* $\Gamma(q) = \{\tilde{g}(\cdot\,; q) = 0\}$ *with* $\tilde{g}$ *smooth (Lemma 7, per* $q$ *along the path), so*

$$v_\Gamma \;=\; -\,\frac{D_q \tilde{g}}{|\nabla_a \tilde{g}|},$$

*and the chain returns the same speed whether run on the constrained coordinate, on the multiplier, or on the gap: three carriers, one motion.*

*At* $d_a = 1$ *the speed is the motion of the crossing point,* $v_\Gamma = D_q a_\Gamma$ *and* $v_\Gamma^{(j)} = D_q^j a_\Gamma$*, and items 1–2 collapse to the labeled scalar displays*

$$D_q a_\Gamma \;=\; -\,\big[D_a \bar{x}^*\big]^{-1}\,[\times_s], \qquad D_q^2\, a_\Gamma \;=\; -\,\big[D_a f\big]^{-1}\Big(\big[D_q^2 f\big] + 2\,\big[D_q D_a f\big]\,\big(D_q a_\Gamma\big) + \big[D_a^2 f\big]\,\big(D_q a_\Gamma\big)^2\Big).$$

*Proof. First order.* Follow any point $x_q$ of the moving interface. The constraint chain of Lemma 10, $[D_q f] = -\,[\partial_{\mathbf{n}} f]\,v_\Gamma$, is the trace identity $f^+(x_q; q) = f^-(x_q; q)$ differentiated once, the tangential slip of $x_q$ killed against $[\nabla f] = [\partial_{\mathbf{n}} f]\,\mathbf{n}$; the first-order response in the direction $q$ is $\times_s$ and the steady-state normal jump is $[\partial_{\mathbf{n}}\,\bar{x}^*]$, which is item 1.

*Second order.* Differentiate $0 = [D_q f] + v_\Gamma\,[\partial_{\mathbf{n}} f]$ once more along the normal flow, every derivative material. Two facts organize the computation. First, the normal itself rotates: in the graph chart of Lemma 10, Step 1, with $\Gamma(q) = \{s = \gamma(a'; q)\}$ and $\nabla'\gamma = 0$ at the base point,

$$\dot{\mathbf{n}} \;=\; -\,\nabla_\Gamma v_\Gamma,$$

a tangent vector, as $|\mathbf{n}| = 1$ forces $\mathbf{n}\cdot\dot{\mathbf{n}} = 0$. Second, at this order the rotation drops out: it enters only through $[\nabla f]\cdot\dot{\mathbf{n}} = [\partial_{\mathbf{n}} f]\,(\mathbf{n}\cdot\dot{\mathbf{n}}) = 0$, killed by the same compatibility that made the speed well defined; it survives from order three on, where jumps of objects *not* continuous across the interface, $[\nabla\times_s]$ with its nonzero tangential part, are differentiated along the flow, the $\nabla_\Gamma v_\Gamma$ family flagged in item 2. With the rotation gone, the two jump factors move by

$$\frac{d}{dq}\big[\partial_{\mathbf{n}} f\big] \;=\; \big[\partial_{\mathbf{n}} D_q f\big] \;+\; v_\Gamma\,\big[\mathbf{n}\cdot\big(D^2 f\big)\,\mathbf{n}\big], \qquad \frac{d}{dq}\big[D_q f\big] \;=\; \big[D_q^2 f\big] \;+\; v_\Gamma\,\big[\partial_{\mathbf{n}} D_q f\big],$$

each one-sided limit riding the trace point and moving in the parameter; collecting,

$$0 \;=\; v_\Gamma^{(2)}\,\big[\partial_{\mathbf{n}} f\big] \;+\; 2\,v_\Gamma\,\big[\partial_{\mathbf{n}} D_q f\big] \;+\; v_\Gamma^2\,\big[\mathbf{n}\cdot\big(D^2 f\big)\,\mathbf{n}\big] \;+\; \big[D_q^2 f\big],$$

and solving for $v_\Gamma^{(2)}$ is item 2's display. At order $j$ the one-variable Faà di Bruno expansion of $q \mapsto f(x_q; q)$ along the normal flow, jumped across the interface, isolates $[\partial_{\mathbf{n}} f]\,v_\Gamma^{(j)}$ as the only term carrying the order-$j$ speed, every other term a jump of a material mixed derivative of total order at most $j$ against speeds of orders below $j$: one inversion serves every order, as with $J$ and $\tilde{J}_X$.

*Three carriers.* The multiplier is continuous across the interface, vanishing identically on the interior regime; the gap likewise, vanishing on the binding regime; the constrained coordinate by construction. Running the first-order chain on each gives three expressions for the one motion of the interface, and on the smooth carrier $\tilde{g}$, whose zero set the moving interface is along the whole path (the multiplier-violation identity of Lemma 7 holds per $q$), the chain is the classical level-set formula of item 3, the orientation $\mathbf{n} = \nabla_a \tilde{g}/|\nabla_a \tilde{g}|$ fixing the sign. At $d_a = 1$ every display collapses onto the moving crossing point, $v_\Gamma = D_q a_\Gamma$, and items 1–2 read as the scalar sensitivities of the statement. □

Applied to the steady-state policy, the calculus reproduces mechanically the generalized derivatives that Bhandari et al. [2023, App. B.2, Claim 5] assemble case by case. The policy is continuous with a jumping normal slope, $[\bar{x}^*] = 0$ and $[\partial_{\mathbf{n}}\,\bar{x}^*] \neq 0$, so the first derivatives carry no layer, Proposition 2 again, and the three second derivatives are

$$\begin{aligned} D^2\,\bar{x}^* &= \{D^2\,\bar{x}^*\} + \big[\partial_{\mathbf{n}}\,\bar{x}^*\big]\,(\mathbf{n}\otimes\mathbf{n})\,\delta_\Gamma, \\ D_q\nabla\,\bar{x}^* &= \{D_q\nabla\,\bar{x}^*\} - \big[\partial_{\mathbf{n}}\,\bar{x}^*\big]\,v_\Gamma\,\mathbf{n}\,\delta_\Gamma, \\ D_{q'}D_{q''}\,\bar{x}^* &= \{D_{q'}D_{q''}\,\bar{x}^*\} + \big[\partial_{\mathbf{n}}\,\bar{x}^*\big]\,v_\Gamma^{(q')}\,v_\Gamma^{(q'')}\,\delta_\Gamma. \end{aligned}$$

The first is (JF) applied to each component of the gradient, whose jump is purely normal by the compatibility of Lemma 10, Step 1: the singular part of the Hessian is rank one, symmetric, purely $\mathbf{n}\otimes\mathbf{n}$, with no tangential-tangential layer (that block's *jump* is curvature-bound by the second-order compatibility, but it is a jump of the classical part, not a layer). The second is (PT) applied to the gradient, whose jump is $[\partial_{\mathbf{n}}\,\bar{x}^*]\,\mathbf{n}$; the third is (PT) applied to $D_{q''}\,\bar{x}^*$, whose jump is $-\,[\partial_{\mathbf{n}}\,\bar{x}^*]\,v_\Gamma^{(q'')}$ by the constraint chain, transported at the speed $v_\Gamma^{(q')}$. The middle line is also (JF) applied to the two-regime function $D_q\,\bar{x}^*$: the policy-jump route and the moving-boundary route return the same layer, which is the statement that the gradient and the parameter derivative commute on the singular part exactly when the constraint chain holds. At $d_a = 1$, with $\Delta \equiv [D_a\bar{x}^*]$ and $v_\Gamma^{(q)} = D_q a_\Gamma$, the three displays reduce to the scalar table

$$\begin{aligned} D_a^2\,\bar{x}^* &= \{D_a^2\,\bar{x}^*\} + \Delta\,\delta, \qquad D_qD_a\,\bar{x}^* = \{D_qD_a\,\bar{x}^*\} - \Delta\,\big(D_q a_\Gamma\big)\,\delta, \\ D_{q'}D_{q''}\,\bar{x}^* &= \{D_{q'}D_{q''}\,\bar{x}^*\} + \Delta\,\big(D_{q'}a_\Gamma\big)\big(D_{q''}a_\Gamma\big)\,\delta, \end{aligned}$$

their table of $\mathring{x}$-plus-$\delta$ corrections, line by line. The weight in the third display is the second-order boundary term of Section 4.5, the $C_{1,1}$ the $P^1$ histogram misses [Bhandari et al., 2023, App. C.1]; note the division of labour, for the display carries the pointwise weight while its contribution to an aggregate pairs that weight against the trace of the interior density along the interface, the mass the moving boundary sweeps. The structure is the second derivative of an integral over a moving domain, the boundary term of the transport formula of shape calculus [Delfour and Zolésio, 2011, Ch. 9, Thm. 4.2, eqs. (4.6)–(4.7), and Sect. 6.1, eq. (6.2), for the second order]. In the frame of Assumption 1 a layer $w\,\partial_{\mathbf{n}}^i\delta_\Gamma$ pairs through the trace of the $i$-th normal derivative of the test function (Definition 6), bounded on the grade $V_{i+1}$ introduced next: the derivatives through order $i$ of a $V_{i+1}$ function are $W^{1,p}$ componentwise, their traces land in the fractional class of Theorem 4's refinement, and contracting against $\mathbf{n}^{\otimes i}$, Lipschitz on the $C^{1,1}$ interface, keeps the class; each further order adds one degree, $\partial_{\mathbf{n}}^{K-2}\delta_\Gamma$ at total order $K$ (Step 6 of the proof).

The aggregate reading deserves its own display, because it is where a layer stops being machinery and becomes a number. Pair the first of the three second derivatives against the cross-section, clause 3 of Assumption 3 supplying the trace of the interior density, and clause 4 of Assumption 1, cited as the open genericity remark it is, keeping the singular carriers off $\Gamma$:

$$\int D^2\,\bar{x}^*\,d\bar{\Omega}^* = \int \{D^2\,\bar{x}^*\}\,d\bar{\Omega}^* + \sum_\theta \int_{\Gamma(\theta)} \big[\partial_{\mathbf{n}}\,\bar{x}^*\big]\,(\mathbf{n}\otimes\mathbf{n})\,\rho^*\big|_\Gamma\,d\Gamma,$$

the per-regime classical part integrated as usual, and the layer read against the trace $\rho^*|_\Gamma$ of the interior density, one surface integral per Markov state, single-valued by the continuity of $\rho^*$ across $\Gamma$ that clause 3 of Assumption 3 supplies; the null-atom condition $\bar{\Omega}^*(\Gamma) = 0$ alone would not do, ruling out a mass sitting on $\Gamma$ but not fixing the boundary value of a density that jumped. That trace is the smooth interior one, and two facts leave no other candidate: the layer weight $[\partial_{\mathbf{n}}\,\bar{x}^*]$ is itself one-sided interior, the constrained coordinate not responding on the binding side (Corollary 3, item 1); and any constrained-side singular structure, the mass piled onto a constraint wall that $\Gamma$ runs close to through the mass-heavy region, is a distinct carrier held in the register's singular part $\sum_f \sigma_f^*$, a layer or, at $d_a = 1$, an atom, and kept off $\Gamma$ by clause 4 of Assumption 1, exactly as Bhandari et al. [2023] carry that mass separately from the continuous density. The $C_{1,1}$ boundary term of the order-two aggregates is the same surface reading of the third display,

$$\sum_\theta \int_{\Gamma(\theta)} \big[\partial_{\mathbf{n}}\,\bar{x}^*\big]\,v_\Gamma^{(q')}\,v_\Gamma^{(q'')}\,\rho^*\big|_\Gamma\,d\Gamma.$$

At $d_a = 1$ the surface integrals collapse to point evaluations, the trace to the value $\rho^*(\theta, a_\Gamma(\theta))$, and the first display reduces to

$$\int D_a^2\, \bar{x}^*\; d\bar{\Omega}^* \;=\; \int \{D_a^2\, \bar{x}^*\}\; d\bar{\Omega}^* \;+\; \sum_{\theta} \big[D_a \bar{x}^*\big](\theta)\; \rho^*\big(\theta, a_\Gamma(\theta)\big),$$

which is Bhandari et al.'s integral display in the same Claim 5. The boundary term is the constrained households' own contribution to the order-two aggregates, the jump in their marginal response weighted by the mass sitting at the binding boundary $\Gamma$, and it is exactly what a per-regime differentiation, or the $P^1$ histogram with $\mathcal{P}_{aa} = 0$, silently drops; the paragraph on what the element must carry returns to it. And the general-$d_a$ surface integrals are not a notational upgrade of a fiberwise scalar reading: the components of $\mathbf{n} \otimes \mathbf{n}$ off the normal axis carry weights no one-dimensional slice produces.

**The distribution at order two.** The transfer side differentiates once more as well, and the computation is the same second difference quotient a Hessian comes from. Perturb the policy along two directions, $\bar{a}^* \rightsquigarrow \bar{a}^* + h\, da + h\, db$, and recall from Proposition 3 that the pairing $\langle \Lambda[\bar{a}]\, \bar{\Omega}^*, \varphi \rangle$ depends on the policy only through the landing point inside the test function, $\varphi(\theta', \bar{a}(\theta, a))$. The second derivative in $h$ at zero is then, pointwise under the integral, the second difference quotient

$$\frac{1}{h^2}\Big[\, \varphi\big(\theta',\ \bar{a}^* + h\, da + h\, db\big) \;-\; \varphi\big(\theta',\ \bar{a}^* + h\, da\big) \;-\; \varphi\big(\theta',\ \bar{a}^* + h\, db\big) \;+\; \varphi\big(\theta',\ \bar{a}^*\big) \,\Big] \;\longrightarrow\; \nabla^2_{a'} \varphi\big(\theta', \bar{a}^*\big)\big[\, da,\ db\,\big],$$

where the limit, written $\nabla^2_{a'}\varphi\,[da, db] = \sum_{j,l} \partial^2_{a'^j a'^l} \varphi\; da^j\, db^l$, is the Hessian of the test function contracted against the two displacement directions; dominated convergence, exactly as in Proposition 3's proof but one derivative higher, carries the limit through the integrals, giving

$$\big\langle \big(D^2_{\bar{a}^*}\Lambda\big)(da, db)\; \bar{\Omega}^*,\ \varphi \big\rangle \;=\; \int\!\!\int \nabla^2_{a'} \varphi\big(\theta', \bar{a}^*(\theta, a)\big)\big[\, da(\theta, a),\ db(\theta, a)\,\big]\; d\pi(\theta' \mid \theta)\; d\bar{\Omega}^*(\theta, a)\,:$$

the second derivative of the transfer operator reads test functions through their Hessians, where the first read them through their gradients. To name the object this pairing defines, introduce the second-order displacement $\mathsf{M}_2$ componentwise,

$$\big\langle \big(\mathsf{M}_2 \cdot (da, db)\big)^{jl},\ \psi \big\rangle \;=\; \int\!\!\int \psi\big(\theta', \bar{a}^*(\theta, a)\big)\; da^j(\theta, a)\; db^l(\theta, a)\; d\pi(\theta' \mid \theta)\; d\bar{\Omega}^*(\theta, a),$$

one component per pair of coordinates, the mass-weighted joint displacement of the arriving cohorts, the order-two sibling of $\mathsf{M}$; and define the weak second divergence of such a matrix-valued object by its pairing, $\langle \nabla^2_{a'} : \mathsf{m},\ \varphi \rangle \equiv \sum_{j,l} \langle \mathsf{m}^{jl},\ \partial^2_{a'^j a'^l} \varphi \rangle$, the double-dot contraction summing both indices, just as the weak divergence of Proposition 3 paired against single derivatives. With these two definitions the display above condenses to

$$\big(D^2_{\bar{a}^*}\Lambda\big)(da, db)\; \bar{\Omega}^* \;=\; \nabla^2_{a'} : \big(\mathsf{M}_2 \cdot (da, db)\big),$$

with a positive sign: moving two derivatives onto the test function costs $(-1)^2$, against the single minus of Proposition 3. Now differentiate the law of motion $\bar{\Omega}_{t+1} = \Lambda[\bar{a}_t]\, \bar{\Omega}_t$ twice along a perturbed path. The pattern is the second derivative of a product, $d^2(AB) = A''B + 2A'B' + AB''$, with $A$ the operator moving through its policy argument and $B$ the distribution path; matching term by term yields

$$\hat{\Omega}_{t+1}{}^{(2)} \;=\; \Lambda\, \hat{\Omega}_t{}^{(2)} \;+\; 2\,\big(D_{\bar{a}^*}\Lambda \cdot \hat{a}_t\big)\, \hat{\Omega}_t \;+\; \big(D_{\bar{a}^*}\Lambda \cdot \hat{a}_t{}^{(2)}\big)\, \bar{\Omega}^* \;+\; \big(D^2_{\bar{a}^*}\Lambda\big)\big(\hat{a}_t, \hat{a}_t\big)\, \bar{\Omega}^*,$$

the propagation of the order-two response, the first-order displacement applied to the moving distribution, the displacement of $\bar{\Omega}^*$ by the order-two policy direction, and the second displacement above; the first and third reuse the order-one operators, the second and fourth are new at order two, and the discrete recursion implementing exactly these four terms is the second-order distribution block of Bhandari et al. [2023, Cor. 2]. The same two rules, the product rule on the pair $(\Lambda, \bar{\Omega})$ and the partition formula on the composition $\Lambda[\bar{a}(\cdot)]$, close the law of motion at every order: differentiating $\bar{\Omega}_{t+1} = \Lambda[\bar{a}_t]\, \bar{\Omega}_t$ along $m$ perturbation directions,

$$\hat{\Omega}_{t+1}{}^{(m)} \;=\; \Lambda\, \hat{\Omega}_t{}^{(m)} \;+\; \sum_{\emptyset \neq S \subseteq \{1, \ldots, m\}}\; \sum_{\pi \vdash S} \Big(D^{|\pi|}_{\bar{a}^*}\Lambda\Big)\big(\hat{a}_t{}^{(B_1)}, \ldots, \hat{a}_t{}^{(B_{|\pi|})}\big)\; \hat{\Omega}_t{}^{(S^c)},$$

each of the $m$ derivatives landing on the operator factor, through a partition $\pi$ of its labels $S$, or on the distribution factor, which keeps the complementary labels $S^c$, with $\hat{\Omega}_t{}^{(\emptyset)} = \bar{\Omega}^*$; at $m = 2$ the subsets $\{1\}, \{2\}, \{1,2\}$ reproduce the four terms and the weight 2 of the display above. Each new order adds one deeper displacement: $D^j_{\bar{a}^*}\Lambda$ pairs test functions through their $j$-th derivatives, the weak $j$-fold divergence of the $j$-linear displacement $\mathsf{M}_j$, by the argument of Proposition 3 run $j-1$ derivatives higher: one more derivative of the test functions per order, the count the grades below organize. Two costs are visible in the four-term display. Its second term displaces the first-order response $\hat{\Omega}_t$, which contains dipole terms of $D$, and displacing a dipole reads one more gradient; the fourth term pairs through Hessians outright. Each order thus consumes one more derivative of the test functions, and the natural frame is graded.

**Definition 7** (Grades)**.** The *grades* are the nested test spaces

$$V_k \;=\; \ell^2(\Theta) \otimes W^{k,p}(\mathcal{A}), \qquad p > d_a, \qquad V = V_1 \supset V_2 \supset \cdots,$$

and the order-$k$ frame for the distribution-side responses is $E_k = V_k' \oplus D_k$, the grade-$k$ dual plus the singular part $D_k$: finitely many carriers, each bearing a weight, the orbit atoms with their dipole directions, finite vectors per point, the layers $\sigma_f^*$ of Assumption 1 with their displacement layers on their carriers, and the binding boundaries $\Gamma_b$ with the layers $w\,\partial_{\mathbf{n}}^i\delta_\Gamma$ of Lemma 10 to degree $k-1$, each weight a density on its carrier in the classes of Definition 6. What is finite at general $d_a$ is the carrier count, not the dimension of $D_k$; at $d_a = 1$ the carriers are points, the weights scalars, and $D_k$ a finite-dimensional span. At $k = 1$, $E_1 = E = V' \oplus D$.

The embedding condition is unchanged, only the continuity of the test functions themselves ever being evaluated. I state the grades as an extension of Assumption 1 and use them; the smooth-case analogue is the graded family of spaces of Tanaka [2022, Sect. 2.1, condition (I)], one grade lost per order of expansion, a device he credits to Gouëzel and Liverani [2006]. Two facts make the grades serviceable, and both run on machinery already in place. First, the conditional-expectation operator $\mathcal{K}$ is bounded on each grade *regime by regime*: on a regime every pinned coordinate of $\bar{a}^*$ is constant, so every derivative of the composition $\varphi \circ \bar{a}^*$ reads $\varphi$ only along the free landing directions, where clause 2 of Assumption 1 runs the change of variables; this is Lemma 12's argument iterated through the Faà di Bruno expansion, and the per-regime reading is essential, the policy Jacobian's jump across $\Gamma$ already denying $\mathcal{K}\varphi$ a global second derivative. The iterates $\mathcal{K}^m\varphi$ are then piecewise-smooth with respect to the partition refined by the pre-images of $\Gamma$ under the per-regime policy, and one clause disciplines those seams, the sibling of clause 4: *the iterated pre-images of* $\Gamma$ *under the per-regime steady-state policy are transversal, hence* $C^k$ *hypersurfaces, finitely many per horizon*; every pairing below reads the iterates off these seams or through the priced traces. Second, the transfer operator respects the register $D_k$ carrier by carrier: the per-regime policy restricted to a carrier is bi-Lipschitz onto its image (the nondegenerate free block restricted to the tangent is injective), the transformed weight is the old one composed with the inverse and multiplied by a tangential Jacobian pinched by clause-2 constants, and the fractional classes $W^{c_f/p,\,p'}$ survive both operations, the Gagliardo seminorm transforming under bi-Lipschitz maps with equivalent distances; layers of degree $i \le k-1$ transform through the rules of Lemma 10 within the per-regime $C^k$ budget. The quantitative decay rates of the iterates are not part of these statements: they are the graded clauses of Appendix E, the spaces here and the sizes there.

**The general-equilibrium closure.** The pieces assemble as at first order, and into the same operator; this is the subsection's central fact. Two objects prepare the statement. First, the *gap kernel*: away from date 0 in every index, the order-$m$ response at date $t$ to innovations dated $s_1, \ldots, s_m$ will turn out to depend only on the *gaps* $u_i = t - s_i$, the time differences between the reading date and the shock dates (no kinship with the bound's gap $g^*$ of Definition 1), and the array read as a function of the gaps, $H_{u_1 \cdots u_m}$, is the response's gap kernel; at $m = 1$ the gap is $t - s$ and the gap kernel is the familiar sequence of impulse-response coefficients. Second, its *transform*: the $m$-variable Laurent series

$$H^{(m)}(z_1, \ldots, z_m) \;=\; \sum_{u_1, \ldots, u_m \in \mathbb{Z}} H_{u_1 \cdots u_m}\, z_1^{u_1} \cdots z_m^{u_m},$$

the direct generalization of the symbol of Definition 4, which is the $m = 1$ case; it converges on a polyannulus, a product of annuli one per variable, whenever the kernel decays geometrically in every index, and the *torus* is the product of unit circles $|z_1| = \cdots = |z_m| = 1$, the $m$-variable stand-in for the unit circle of the determinacy analysis.

**Proposition 12** (One factorization serves every order)**.** *Let Assumptions 1 and 2 and the settings of Propositions 2 and 11 be in force; let* $\det j_X \neq 0$ *on the unit circle with* $\mathrm{wind}(\det j_X) = 0$ *and trivial kernel, so that* $\tilde{J}_X$ *is invertible (Proposition 7) and* $j_X$ *admits the canonical factorization* $j_X = j_- j_+$ *of Proposition 9; and assume the* order-$k$ decay hypothesis*, stated on kernel classes patterned on the bounds of Proposition 6: with* $r(u) = \lambda_m^u$ *for* $u \geq 0$ *and* $r(u) = \rho^{-u}$ *for* $u < 0$*, built on the anticipation rate* $\rho$ *of that proposition and, per order, on the root enlargement* $\lambda_m \equiv \lambda^{1/m}$ *of its propagation rate* $\lambda$*, an enlargement Appendix E derives and shows exact, the corner part below keeping the first order's own* $\lambda$*, and*

$$\mathcal{K}_m \;\equiv\; \Big\{ \big(H_{t;\,s_1\cdots s_m}\big)_{t\geq 0} \;:\; \big\|H_{t;\,s_1\cdots s_m}\big\| \;\leq\; C\Big[\; \prod_{i=1}^{m} r(t-s_i) \;+\; \lambda^t \prod_{i=1}^{m} \rho^{\,s_i} \;\Big] \text{ for some } C \Big\},$$

*an interior part decaying at the rates* $r$ *in every gap plus a part decaying like* $\mathsf{C}_X$ *in the row and in every column, assume that the multilinear derivatives of the stacked system* $\mathcal{G}$ *of Corollary 2, taken at the steady state, preserve the classes: for every* $p$ *and* $m_1 + \cdots + m_p = m \leq k$*,*

$$K_i \in \mathcal{K}_{m_i},\; i = 1,\dots,p \qquad \Longrightarrow \qquad D^p\mathcal{G}\big(K_1,\dots,K_p\big) \;\in\; \mathcal{K}_m.$$

*The hypothesis is not bare: Appendix E derives it at* $m = 2$ *for the three families the source actually contains, from two graded strengthenings of Assumption 2 stated there, the general order following by the same induction; I keep the class form as the hypothesis because the graded clauses are quantitative: the grade frame itself is established after Definition 7, and what remains assumed is the rates. Then, for every order* $2 \leq m \leq k$*:*

1. (closure) *for each tuple of shock dates* $(s_1,\dots,s_m)$*, the order-*$m$ *response path solves the first-order stacked operator against a known source,*

$$\tilde{J}_X \left( \widehat{X}^{(0)}_{t-s_1\,\cdots\,t-s_m} \right)_{t\geq 0} \;=\; -\,\Xi^{(m)}\big(\cdot\,;\,s_1,\dots,s_m\big),$$

*the source the coefficient of the full monomial* $\bar{\mathcal{E}}_{s_1}\cdots\bar{\mathcal{E}}_{s_m}$ *in the partition sum of the* $\sigma$*-differentiated stacked equations, every slot of which carries kernels of orders below* $m$ *only,*

$$\Xi^{(m)}\big(\cdot\,;\,s_1,\dots,s_m\big) \;=\; \sum_{\substack{\pi\vdash\{1,\dots,m\}\\ |\pi|\geq 2}} D^{|\pi|}\mathcal{G}\Big(\hat{h}_{s_{B_1}},\;\dots,\;\hat{h}_{s_{B_{|\pi|}}}\Big), \qquad \hat{h}_{s_B} \;\equiv\; \begin{cases} \Big( \big(\widehat{X}^{(0)}_{t-s_i}\big)_{t\geq 0},\; \mathbf{1}\{\cdot = s_i\} \Big) & B = \{i\}, \\ \Big( \big(\widehat{X}^{(0)}_{t-s_B}\big)_{t\geq 0},\; 0 \Big) & |B| \geq 2, \end{cases}$$

*each slot fed the whole date-path of its block's kernel, the order-*$|B|$ *family* $\widehat{X}^{(0)}_{t-s_B}$ *carrying* $q = 0$ *risk slots and* $|B|$ *lag indices, next to the block's own innovation date; unpacking the rows of* $\mathcal{G}$ *sorts the terms into the multilinear forms of the aggregate equations, the heterogeneous residual's terms through Lemma 8, and the distribution terms of the preceding paragraph with their ledgers* $\mathsf{L}_\Gamma$ *(Lemma 10);*

2. (interior formula) *away from date* $0$ *in every index, source and solution depend on the gaps* $(t - s_1,\dots,t-s_m)$ *alone; writing* $\Xi^{(m)}(z_1,\dots,z_m)$ *and* $H^{(m)}(z_1,\dots,z_m)$ *for the transforms of the gap kernels, the whole-line comparison problem solves*

$$j_X\big(z_1\cdots z_m\big)\; H^{(m)}\big(z_1,\dots,z_m\big) \;=\; -\,\Xi^{(m)}\big(z_1,\dots,z_m\big)$$

*on the torus* $|z_1| = \cdots = |z_m| = 1$*: the first-order symbol, evaluated at the product of the transform variables, is the only object ever inverted;*

3. (the correction terms, exact) *on the half line the inversion is exact in two triangular passes and departs from the whole-line reading by a product of two Hankel operators,*

$$\begin{aligned} \mathsf{T}(j_X)^{-1} \;&=\; \mathsf{T}\big(j_+^{-1}\big)\,\mathsf{T}\big(j_-^{-1}\big) \;=\; \mathsf{T}\big(j_X^{-1}\big) \;-\; \mathsf{H}\big(j_+^{-1}\big)\,\mathsf{H}\big(\tilde{j}_-^{-1}\big), \\ \tilde{J}_X^{-1} \;&=\; \mathsf{T}(j_X)^{-1} \;-\; \mathsf{T}(j_X)^{-1}\,\mathsf{C}_X\big(I + \mathsf{T}(j_X)^{-1}\,\mathsf{C}_X\big)^{-1}\mathsf{T}(j_X)^{-1}, \end{aligned}$$

*with* $[\mathsf{H}(a)]_{t,s} = a_{t+s+1}$ *and* $\tilde{j}_-(z) \equiv j_-(1/z)$*; the order-*$m$ *solution is therefore the interior formula plus four correction terms, the Hankel product applied to the source, the pre-date-*$0$ *tail of the interior convolution, the image of the source's own edge, and the resolvent term in* $\mathsf{C}_X$*; the pre-date-*$0$ *tail keeps the full corner profile* $C\,\lambda^t \prod_i \rho^{\,s_i}$*, while the other three carry the bound* $C\,\lambda^t \rho^{\max_i s_i}$*, geometric in the row and in the latest shock date, which is all the summability below uses (Appendix E);*

4. (order-free determinacy) *if* $|z_i| = 1$ *for every* $i$ *then* $|z_1 \cdots z_m| = 1$*, so the hypothesis on the circle is exactly what makes the interior inversion well defined on the whole torus at every order; and the factorization* $j_- \, j_+$*, computed once, solves every source by one anti-causal and one causal pass.*

*Proof. Step 1: closure.* Recall first what the object being differentiated is: the stacked system of Corollary 2 carries one row per aggregate equation and date, each row an equation in the whole aggregate path and the whole innovation path, with the heterogeneous block solved inside it; along the $\sigma$-path both arguments move, the aggregate path as $X(\sigma)$ and the innovation path as $\sigma\bar{\mathcal{E}}$, and every row holds identically in $\sigma$, so all $\sigma$-derivatives of every row vanish. Differentiate $m$ times and sort the resulting terms, exactly as in Lemma 8. Run $m = 2$ first: write $\mathcal{G}$ for the stacked system as a map of the pair, and note that the direction of travel is $h(\sigma) = \big(X(\sigma),\ \sigma\bar{\mathcal{E}}\big)$, whose first $\sigma$-derivative is $(X^{(1)}, \bar{\mathcal{E}})$ and whose second is $(X^{(2)}, 0)$, the innovation component being linear in $\sigma$ and contributing nothing beyond the first order. The second $\sigma$-derivative of $\mathcal{G}(h(\sigma)) = 0$ is then, by Step 1 of that lemma,

$$D\mathcal{G} \cdot \big(X^{(2)},\ 0\big) \ + \ D^2\mathcal{G}\Big(\big(X^{(1)}, \bar{\mathcal{E}}\big),\ \big(X^{(1)}, \bar{\mathcal{E}}\big)\Big) \ = \ 0,$$

the slope of the stacked system against the order-two path plus its bilinear form against two copies of the first-order pair. The slope, restricted to directions moving the aggregate path alone, is by construction the operator of Corollary 2, $\tilde{J}_X$: the same linear map whether the direction it acts on is a first-order response or an order-two kernel, because a derivative is taken at a base point and the base point has not moved. At order $m$ the partition formula of Lemma 8, Step 2, puts the single-block partition on the left and every coarser partition on the right, the blocks fed by the derivatives of the direction of travel,

$$\tilde{J}_X \ X^{(m)} \ = \ - \sum_{\substack{\pi \vdash \{1,\dots,m\} \\ |\pi| \geq 2}} D^{|\pi|}\mathcal{G}\Big(h^{(|B_1|)},\ \dots,\ h^{(|B_{|\pi|}|)}\Big), \qquad h^{(1)} = \big(X^{(1)},\ \bar{\mathcal{E}}\big), \quad h^{(i)} = \big(X^{(i)},\ 0\big) \ \text{ for } i \geq 2,$$

every block on the right of order below $m$: the source. It remains to split by lag tuples. By Proposition 11, the $m$-th $\sigma$-derivative of the path is a polynomial in the realized unit innovations,

$$X_t^{(m)} \ = \ \sum_{\substack{q+n=m \\ q \text{ even}}} \binom{m}{q} \sum_{s_1,\dots,s_n} \widehat{X}^{(q)}_{t-s_1 \,\cdots\, t-s_n} \ \bar{\mathcal{E}}_{s_1} \cdots \bar{\mathcal{E}}_{s_n},$$

the $q = 0$ term carrying the full monomials, the shock kernels, the even $q \geq 2$ terms the shorter monomials, the risk rungs, and the odd $q$ absent by item 2 of that proposition; the source, built from lower-order objects, has the same polynomial shape. The realized innovations are free variables, so the identity must hold coefficient by coefficient: the coefficient of the full monomial $\bar{\mathcal{E}}_{s_1} \cdots \bar{\mathcal{E}}_{s_m}$ is one stacked path equation per tuple $(s_1, \dots, s_m)$ for the shock kernels, which is item 1, and the shorter monomials yield the risk-rung equations that Proposition 13 takes up. Existence and uniqueness of each column in $\ell^2$ is the assumed invertibility of $\tilde{J}_X$, granted $\Xi^{(m)} \in \ell^2$, which the decay hypothesis provides. At order two and in the state space this closure, for the shock kernels and for the risk block alike, is Proposition $2^{HA}$ of Bhandari et al. [2023, eqs. (47)–(48)]; the order-$k$ statement, the symbol form, and the correction terms are what the present framework adds.

*Step 2: gap kernels.* The *whole-line comparison problem* is the same stacked system with dates running over all of $\mathbb{Z}$ rather than $t \geq 0$, the economy sitting at the steady state in the infinitely remote past; it is the natural comparison because, at first order, it is exactly the object the diagonal limits of Proposition 6, item 2, converge to, the half-line problem forgetting its date-0 corrections at the two geometric rates. On the whole line the only datum breaking translation invariance is the innovations themselves: the backward operator $\mathfrak{x}$ and the transfer operator $\Lambda$ carry no date index, so shifting every innovation date and the reading date by one period reproduces the same problem verbatim, and the response coefficient cannot change. Concretely at $m = 2$, write $X_t^{(2)}\big|_{\mathcal{E}_{s_1}\mathcal{E}_{s_2}}$ for the coefficient of $\mathcal{E}_{s_1}\mathcal{E}_{s_2}$ in the date-$t$ response; the one-period shift gives

$$X_t^{(2)}\Big|_{\mathcal{E}_{s_1}\,\mathcal{E}_{s_2}} \ = \ X_{t+1}^{(2)}\Big|_{\mathcal{E}_{s_1+1}\,\mathcal{E}_{s_2+1}} \qquad \Longrightarrow \qquad X_t^{(2)}\Big|_{\mathcal{E}_{s_1}\,\mathcal{E}_{s_2}} \ = \ H_{t-s_1,\,t-s_2},$$

the coefficient depending only on the gaps, so the response has a gap kernel in the sense defined above; the same shift argument applies to the source, whose ingredients are steady-state forms contracted against lower-order kernels, themselves gap kernels by induction on $m$. The decay hypothesis bounds both at

the two rates in every index, so their transforms converge on a polyannulus containing the torus. In the running illustration the first-order gap kernel is the anticipation sequence read off earlier, $x_u = (\beta\varphi')^{-u}$ for gaps $u \le 0$ and zero for $u > 0$: the date-$t$ response to an innovation $|u|$ periods ahead.

*Step 3: the diagonal convolution.* Now apply the Toeplitz part of the operator to a response with gap kernel $H_{u_1 \cdots u_m}$ and read, at each date, the coefficient of the innovation product $\mathcal{E}_{s_1} \cdots \mathcal{E}_{s_m}$. The Toeplitz action convolves over the reading date,

$$\Big(\mathsf{T}(j_X)\, X^{(m)}\Big)_t\Big|_{\mathcal{E}_{s_1}\cdots\mathcal{E}_{s_m}} \;=\; \sum_v j_{X,v}\; H_{t-v-s_1,\,\ldots,\,t-v-s_m},$$

and here is the crucial geometric fact: shifting the reading date by $v$ shifts *all $m$ gaps together*, so the convolution runs along the diagonal of the gap kernel. To transform it, multiply by $z_1^{t-s_1}\cdots z_m^{t-s_m}$, sum over the gaps, and split each power along the diagonal,

$$z_1^{u_1}\cdots z_m^{u_m} \;=\; \big(z_1 z_2 \cdots z_m\big)^{v}\; z_1^{u_1-v}\cdots z_m^{u_m-v},$$

one factor carrying the diagonal shift $v$ and the other the shifted gaps; summing first over the shifted gaps at fixed $v$, then over $v$, the double sum factorizes,

$$\begin{aligned}\sum_{u_1,\ldots,u_m} z_1^{u_1}\cdots z_m^{u_m} \sum_v j_{X,v}\, H_{u_1-v,\ldots,u_m-v} \;&=\; \Big(\sum_v j_{X,v}\,\big(z_1\cdots z_m\big)^{v}\Big)\Big(\sum_{u_1,\ldots,u_m} z_1^{u_1}\cdots z_m^{u_m}\, H_{u_1\cdots u_m}\Big)\\ &=\; j_X\big(z_1\cdots z_m\big)\; H^{(m)}\big(z_1,\ldots,z_m\big),\end{aligned}$$

the one-variable symbol evaluated at the product of the arguments: the whole-line equation transforms into item 2's display. In continuous time and for Volterra kernels this is the classical rule that a linear system acting on the shared time axis multiplies the $m$-variable transform by its transfer function at the *sum* of the frequencies [Rugh, 1981, Thm. 2.3]; the product of $z$-arguments is the discrete image of that sum, the exponential map carrying frequencies to the circle. In the running illustration the whole computation can be checked by hand at order two: differentiating $X_t = \beta\varphi(X_{t+1}) + Z_t$ twice against innovations at dates $s_1$ and $s_2$ gives the gap-kernel recursion

$$H_{u_1,u_2} \;-\; \beta\varphi'\, H_{u_1+1,\,u_2+1} \;=\; \beta\varphi''\, x_{u_1+1}\, x_{u_2+1},$$

next period seeing both gaps advanced by one, and the two-variable transform of "advance both gaps" is multiplication by $(z_1 z_2)^{-1}$, so the transform of the left side is $\big(1 - \beta\varphi'\,(z_1z_2)^{-1}\big)H^{(2)}(z_1,z_2) = j_X(z_1z_2)\,H^{(2)}(z_1,z_2)$, the interior formula verbatim.

*Step 4: the correction terms.* The half-line operator is not the whole-line convolution, and the departures are what this step computes, exactly. The tool is a second operator built on a symbol, so I define it first. Where the Toeplitz operator reads the coefficient at the *difference* of its indices, the *Hankel* operator reads it at their sum,

$$\big[\mathsf{T}(a)\big]_{t,s} \;=\; a_{t-s}, \qquad \big[\mathsf{H}(a)\big]_{t,s} \;=\; a_{t+s+1}, \qquad t,s \ge 0,$$

constant along diagonals for the first, along anti-diagonals for the second; their upper-left corners are

$$\mathsf{T}(a) \;=\; \begin{pmatrix} a_0 & a_{-1} & a_{-2} & \\ a_1 & a_0 & a_{-1} & \cdots \\ a_2 & a_1 & a_0 & \\ & \vdots & & \ddots \end{pmatrix}, \qquad \mathsf{H}(a) \;=\; \begin{pmatrix} a_1 & a_2 & a_3 & \\ a_2 & a_3 & a_4 & \cdots \\ a_3 & a_4 & a_5 & \\ & \vdots & & \ddots \end{pmatrix}.$$

A Toeplitz operator shifts and weighs; a Hankel operator reads its input *mirrored*, pairing the $s$-th entry of the input with ever deeper coefficients, and it is the natural bookkeeper for whatever couples the two sides of a time boundary. The one identity needed about them is checked entrywise in three lines. The half-line product has entries

$$\big[\mathsf{T}(a)\,\mathsf{T}(b)\big]_{t,s} \;=\; \sum_{r\ge 0} a_{t-r}\, b_{r-s},$$

the sum running over the half-line index $r \ge 0$ only, while the coefficient of the product symbol sums over all of $\mathbb{Z}$, $[\mathsf{T}(ab)]_{t,s} = (ab)_{t-s} = \sum_{r\in\mathbb{Z}} a_{t-r}b_{r-s}$; the difference is the missing half, and re-indexing it by $r = -r' - 1$ with $r' \ge 0$,

$$\sum_{r<0} a_{t-r}\, b_{r-s} \;=\; \sum_{r'\ge 0} a_{t+r'+1}\; b_{-(s+r'+1)} \;=\; \sum_{r'\ge 0} \big[\mathsf{H}(a)\big]_{t,r'}\,\big[\mathsf{H}(\tilde b)\big]_{r',s},$$

using $\tilde{b}_v = b_{-v}$, so that

$$\mathsf{T}(a)\,\mathsf{T}(b) \;=\; \mathsf{T}(ab) \;-\; \mathsf{H}(a)\,\mathsf{H}(\tilde{b}),$$

the product identity of Böttcher and Silbermann [1999, Prop. 1.12], stated there with the Hankel product on the other side. Now specialize it three times. For the pair $(j_-, j_+)$: the factor $j_-$ has no positive-index coefficients, so every entry $[\mathsf{H}(j_-)]_{t,s} = (j_-)_{t+s+1}$, index at least one, vanishes, the Hankel term dies, and $\mathsf{T}(j_X) = \mathsf{T}(j_-)\,\mathsf{T}(j_+)$ exactly. For the pairs $(j_\pm^{-1}, j_\pm)$: the inverses of the factors are one-sided on the same side as the factors themselves (Proposition 9), the same argument kills the Hankel term, and each triangular factor inverts inside its own algebra, $\mathsf{T}(j_\pm)^{-1} = \mathsf{T}(j_\pm^{-1})$, whence the two-pass inverse $\mathsf{T}(j_X)^{-1} = \mathsf{T}(j_+^{-1})\,\mathsf{T}(j_-^{-1})$. For the pair $(j_+^{-1}, j_-^{-1})$, finally, *both* Hankel factors survive, since $j_+^{-1}$ has positive-index coefficients and $\tilde{j}_-^{-1}$ likewise, and the identity delivers

$$\mathsf{T}\big(j_+^{-1}\big)\,\mathsf{T}\big(j_-^{-1}\big) \;=\; \mathsf{T}\big(j_X^{-1}\big) \;-\; \mathsf{H}\big(j_+^{-1}\big)\,\mathsf{H}\big(\tilde{j}_-^{-1}\big), \qquad j_+^{-1}\,j_-^{-1} \;=\; j_X^{-1},$$

the first display of item 3. Its reading: $\mathsf{T}(j_X^{-1})$ is the whole-line convolution, the interior formula; the Hankel product is what the half line subtracts from it, one anti-causal Hankel read of the source, reaching into what would have been the pre-date-0 rows, composed with one causal Hankel write of the response. It is the order-$m$ image of the missing anticipation that Proposition 6 isolated in $\mathsf{C}$: the whole-line formula presumes shocks that could have been announced before date 0, and the Hankel product removes precisely that anticipation. In the running illustration $j_+ = 1$, a constant symbol with no positive-index coefficients at all, so $\mathsf{H}(j_+^{-1}) = \mathsf{H}(1) = 0$ and the correction vanishes identically: a purely anticipatory economy has no missing anticipation to remove, because the correction lives exactly where propagation meets anticipation, and the toy has nothing to propagate. The remaining three correction terms are quicker. At a row $t \geq 0$ the interior convolution and its half-line read differ by the source's rows at negative dates, which the half-line problem does not have,

$$\sum_{s\in\mathbb{Z}} \big(j_X^{-1}\big)_{t-s}\,\Xi^{(m)}{}_s \;=\; \Big(\mathsf{T}\big(j_X^{-1}\big)\,\Xi^{(m)}\Big)_t \;+\; \sum_{s<0} \big(j_X^{-1}\big)_{t-s}\,\Xi^{(m)}{}_s,$$

and the pre-date-0 tail on the right fades like the correction, $C\,\lambda^t\rho^s$ entrywise: its coefficients read ever deeper into the symbol, $t - s > t$ when $s < 0$, and the source decays. Near date 0 the multilinear forms differ from their translation-invariant limits, their pre-date-0 slots pinned at the steady state; this edge of the source is bounded exactly as the correction of Proposition 6, item 3, missing anticipation once more. And the correction $\mathsf{C}_X$ enters through the second displayed identity of item 3, which is two lines of algebra: abbreviating $T \equiv \mathsf{T}(j_X)$,

$$\tilde{J}_X^{-1} \;=\; \Big(T\,\big(I + T^{-1}\,\mathsf{C}_X\big)\Big)^{-1} \;=\; \big(I + T^{-1}\,\mathsf{C}_X\big)^{-1}\,T^{-1} \;=\; T^{-1} \;-\; T^{-1}\,\mathsf{C}_X\,\big(I + T^{-1}\,\mathsf{C}_X\big)^{-1}\,T^{-1},$$

the last equality the identity $(I+M)^{-1} = I - M\,(I+M)^{-1}$ at $M = T^{-1}\mathsf{C}_X$. It is algebraically valid once $\tilde{J}_X$ and $\mathsf{T}(j_X)$ are both invertible, the former by hypothesis and the latter by the factorization just established, and its middle factor inherits the $\lambda^t\rho^s$ decay and the stored low-rank compression of $\mathsf{C}_X$ from Section 5.1.

*Step 5: every order at once.* On the torus the product $z_1 \cdots z_m$ has modulus one, so every value the interior formula inverts, $j_X(z_1 \cdots z_m)$, is a value of the symbol on the unit circle, invertible by hypothesis; and each source column is solved by the two triangular passes of item 3 plus the correction terms, the factorization computed once for all orders and all columns. □

In practice the interior term reuses the determinacy grid: the values $j_X(e^{i\omega})$ tabulated for the winding count of Section 5.1 are read at sum frequencies, $j_X(e^{i(\omega_1+\cdots+\omega_m)})$, so the order-two solve is one two-dimensional fast Fourier transform of the source, a pointwise multiplication by a one-variable object, and an inverse transform, zero-padded against aliasing; the correction terms apply through the stored factors, a Hankel product being as cheap as a Toeplitz one. What grows with the order is the count of source columns, not the algebra applied to each.

**The risk corrections.** The even-$\sigma$ kernels are the genuinely stochastic content of Proposition 11, and they obey the same discipline: the same stacked operator $\tilde{J}_X$ closes every rung, and the sources reach the orders below only through their one-period-ahead slots, contracted against the moments of the unit innovation $\bar{\mathcal{E}}$.

**Proposition 13** (The risk rungs)**.** *In the setting of Proposition 12, and with every conditional expectation of the model one period ahead, as the blocks of Section 1 prescribe:*

1. *if the odd moments of the unit innovation vanish through the order of the approximation, every kernel odd in the risk scale vanishes;*

2. *each even risk kernel solves the first-order stacked operator,*

$$\tilde{J}_X\left(\widehat{X}^{(q)}_{t-s_1\,\cdots\,t-s_n}\right)_{t\geq 0} \;=\; -\,\Xi^{(q,n)},$$

*the source given, row by row, by the finite sum, $\mathcal{G}_t$ denoting the date-$t$ row of the stacked system of Corollary 2,*

$$\Xi^{(q,n)}{}_t \;=\; \sum_{p\geq 1}\frac{1}{p!}\sum_{(q_i,\,r_i,\,N_i)_{i\leq p}}{}'\;\frac{q!}{q_1!\cdots q_p!\;r_1!\cdots r_p!}\;D^p\mathcal{G}_t\big(K_1,\ldots,K_p\big)\Big|_{\mathcal{Z}^*}\;\cdot\;\mathbb{E}\big[\,\bar{\mathcal{E}}^{\,r_1+\cdots+r_p}\,\big]$$

*over the assignments distributing the slots among the row's argument paths, $\sum_i(q_i+r_i)=q$ and $N_1\sqcup\cdots\sqcup N_p=\{s_1,\ldots,s_n\}$: the argument $K_i$ collects the derivatives landing on the $i$-th path, a kernel carrying $q_i$ risk slots, $r_i$ slots pinned at the row's own integrated innovation, possible only under an expectation and sitting at the one-period-ahead path's impact positions, $\widehat{X}^{(q_i)}_{0\cdots 0\;\;t+1-s,\;\,s\in N_i}$ with $r_i$ zeros, and its realized slots at the dates $N_i$. The moment is a single expectation because every integrated innovation of row $t$ is the same $\bar{\mathcal{E}}_{t+1}$, the one-period leads of Section 1 allowing no deeper reach (with a vector innovation it is the joint moment of the components left behind); the primed sum excludes the top assignment, $p=1$ with $(q_1,r_1)=(q,0)$, whose term carries the unknown kernel and sits on the left inside $\tilde{J}_X$; and every remaining argument has either strictly lower total order $q_i+r_i+|N_i|<q+n$ or, at equal total order, an integrated slot $r_i\geq 1$ and hence strictly fewer risk slots, so the ladder is triangular and consumes the lower kernels only through their impact slots;*

3. *the bottom risk rung, $q=2$ and $n=0$, has a source constant in the date up to the correction terms; its path is a transition, and the limit solves*

$$j_X(1)\;\widehat{X}^{(2)}_{\infty} \;=\; -\,\Xi^{(2,0)}{}_{\infty},$$

*the first-order symbol at $z=1$ applied to the limiting source: the shift of the risky steady state.*

*Proof. Step 1: the odd rungs vanish, by counting.* Differentiate the stacked system $q$ times in $\sigma$ and $n$ times in realized-innovation slots at the base point, and collect the equation of the $(q,n)$ rung as in Proposition 12, the display of item 2, $\tilde{J}_X\left(\widehat{X}^{(q)}_{t-s_1\,\cdots\,t-s_n}\right)_{t\geq 0}=-\,\Xi^{(q,n)}$. Track where the $q$ risk derivatives can land. Each lands either on a $\sigma$-slot of some lower kernel appearing in the source, or on one of the integrated one-period-ahead innovations $\sigma\bar{\mathcal{E}}$ inside an expectation, each such landing leaving one power of $\bar{\mathcal{E}}$ behind, to be absorbed into a moment. A generic source term is therefore a steady-state multilinear form of the stacked system contracted against kernels $K_1,\ldots,K_p$, kernel $K_i$ carrying $q_i$ risk slots, $r_i$ integrated-innovation slots, and $n_i$ realized slots, times the moment left by the integrated powers,

$$D^p\mathcal{G}\big(K_1,\ldots,K_p\big)\Big|_{\mathcal{Z}^*}\;\cdot\;\mathbb{E}\big[\,\bar{\mathcal{E}}^{\,r_1}\;\cdots\;\bar{\mathcal{E}}^{\,r_p}\,\big],$$

the innovation powers sitting at the dates of the expectations that integrated them, and the chain rule conserves the counts:

$$q\;=\;\sum_i q_i\;+\;\sum_i r_i,\qquad n\;=\;\sum_i n_i.$$

Rank the kernels lexicographically, by total order $q+n$ first and by $q$ within a total order. Every kernel in a source is of strictly lower rank: with two or more factors each has total order below $q+n$ outright, and a single factor of the same total order must have been reached through at least one integrated innovation, $r_1\geq 1$, hence carries $q_1=q-r_1<q$. Now let $q$ be odd and run the induction along the ranking. By the counting identity, either $\sum_i r_i$ is odd, and the moment factor vanishes, the innovations being independent across dates so that their joint moments factorize date by date with at least one factor of odd order, zero by hypothesis; or some $q_i$ is odd, and that kernel is zero by the induction. Every source term dies, the

rung equation is $\tilde{J}_X\left(\widehat{X}^{(q)}_{t-s_1\cdots t-s_n}\right)_{t\geq 0}=0$, linear and homogeneous, and the assumed invertibility gives zero. The base case, $(q,n)=(1,0)$, is Step 2 of Proposition 11, and at the representative-agent level the argument at orders one and three is eqs. (20)–(22) and (46) of Lan and Meyer-Gohde [2013] verbatim, the technique that of Theorem 1 of Schmitt-Grohé and Uribe [2004].

*Step 2: what a $\sigma$-derivative sees.* The risk scale enters the equilibrium conditions in one place only, the conditional expectations, so the whole computation reduces to differentiating one expectation carefully. Let $V_{t+1}$ stand for any next-dated argument of a date-$t$ conditional expectation $\mathbb{E}_t\left[\Phi(V_{t+1})\right]$ of the model. As a function of the risk scale, $V_{t+1}$ depends on $\sigma$ through two distinct channels, and separating them is the heart of the step: through the *anticipated* scale of all innovations beyond $t+1$, which its own expectations integrate, and through its *realized* innovation $\mathcal{E}_{t+1}=\sigma\,\bar{\mathcal{E}}_{t+1}$, the one the date-$t$ expectation integrates, and by the one-period leads of Section 1 the only one. Write both channels explicitly, $V_{t+1}=V(\sigma,\ \sigma\bar{\mathcal{E}}_{t+1};\ \cdot)$, and set $\psi(\sigma)\equiv\mathbb{E}_t\left[\Phi\big(V(\sigma,\sigma\bar{\mathcal{E}}_{t+1})\big)\right]$. The chain rule gives the first derivative through both channels,

$$\psi'(\sigma)\ =\ \mathbb{E}_t\Big[\ \Phi'\big(V\big)\ \big(V_\sigma\ +\ V_e\ \bar{\mathcal{E}}_{t+1}\big)\ \Big],$$

with $V_e$ the derivative in the realized-innovation slot; and differentiating once more, product rule on the two factors,

$$\psi''(\sigma)\ =\ \mathbb{E}_t\Big[\ \Phi''\big(V\big)\ \big(V_\sigma+V_e\,\bar{\mathcal{E}}_{t+1}\big)^2\ +\ \Phi'\big(V\big)\ \big(V_{\sigma\sigma}\ +\ 2\,V_{\sigma e}\,\bar{\mathcal{E}}_{t+1}\ +\ V_{ee}\,\bar{\mathcal{E}}_{t+1}{}^2\big)\ \Big].$$

Now evaluate at the base point. Step 1 clears two entries: $V_\sigma=0$ and $V_{\sigma e}=0$, the $(1,0)$ and $(1,1)$ rungs of the next-dated object, both odd in the risk scale. The moments clear the rest: $\mathbb{E}\,\bar{\mathcal{E}}=0$ kills the surviving cross terms and $\mathbb{E}\,\bar{\mathcal{E}}^2=1$ normalizes the squares, leaving

$$\psi''(0)\ =\ \Phi'\left(V_{\sigma\sigma}\ +\ V_{ee}\right)\ +\ \Phi''\,V_e^2,$$

every function now evaluated at the steady state. Read the three survivors. $V_e$ and $V_{ee}$ are the first and second derivatives of the next-dated object with respect to its own innovation, its *impact* loading and *impact* diagonal kernel, order-one and order-two objects already computed: they are the source, and they are read only at the one-period-ahead slot, because the expectation integrated only one innovation. $V_{\sigma\sigma}$ is the unknown itself, the risk kernel of the next-dated object, and its coefficient $\Phi'$ at the steady state is precisely the entry with which the lead block of the stacked operator weighs next period's path, the $Q_+$ column of Corollary 2: the term joins the left-hand side inside $\tilde{J}_X$, and no new operator appears. Higher even $\sigma$-derivatives iterate the same computation with higher moments of the unit innovation replacing $\mathbb{E}\,\bar{\mathcal{E}}^2$; the slots never deepen, because no expectation of the model reaches further than one period. The mixed rungs, $n\geq 1$, differentiate additionally in realized innovations and inherit the closure of Proposition 12 verbatim, which is item 2. In the running illustration the whole step is five lines: the one condition gives $\psi(\sigma)=\beta\,\mathbb{E}_t[\varphi(X_{t+1})]$, so with $x_0$ the impact loading and $\kappa_0\equiv\widehat{X}^{(0)}_{00}$ the impact diagonal kernel,

$$\widehat{X}^{(2)}_t\ =\ \beta\left[\varphi'\left(\widehat{X}^{(2)}_{t+1}+\kappa_0\right)+\varphi''\,x_0^2\right]\qquad\Longrightarrow\qquad\left(1-\beta\varphi'\right)\widehat{X}^{(2)}_\infty\ =\ \beta\left[\,\varphi'\,\kappa_0+\varphi''\,x_0^2\,\right]$$

on the constant tail: the risk kernel solves the same one-equation operator, sourced by the impact slices alone; nothing at deeper lags enters the source, the future's variance reaching $\widehat{X}^{(2)}$ through the fixed point on the left, not through the sum on the right.

*Step 3: the bottom rung.* At $q=2$, $n=0$ the source's ingredients, the steady-state forms of the stacked system and the impact slices $V_e, V_{ee}$, carry no date, so away from date 0 the source repeats one constant column. What does the operator do to a constant path? Each interior row convolves the constant against the symbol coefficients,

$$\Big(\mathsf{T}(j_X)\ (c,c,c,\dots)\Big)_t\ =\ \sum_v j_{X,v}\ c\ =\ j_X(1)\ c,$$

the symbol summed over all its coefficients, which is its value at $z=1$; the correction terms perturb the early rows and decay in the date at the rate $\lambda^t$ (Proposition 12, item 3). The path $\widehat{X}^{(2)}_t$ therefore converges as $t$ grows, and its limit solves the displayed equation; $j_X(1)$ is invertible because $z=1$ lies on the unit circle, where $\det j_X\neq 0$ by hypothesis. The transition itself is the work of the correction terms: the economy starts at the deterministic cross-section, and the corrections carry it to the risky rest point.

In the running illustration $j_X(1) = 1 - \beta\varphi'$, and the risky-steady-state shift reads as economics: the impact-variance source divided by one minus the discounted marginal feedback, a long-run multiplier applied to the standing risk. □

Three stationary objects emerge, and deserve distinct names. The deterministic steady state $\bar{X}^*$ anticipates no risk and receives none. The risky steady state, $\bar{X}^* + \frac{1}{2}\sigma^2 \widehat{X}^{(2)}_\infty$, the term of Coeurdacier et al. [2011], anticipates risk that never materializes: the rest point of item 3. The ergodic mean adds the risk that does materialize, the expectation of the quadratic terms,

$$\bar{X}^* \;+\; \tfrac{1}{2}\sigma^2\Big[\; \widehat{X}^{(2)}_\infty \;+\; \sum_{j\geq 0} \widehat{X}^{(0)}_{jj} \;\Big],$$

the long-run average of the simulated economy. All three come out of the same operators, and the expansion stays anchored at the first throughout: the risky steady state is a computed shift, not a new expansion point. One rung up, at total order three, the ladder corrects the impulse response itself: the kernels $\widehat{X}^{(2)}_j$, solving the same operator with a source built from the impact slices of the order-two objects, measure how standing risk bends the response to a realized innovation. At total order four the pattern closes on itself: the pure rung $\widehat{X}^{(4)}$ consumes the fourth moment of the unit innovation against the impact loadings, the rungs $\widehat{X}^{(2)}_{j_1 j_2}$ and $\widehat{X}^{(2)}_j$ already solved, and the quartic impact kernels; the ladder is triangular in the total order, each rung needing only rungs already computed. The moments that enter are the Gaussian ones, and their bookkeeping has a standard name: by Isserlis's theorem, often called Wick's, every odd moment of a centered Gaussian vanishes and every even moment counts the ways of pairing its slots, $\mathbb{E}\,\bar{\mathcal{E}}^4 = 3$ being the three pairings of four slots; a non-Gaussian innovation with vanishing odd moments changes these weights but not the structure. The bottom rung also admits a fixed-point implementation with no new machinery: perturb every conditional expectation of the steady-state routine of Section 3 by its impact variance correction, $\frac{1}{2}\sigma^2 V_{ee}$, and re-solve the steady state; the fixed point of the modified routine is the risky steady state up to $O(\sigma^4)$, the iteration implementing the inversion of item 3.

**The algorithm.** Collecting the pieces (Figure 3), the order-$k$ pass reads, with the kernels of orders below $k$ in hand:

1. *Assemble the heterogeneous right-hand side*, horizon tuple by horizon tuple: the composite slices $w_{s_B}$ and the bracket $R_{s_1\cdots s_k}$ of the display after the order-three ledger,

$$-\sum_{\substack{\pi\vdash\{1,\dots,k\}\\ |\pi|\geq 2}} D^{|\pi|}\mathcal{F}\big(w_{s_{B_1}},\dots,w_{s_{B_{|\pi|}}}\big) \;-\; \mathcal{F}_{x^+}\,\mathbb{E}\big[\,R_{s_1\cdots s_k}\,\big|\,\theta,a\big]\,:$$

   the multilinear derivatives of $\mathcal{F}$ are the preprocessor's arrays, contracted against the lower-order slices and symmetrized over the shock arguments, at order two by the symmetrizer $N = \frac{1}{2}(I+K)$ built on the commutation matrix $K$ of Magnus and Neudecker [2019, Ch. 3, §7, eq. (20) and Thm. 3.11], at higher orders by averaging the commutation matrices over the permutations of the slots; on the binding boundary $\Gamma$, the layer weights follow by Lemma 10 from the jumps $[\partial_{\mathbf{n}}\times_s], [\partial^2_{\mathbf{n}}\bar{x}^*], \dots$ of the lower-order objects and the normal speeds $v^{(j)}_\Gamma$, $j<k$, of Corollary 3, each solving against the same leading normal jump $[\partial_{\mathbf{n}}\,\bar{x}^*]$. The same jumps also ride the *seams*, the pre-images of $\Gamma$ under the steady-state policy of Section 5.2: wherever a lower-order slice enters through $\mathbb{E}[\,\cdot\mid\theta,a]$ its state derivatives are read at the landing point $\bar{a}^*(\theta,a)$, singular on $\{(\theta,a):\bar{a}^*(\theta,a)\in\Gamma\}$ and not on $\Gamma$ itself, and the layer carried there is the interface weight of Lemma 10 composed through the chain rule of Corollary 1 at $\bar{a}^*$; a source assembled on $\Gamma$ alone would silently drop these seam layers.

2. *Solve backward* in the horizon tuple, from zero at large horizons: one linear solve per node and tuple in the same factored operator that served the first order,

$$J \;=\; F_x + F_{x^+}\,\mathbb{E}[D_a\bar{x}^*\,|\,\theta,a]\,p$$

   (Proposition 1), applied to the right-hand side of step 1 exactly as Lemma 1 ran the first order, returning the order-$k$ policy kernels ${\times_{s_1\cdots s_k}}^{(k)}$.

3. *Deposit and propagate*: push the order-$k$ policy kernels into the cross-section by the order-$m$ law of motion,

$$\hat{\Omega}_{t+1}{}^{(m)} \;=\; \Lambda\,\hat{\Omega}_t{}^{(m)} \;+\; \sum_{\emptyset\neq S\subseteq\{1,\dots,m\}}\;\sum_{\pi\vdash S}\Big(D^{|\pi|}_{\bar{a}^*}\Lambda\Big)\big(\hat{a}_t{}^{(B_1)},\dots,\hat{a}_t{}^{(B_{|\pi|})}\big)\;\hat{\Omega}_t{}^{(S^c)},$$

at $m=2$ the four-term display of the distribution paragraph, through the displacements $\mathsf{M},\mathsf{M}_2,\dots$ on the grade $V_k$ the order requires; the layers deposit through the traces of Theorem 4.

4. *Close in general equilibrium*: assemble the source by the partition display of Proposition 12, item 1,

$$\Xi^{(k)}\big(\cdot\,;\,s_1,\dots,s_k\big) \;=\; \sum_{\substack{\pi\vdash\{1,\dots,k\}\\ |\pi|\geq 2}} D^{|\pi|}\mathcal{G}\Big(\hat{h}_{s_{B_1}},\;\dots,\;\hat{h}_{s_{B_{|\pi|}}}\Big),$$

and solve $\tilde{J}_X\,\big(\widehat{X}^{(0)}_{t-s_1\,\cdots\,t-s_k}\big)_{t\geq 0} = -\,\Xi^{(k)}\big(\cdot\,;\,s_1,\dots,s_k\big)$ column by column with the stored factorization $j_X = j_-\,j_+$, one anti-causal pass $\mathsf{T}(j_+^{-1})$ and one causal pass $\mathsf{T}(j_-^{-1})$, plus the four correction terms of item 3 of the same proposition.

5. *Read out*: pair any aggregate kernel against the cross-section, $\langle\bar{\Omega}^*,\cdot\rangle$; at even orders add the risk rungs, $\tilde{J}_X\,\big(\widehat{X}^{(q)}_{t-s_1\,\cdots\,t-s_n}\big)_{t\geq 0} = -\,\Xi^{(q,n)}$ (Proposition 13); when reporting impulse responses, recenter the diagonal kernels by the standing variance, the order-two diagonal entering as $\frac{1}{2}\,\widehat{X}^{(0)}_{tt}\,\big(\mathcal{E}_0{}^2-\sigma^2\big)$, so that the response measures news relative to standing risk.

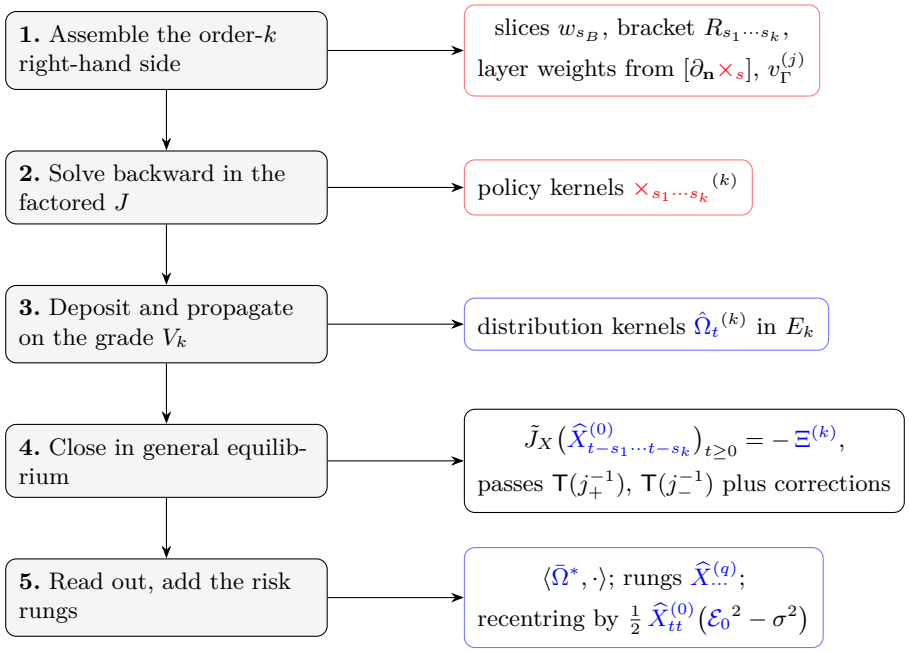


Figure 3: The order-$k$ cascade, mirroring the first-order cascade of Figure 2: the same steps one order up, every solve reusing the factored objects of the first order, $J$ at the household step and $j_X = j_-\,j_+$ at the general-equilibrium step, the sources assembled from the kernels of orders below $k$.

Because the shock kernels are perfect-foresight objects, every step admits a derivative-free check: nonlinear perfect-foresight transitions on the stencil $\{-h,+h\}^k$ of evaluation points, mixed differences across transitions relaunched from stored intermediate cross-sections, closed in general equilibrium by the same first-order solve. The check has two blind spots the analytic route does not. A symmetric difference straddling the binding boundary $\Gamma$ halves the jump, an $O(1)$ pointwise artifact on the binding set and $O(h)$ in aggregates, so the stencil must respect the regime; and order $k$ needs the stencil's $2^k$ vertices against a division by $h^k$. It remains what it is, a check; the recursion is the production route.

**What the element must carry.** The order of the expansion and the degree of the finite element are one lever seen twice. The pointwise $C_{1,1}$ weight, the third of the second-derivative displays after Corollary 3, needs a second state derivative of the basis to land on, and the $P^1$ hat has none, $\mathcal{P}_{aa} = 0$: the $C_{1,1}$ that Section 4.5 records as missed is the order-two shadow of a basis one degree too low, however fine the grid [Bhandari et al., 2023, App. C.1]. And the pairings of order $k$ read test functions through $k$ derivatives, so a basis reproducing polynomials of degree $d$ carries the expansion to order $d$ and no further: a biquadratic or Hermite element restores order two, the interpolant's exactness on its shape space doing at degree two what Section 4.5 recorded at degree one [Brenner and Scott, 2008, (3.1.3), (3.3.1), the exactness their Prop. 3.3.7]. The layers add one demand of their own, element nodes or element faces aligned with the binding boundary $\Gamma$, so that the traces the ledger $\mathsf{L}_\Gamma$ pairs against are grid objects.

**Truncation and consumption.** Truncating the expansion at order $K$ and never feeding the discarded orders back through a state is the pruning of the representative-agent literature, in its moving-average reading: the kernels through $K$ are kept, the rest dropped, and no invariant-manifold construction is asked for [Lan and Meyer-Gohde, 2013, Andreasen et al., 2018]. What remains is directly consumable. Simulation is one polynomial evaluation per date. Unconditional moments are contractions of the kernels against the innovation moments: the mean shift is the ergodic display above, and the covariances add the kernel-squared Wick contractions. Shock recovery from observed aggregates is least squares in the polynomial map. What Dynare 7 would report next to `dr.G` is the kernel family itself, one array per order and rung, every entry produced by the two operators the first-order solve already factored.

# Conclusion

What Dynare 7 implements is the first-order pipeline: a `.mod` syntax for the two-block class of Section 1, two steady-state commands with their time-iteration, forward-iteration, and calibration algorithms, and a first-order solver whose product is the sequence-space Jacobian `dr.G`, consumed by `heterogeneity_simulate` for impulse responses and simulated paths. The mathematical content of that pipeline is the pair of operators of Section 4, the backward operator $\mathfrak{r}$ of the heterogeneous block and the transfer operator $\Lambda$ of the cross-section, differentiated at the steady state and assembled in the sequence space; the finite-element reading of Section 4.5 says exactly in which norm the histogram pipeline converges to those objects, at rate $O(h)$ in the pairings the aggregates consume. Read at this level, the state-space recursion of Bhandari et al. [2023] and the sequence-space Jacobian of Auclert et al. [2021] are one Fréchet derivative in two coordinate systems, the reading this document takes throughout.

The framework pays for itself beyond the solver, on arrays already formed. The determinacy and existence test of Section 5.1 sums the stored arrays $\mathcal{F}_{t,s}$ along their diagonals, evaluates the symbol $j_X$ on a frequency grid by one fast Fourier transform, and reads the winding number of $\det j_X$; it carries the criterion of Auclert et al. [2025] to the continuous model, with a finite-sample guarantee that the discretized verdict is that model's, at a cost negligible next to the solve it certifies, and the truncated solve alone cannot provide it. The order-$k$ cascade of Section 5.2 carries the second-order heterogeneous construction of Bhandari et al. [2023] to arbitrary order, solving every higher order, shock kernels and risk corrections alike, in the one operator the first order already factored, with a right-hand side assembled from orders below; the second derivatives the order-two step consumes are already emitted by the preprocessor. And the frequency-domain algebra behind both, the symbol as a transfer function evaluated at products of its arguments, prices moments and spectral objects without simulation. None of the three is in the code yet; each is a bounded step from what is.

The assumption set deserves its one-paragraph summary. At first order: strong regularity of the unfolded household problem in the sense of Robinson [1980], strict complementarity off a $\bar{\Omega}^*$-null binding boundary $\Gamma$, the four clauses of Assumption 1 on the stationary distribution and the test space, finitely many singular carriers originating at the bounds among them, and the two geometric-stability clauses of Assumption 2, discounting and mixing, the mixing bound read as an essential supremum over the seam-free set, all stated rather than derived from primitives. For determinacy and for the truncation Dynare performs: a zero winding number, and, for the finite-section corner to return the infinite inverse behind `dr.G`, invertibility of the reflected operator $z \mapsto j_X(1/z)$ as well; both hold generically, but in operator space, whose translation to model primitives Auclert et al. [2025] leave open. The two-pass inverse and the cheap higher-order cascade ask strictly more than $w = 0$, a *canonical* factorization with every partial index zero (Proposition 9), $w = 0$ being necessary but not sufficient. At the higher orders: clauses 1

and 3 of Assumption 3, virtual-saving non-degeneracy and the interior-density trace, continuous across the binding boundary $\Gamma_b$, the shadow-price positivity of clause 2 now derived (Lemma 7), with their uniformity at intersections; the uniform-in-$q$ strong regularity of Lemma 9 and the transversality of the iterated pre-images of $\Gamma$; existence and $k$-fold continuous differentiability of the stochastic solution near $\sigma = 0$ (Proposition 11, the deep hypothesis it shares with Schmitt-Grohé and Uribe [2004] and Lan and Meyer-Gohde [2013]); and the graded decay clauses (i)–(ii) of Appendix E that strengthen Assumption 2, whose qualitative frame is established in the text and whose leading case is derived there. Everything else in the document is proved from these.

# References


Ralph Abraham, Jerrold E. Marsden, and Tudor Ratiu. *Manifolds, Tensor Analysis, and Applications*, volume 75 of *Applied Mathematical Sciences*. Springer, New York, second edition, 1988. doi: 10.1007/978-1-4612-1029-0.

Robert A. Adams and John J. F. Fournier. *Sobolev Spaces*, volume 140 of *Pure and Applied Mathematics*. Academic Press, Amsterdam, second edition, 2003.

Stéphane Adjemian, Houtan Bastani, Michel Juillard, Frédéric Karamé, Junior Maih, Ferhat Mihoubi, Willi Mutschler, George Perendia, Johannes Pfeifer, Marco Ratto, and Sébastien Villemot. Dynare: Reference Manual Version 4. Dynare Working Papers 1, CEPREMAP, 2011.

S. Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684, 1994. doi: 10.2307/2118417.

Martin M. Andreasen, Jesús Fernández-Villaverde, and Juan F. Rubio-Ramírez. The pruned state-space system for non-linear DSGE models: Theory and empirical applications. *The Review of Economic Studies*, 85(1):1–49, 2018. doi: 10.1093/restud/rdx037.

Adrien Auclert, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408, 2021. doi: https://doi.org/10.3982/ECTA17434. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA17434.

Adrien Auclert, Evan Majic, Matthew Rognlie, and Ludwig Straub. Determinacy and large-scale solutions in the sequence space. *Journal of Political Economy: Macroeconomics*, 2025.

Anmol Bhandari, Thomas Bourany, David Evans, and Mikhail Golosov. A perturbational approach for approximating heterogeneous agent models. Working Paper 31744, National Bureau of Economic Research, September 2023. URL http://www.nber.org/papers/w31744.

Dario A. Bini, Stefano Massei, and Leonardo Robol. Quasi-Toeplitz matrix arithmetic: A MATLAB toolbox. *Numerical Algorithms*, 81:741–769, 2019.

Olivier Jean Blanchard and Charles M. Kahn. The solution of linear difference models under rational expectations. *Econometrica*, 48(5):1305–1311, 1980.

Albrecht Böttcher and Bernd Silbermann. *Introduction to Large Truncated Toeplitz Matrices*. Universitext. Springer, New York, 1999. doi: 10.1007/978-1-4612-1426-7.

Susanne C. Brenner and L. Ridgway Scott. *The Mathematical Theory of Finite Element Methods*, volume 15 of *Texts in Applied Mathematics*. Springer, New York, third edition, 2008. doi: 10.1007/978-0-387-75934-0.

James Ward Brown and Ruel V. Churchill. *Complex Variables and Applications*. McGraw-Hill Education, New York, ninth edition, 2014.

Nicolas Coeurdacier, Hélène Rey, and Pablo Winant. The risky steady state. *American Economic Review*, 101(3):398–401, 2011.

Wilbur John Coleman, II. Solving the stochastic growth model by policy-function iteration. *Journal of Business & Economic Statistics*, 8(1):27–29, 1990. doi: 10.1080/07350015.1990.10509766.

Fabrice Collard and Michel Juillard. A higher-order Taylor expansion approach to simulation of stochastic forward-looking models with an application to a nonlinear Phillips curve model. *Computational Economics*, 17(2–3):125–139, 2001a. doi: 10.1023/A:1011624124377.

Fabrice Collard and Michel Juillard. Accuracy of stochastic perturbation methods: The case of asset pricing models. *Journal of Economic Dynamics and Control*, 25(6–7):979–999, 2001b. doi: 10.1016/S0165-1889(00)00064-6.

Ruth F. Curtain and Hans Zwart. *An Introduction to Infinite-Dimensional Linear Systems Theory*, volume 21 of *Texts in Applied Mathematics*. Springer, New York, 1995.

M. C. Delfour and J.-P. Zolésio. *Shapes and Geometries: Metrics, Analysis, Differential Calculus, and Optimization*, volume 22 of *Advances in Design and Control*. Society for Industrial and Applied Mathematics, Philadelphia, second edition, 2011. doi: 10.1137/1.9780898719826.

Lawrence C. Evans and Ronald F. Gariepy. *Measure Theory and Fine Properties of Functions*. Textbooks in Mathematics. CRC Press, Boca Raton, second edition, 2025.

Francisco Facchinei and Jong-Shi Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems*. Springer Series in Operations Research. Springer, New York, 2003. doi: 10.1007/b97543.

Alessio Figalli. Free boundary regularity in obstacle problems, 2018. URL https://arxiv.org/abs/1807.01193. Lecture notes, Journées Équations aux Dérivées Partielles 2018; arXiv:1807.01193.

L. E. Fraenkel. Formulae for high derivatives of composite functions. *Mathematical Proceedings of the Cambridge Philosophical Society*, 83(2):159–165, 1978. doi: 10.1017/S0305004100054402.

F. G. Friedlander and M. Joshi. *Introduction to the Theory of Distributions*. Cambridge University Press, Cambridge, second edition, 1998.

I. M. Gel'fand and G. E. Shilov. *Generalized Functions, Volume 1: Properties and Operations*. Academic Press, New York, 1964.

Israel Gohberg, Seymour Goldberg, and Marinus A. Kaashoek. *Classes of Linear Operators, Vol. I*, volume 49 of *Operator Theory: Advances and Applications*. Birkhäuser, Basel, 1990.

Sébastien Gouëzel and Carlangelo Liverani. Banach spaces adapted to Anosov systems. *Ergodic Theory and Dynamical Systems*, 26(1):189–217, 2006.

Mark Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5–6):953–969, 1993. doi: 10.1016/0165-1889(93)90024-M.

Ram P. Kanwal. *Generalized Functions: Theory and Technique*. Birkhäuser, Boston, second edition, 1998.

Per Krusell and Anthony A. Smith, Jr. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, October 1998. doi: 10.1086/250034.

Hong Lan and Alexander Meyer-Gohde. Solving DSGE models with a nonlinear moving average. *Journal of Economic Dynamics and Control*, 37(12):2643–2667, 2013. doi: 10.1016/j.jedc.2013.06.005.

Andrzej Lasota and Michael C. Mackey. *Chaos, Fractals, and Noise: Stochastic Aspects of Dynamics*, volume 97 of *Applied Mathematical Sciences*. Springer, New York, second edition, 1994.

Francesco Maggi. *Sets of Finite Perimeter and Geometric Variational Problems: An Introduction to Geometric Measure Theory*, volume 135 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2012.

Jan R. Magnus and Heinz Neudecker. *Matrix Differential Calculus with Applications in Statistics and Econometrics*. Wiley, Chichester, third edition, 2019. doi: 10.1002/9781119541219.

Alexander Meyer-Gohde. Solving and analyzing DSGE models in the frequency domain. *Journal of Economic Dynamics and Control*, 185:105281, 2026.

Alexei Onatski. Winding number criterion for existence and uniqueness of equilibrium in linear rational expectations models. *Journal of Economic Dynamics and Control*, 30(2):323–345, 2006.

Arshak Petrosyan, Henrik Shahgholian, and Nina Uraltseva. *Regularity of Free Boundaries in Obstacle-Type Problems*, volume 136 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, 2012.

Michael Reiter. Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, 33(3):649–665, 2009. doi: 10.1016/j.jedc.2008.08.010.

Stephen M. Robinson. Strongly regular generalized equations. *Mathematics of Operations Research*, 5 (1):43–62, 1980. doi: 10.1287/moor.5.1.43.

K. Geert Rouwenhorst. Asset pricing implications of equilibrium business cycle models. In Thomas F. Cooley, editor, *Frontiers of Business Cycle Research*, chapter 10, pages 294–330. Princeton University Press, 1995.

Wilson J. Rugh. *Nonlinear System Theory: The Volterra/Wiener Approach*. Johns Hopkins University Press, Baltimore, 1981.

Stephanie Schmitt-Grohé and Martín Uribe. Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control*, 28(4): 755–775, 2004.

Haruyoshi Tanaka. General asymptotic perturbation theory in transfer operators, 2022. URL https://arxiv.org/abs/2205.12561. arXiv:2205.12561.

Eric R. Young. 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, 2010. doi: 10.1016/j.jedc.2008.11.010.

Eberhard Zeidler. *Nonlinear Functional Analysis and its Applications. II/A: Linear Monotone Operators*. Springer, New York, 1990.

# A User-facing helper functions

## A.1 The `rouwenhorst` helper

For the common case where a heterogeneous exogenous variable $e$ follows an AR(1) process

$$e_{i,t} = \rho\, e_{i,t-1} + \epsilon^{e}_{i,t}, \qquad \epsilon^{e}_{i,t} \sim \mathcal{N}(0, \sigma^2),$$

producing the (`initial_guess.shocks.grids.e`, `initial_guess.shocks.Pi.e`) pair documented in Subsection 3.1 is routine enough that Dynare ships a helper. Once Dynare is on the MATLAB path (so that `dynare/matlab` has been added via `addpath`), the function

```
[y, p_vec, P_mat] = rouwenhorst(rho, sigma, N, tol, max_iter)
```

becomes callable from the `verbatim` block of a `.mod` file. The arguments are the AR(1) persistence $\rho$, the standard deviation $\sigma$ of the innovation, the desired number of nodes $N$, a tolerance `tol` on the stationary-distribution iteration, and a maximum iteration count `max_iter` for that same step. The helper constructs an $N$-point Markov chain approximating the AR(1) process following Rouwenhorst [1995]. The output `y` is the grid returned in levels (exponentiated and renormalized so that $\sum_i p_i\, y_i = 1$), `p_vec` is the stationary distribution of the chain, and `P_mat` is the transition matrix; the user wires `y` and `P_mat` directly into `initial_guess.shocks.grids.e` and `initial_guess.shocks.Pi.e`. The helper is implemented as a Fortran MEX function under `mex/sources/heterogeneity/rouwenhorst.f08` and compiled at Dynare build time; it is not part of base MATLAB.

# B Building the `initial_guess` structure in the `.mod` file

In the running Krusell–Smith economy, the construction of the initial-guess structure is the following.

```
delta = 0.025;
eis = 1;
L = 1;
alpha = 0.11;
beta = 0.98;
rho_Z = 0.80;

verbatim;
r = 0.01;
rk = r + delta;
end;

% Normalize so that Y = 1
Z_ss = (rk / alpha) ^ alpha;

verbatim;
K = (alpha * Z_ss / rk) ^ (1 / (1 - alpha));
Y = Z_ss * K ^ alpha;
w = (1 - alpha) * Z_ss * (alpha * Z_ss / rk) ^ (alpha / (1 - alpha));

initial_guess = struct;
initial_guess.agg.w = w;
initial_guess.agg.r = r;
initial_guess.agg.Y = Y;
initial_guess.agg.K = K;
initial_guess.agg.Z = 0;

initial_guess.free_parameters.beta.initial_guess = 0.98;
initial_guess.free_parameters.beta.lower_bound = 0.97;
initial_guess.free_parameters.beta.upper_bound = 0.99;

rho_e = 0.966;
sig_e = 0.5;
[grid_e, ~, Pi_e] = rouwenhorst(rho_e, sig_e, 3, 1e-12, 1e5);
initial_guess.shocks.grids.e = grid_e;
initial_guess.shocks.Pi.e = Pi_e;

grid_a = logspace(log10(0.25), log10(200.25), 30) - 0.25;
initial_guess.pol.grids.a = grid_a;
coh = (1 + r) * grid_a + w * grid_e;
```

```
c = 0.1 * coh;
a = coh - c;
initial_guess.pol.values.c = c;
initial_guess.pol.values.a = a;
initial_guess.pol.order = {'e', 'a'};
end;
```

The listing maps onto the discrete objects of Subsection 3.1 field for field. The five `initial_guess.agg.X` assignments give starting values for the aggregate scalars $\bar{X}^*_{[Y]}$, $\bar{X}^*_{[r]}$, $\bar{X}^*_{[w]}$, $\bar{X}^*_{[K]}$, $\bar{X}^*_{[Z]}$ (the TFP deviation $Z_t$ is zero at the steady state, so its entry is exactly 0 rather than a guess); the `free_parameters.beta` block declares $\beta$ as a free parameter to be calibrated, with an initial value and bounds, and its presence switches the command from residual-only mode to calibration mode (Subsection 3.2).

The `rouwenhorst` helper, documented in Appendix A.1, produces the discrete shock grid $\boldsymbol{\theta}_{[e]}$ and the transition matrix $\boldsymbol{\pi}_{[e]}$ from the AR(1) persistence and innovation standard deviation supplied above it; `logspace` produces a log-spaced predetermined-state policy grid $\boldsymbol{a}^p_{[a]}$ of 30 nodes over $[0, 200]$ after a shift; and the cash-on-hand initialization `coh = (1 + r) * grid_a + w * grid_e`, `c = 0.1 * coh`, `a = coh - c` populates the policy arrays $\boldsymbol{x}_{[c]}$ and $\boldsymbol{x}_{[a]}$ at every $(\boldsymbol{\theta}_{[e]}, \boldsymbol{a}^p_{[a]})$ node by broadcasting the column vector `grid_e` against the row vector `grid_a`.

The cell array `pol.order = {'e','a'}` fixes the dimension ordering with the shock axis first and the predetermined-state axis last. The listing intentionally omits `d.hist` (ignored here) and `d.grids` (reused from `pol.grids` when absent), and supplies only the five aggregates that the user has declared explicitly in the `.mod` file.

The listing above mixes inline statements with two `verbatim;` blocks. The six parameter assignments at the top (`delta = 0.025;` through `rho_Z = 0.80;`) sit inline; a first `verbatim;` block contains `r = 0.01; rk = r + delta;`; the single line `Z_ss = (rk/alpha)^alpha;` sits inline between the two blocks; and a second, longer `verbatim;` block builds the `initial_guess` structure. Where each region sits is dictated by what the Dynare preprocessor does to each statement form.

The preprocessor reads the `.mod` file and emits a MATLAB driver file `<modfile>/driver.m`. At the top level, statements the parser recognizes (declarations, model blocks, parameter assignments, command invocations) are transpiled into specific MATLAB calls; anything else is passed through as a literal MATLAB statement; a `verbatim;` block forces this pass-through for code containing tokens the parser would otherwise interpret. The four statement forms appearing in the listing transpile as follows.

| **Statement on the `.mod` side** | **Outside `verbatim;`** | **Inside `verbatim;`** |
|---|---|---|
| `delta = 0.025;` (`delta` declared as `parameter`) | `M_.params(4) = 0.025;` `delta = M_.params(4);` *(slot set, MATLAB local bound)* | `delta = 0.025;` *(only MATLAB local set; `M_.params(4)` stays `NaN`)* |
| `r = 0.01;` (`r` declared as aggregate `var`) | syntax error: a bare top-level assignment to a declared variable is not an admissible `.mod` statement | `r = 0.01;` *(MATLAB local set)* |
| `rk = r + delta;` (`rk` undeclared) | `rk = r + M_.params(4);` *(passed through; parameter `delta` substituted on the RHS)* | `rk = r + delta;` *(passed through literally)* |
| `Z_ss = (rk/alpha)^alpha;` (`Z_ss`, `alpha` declared as `parameter`; `rk` undeclared) | `M_.params(7) = (rk/M_.params(2))^M_.params(2);` `Z_ss = M_.params(7);` *(slot set; RHS uses substitutions and the MATLAB local `rk`)* | `Z_ss = (rk/alpha)^alpha;` *(only MATLAB local set; `M_.params(7)` stays `NaN`)* |

Two consequences read off the table. An assignment whose left-hand side is a declared *parameter must* be inline: only the inline form triggers the `M_.params(i)` update, and an identical statement inside `verbatim;` would set only the MATLAB local, leaving Dynare's internal parameter slot at its default value. An assignment whose left-hand side is a declared *variable* (`var`, `varexo`, or `var(heterogeneity=NAME)`) *must* be inside `verbatim;`: outside, the parser would refuse it. Statements assigning to undeclared identifiers are indifferent to wrapping, and `verbatim;` only serves to group them with the surrounding code.

Applied to the KS listing, the six inline parameter assignments at the top trigger the corresponding `M_.params(i)` updates by the first rule. The first `verbatim;` block is required because `r` is a declared aggregate variable; `rk` is undeclared and travels with it for grouping and so that `rk` exists as a MATLAB local for the subsequent line. The inline `Z_ss = (rk/alpha)^alpha;` is then in turn required to be inline because `Z_ss` is a parameter. The second `verbatim;` block is required for the assignments to `K`, `Y`, `w` (aggregate variables) and `c`, `a` (heterogeneous variables); everything else inside it (the `struct` constructor, the `rouwenhorst` call, `logspace`, the cash-on-hand initialization, every `initial_guess.*` field assignment) is undeclared and would compile inline just as well, but is wrapped in the same block so that the entire initial-guess construction lives in one place and reads as ordinary MATLAB code.

The two-block structure of the KS listing is therefore dictated by the intermediate parameter assignment to `Z_ss`. When no such parameter dependency on scaffold scalars arises, a single `verbatim;` block suffices (the pattern of `hank_one_asset_steady_state.mod`); when several scaffold scalars feed several parameter assignments, the listing alternates a `verbatim;` block, a run of inline parameter assignments, and a second `verbatim;` block (the pattern of `hank_two_assets_steady_state.mod`).

This dual-namespace bookkeeping is the price of keeping the entire initial-guess construction in a single source file alongside the model definition, which is convenient for small examples. Building the structure in a standalone script and passing the resulting `.mat` file through the `filename` option avoids it entirely; a dedicated interface that lets the command read a structure-building `.m` file directly is envisioned for a future release.

# C A functional-analysis primer

The dynamics of Section 4 differentiate maps between spaces of functions: the recursive map $\bar{\mathcal{Z}}$ of Section 4.1, the household residual $\mathcal{F}$ of Section 4.2, and the transfer operator $\Lambda$ of Section 4.3. This appendix first fixes what differentiating such a map means, the symbol $D\cdot$ the section uses throughout, and then collects the distribution-side machinery that carries the same calculus through while a positive mass of households sits exactly at the constraint, without ever differentiating the kinked policy: the mass points move, and their motion is carried by the finite dipole span $D$ of Section 4.3, not by a calculus of generalized functions. The distribution objects form a chain, each brought in because the previous step needs it. I state each term formally, with a reference to Abraham et al. [1988] or Brenner and Scott [2008] by result number, and say why it is needed before defining it. A reader after only the economics can skip this; a reader who wants to follow the proofs will find every term defined.

**Derivatives of maps between function spaces.** Each step of the dynamics linearizes a map between normed spaces of functions, and the kink the constraint puts in the policy (Section 4.2) makes the sense of "derivative" matter, so I fix two of them.

**Definition 8** (Gâteaux derivative)**.** A map $\mathcal{G}$ between normed spaces has a Gâteaux, or directional, derivative at $u$ in the direction $h$ if the limit

$$\left.\frac{d}{dt}\right|_0 \mathcal{G}(u+th)$$

exists; it probes $\mathcal{G}$ along the single ray through $h$ [Abraham et al., 1988, Def. 2.4.5].

**Definition 9** (Fréchet derivative)**.** A map $\mathcal{G}$ is Fréchet differentiable at $u$ if a single bounded linear map $D\mathcal{G}(u)$ approximates it to first order uniformly over all directions at once,

$$\|\mathcal{G}(u+h)-\mathcal{G}(u)-D\mathcal{G}(u)\,h\| \;=\; o(\|h\|) \quad \text{as } \|h\|\to 0$$

[Abraham et al., 1988, Def. 2.3.3]. This $D\cdot$ is the derivative the whole section uses, always read at a base point $u$; a subscript, as in $D_X\mathcal{G}(u)$, marks the partial Fréchet derivative in the argument $X$ alone. In the dynamics that base point is the steady state unless stated otherwise, so $D_{\mathcal{Z}}$ in Section 4.1, $D_{\bar{x}}\mathcal{F}$ in Section 4.2, and $D_{\bar{a}^*}\Lambda$ in Section 4.3 are all partial derivatives of this kind, taken at the steady state.

The Fréchet derivative is the stronger of the two: it controls the approximation in every direction at once, where the Gâteaux derivative controls it only along each ray separately. The two need not agree where a map is kinked, which is exactly the difficulty the constraint raises in Section 4.2, but they coincide as soon as the directional derivative varies continuously.

**Proposition 14** (Gâteaux equals Fréchet)**.** *If $\mathcal{G}$ is Gâteaux differentiable on a neighbourhood of $u$, its Gâteaux derivative is a bounded linear map at each point, and that map depends continuously on the base point, then $\mathcal{G}$ is continuously Fréchet differentiable there and the two derivatives coincide [Abraham et al., 1988, Cor. 2.4.10].*

**Test functions and the dual pairing.** One never observes the cross-sectional distribution $\mu$ pointwise. What the aggregate equations use are integrals against a function, a mean, a moment, a marginal-propensity-weighted sum, $\int \varphi \, d\mu(\theta, a)$. So the right notion of a distribution is not a density but the rule that eats a test function $\varphi$ and returns the number $\int \varphi \, d\mu$. That rule is a linear functional, and this one change of viewpoint is what lets a point mass at the constraint in without trouble. For the rule to be continuous I need the right test space, and building it calls for a generalized, integration-by-parts notion of derivative, which I set up first.

**Weak derivatives.** Two later steps need a derivative that survives a kink: building the test space just promised, and, in Section 4.3, differentiating the law of motion through the constrained policy, which bends at the limit so its classical derivative fails there. The escape is integration by parts: move the derivative off the rough object onto a smooth test function, where it always exists. The test functions are smooth and *compactly supported*, vanishing outside a bounded set, so they and their derivatives are controlled at the edges and the boundary terms of the integration by parts drop; the rough objects they probe need only be *locally integrable*, that is integrable on every bounded set, which is the least that makes the pairing $\int f \, \varphi \, dx$ a finite number, the integral running only over the bounded support of $\varphi$, where $f$ is integrable.

**Definition 10** (Weak derivative)**.** Let $\alpha$ be a multi-index, one non-negative integer $\alpha_i$ per state coordinate $i$, of order $|\alpha| = \sum_i \alpha_i$; the symbol $\varphi^{(\alpha)}$ denotes the classical mixed partial that differentiates $\varphi$ a total of $|\alpha|$ times, $\alpha_i$ times in coordinate $i$. A locally integrable $g$ is the weak derivative $D^\alpha_w f$ of a locally integrable $f$ when

$$\int g \, \varphi \, dx \;=\; (-1)^{|\alpha|} \int f \, \varphi^{(\alpha)} \, dx$$

for every smooth compactly supported $\varphi$ [Brenner and Scott, 2008, (1.2.4)].

**Proposition 15** (Weak extends classical)**.** *Where the classical derivative exists, the weak derivative agrees with it [Brenner and Scott, 2008, (1.2.7)].*

The example to keep in mind [Brenner and Scott, 2008, (1.2.5)] is $f(x) = 1 - |x|$, the local shape of a saving policy bending into a constraint: its weak derivative is minus the sign function, a bounded $L^\infty$ jump.

**Sobolev spaces.** By point evaluation I mean the linear functional $\varphi \mapsto \varphi(x)$ that returns a function's value at a fixed point $x$ of the domain; it is *bounded* on a normed space of functions when $|\varphi(x)| \le C \, \|\varphi\|$ for some constant $C$ and all $\varphi$, so that the value at a point is controlled by the norm. It is the same object as the Dirac $\delta_x$ named in the dual paragraph below, read from the side of the function it acts on. The functional viewpoint needs a test space on which this point evaluation is bounded, since that is what will let a point mass count as a functional. I cut that space out by counting weak derivatives and the integrability in which they are measured: ask for one derivative in a high enough power and evaluating at a point becomes a bounded operation.

**Definition 11** (Sobolev space)**.** For a non-negative integer $k$ and an exponent $1 \le p < \infty$, the Sobolev space $W^{k,p}$ on a given domain consists of the functions whose weak derivatives up to order $k$ lie in $L^p$ [Brenner and Scott, 2008, (1.3.1)]. The square-integrable case is the Hilbert space $H^k = W^{k,2}$. The dynamics need only the first-order spaces $W^{1,p}$ on the continuous predetermined coordinates $a \in \mathcal{A}$.

**Theorem 1** (Sobolev embedding)**.** *Let $d_a = \dim \mathcal{A}$ be the number of continuous predetermined coordinates; the idiosyncratic state $\theta$ is a finite Markov index, on which no derivative is ever taken, so it spends no part of this budget. If $p > d_a$, then every $u \in W^{1,p}(\mathcal{A})$ has a continuous representative, and*

$$\|u\|_{L^\infty} \;\le\; C \, \|u\|_{W^{1,p}};$$

*in particular $W^{1,p} \hookrightarrow C^0$ and point evaluation is a bounded linear functional [Brenner and Scott, 2008, Thm. 1.4.6].*

That continuous representative is what the construction rests on. An element of $W^{1,p}$, like any element of an $L^p$-based space, is not one function but a class of functions that agree outside a set of measure zero. A single point is such a set, so the value $u(\theta, a)$ is on its face undefined. The embedding repairs this: when $p > d_a$, the class of $u$ contains a continuous function, and only one, since two continuous functions that agree almost everywhere agree everywhere. Identifying $u$ with that continuous representative makes $u(\theta, a)$ unambiguous, and the displayed bound makes it move continuously with $u$. This is exactly what lets the Dirac below act as a bounded functional. I fix the test space $V = \ell^2(\Theta) \otimes W^{1,p}(\mathcal{A})$ with $p > d_a$: one weak derivative, in integrability high enough to embed in $C^0$. Why this order and integrability, and why no more economical space serves, is the subject of the hierarchy of test spaces at the end of this appendix.

**What goes wrong without the derivative.** In finite dimension every linear map $A$ is bounded, so $x \mapsto Ax$ is automatically continuous; in infinite dimension this can fail, and the instance that fails is point evaluation $u \mapsto u(x_0)$, whose breakdown fixes intuition for everything below. Take on the line the tents $f_n(x) = \max(0,\, 1 - n|x|)$, of height one on the shrinking base $[-1/n, 1/n]$. Each has $f_n(0) = 1$, yet

$$\|f_n\|_{L^2}^2 \;=\; \int_{-1/n}^{1/n} \big(1 - n|x|\big)^2 \, dx \;=\; \frac{2}{3n} \;\longrightarrow\; 0,$$

so $f_n(0)/\|f_n\|_{L^2} = \sqrt{3n/2} \to \infty$: point evaluation is an *unbounded* functional on $L^2$, no element of its dual. One derivative rules the spikes out, since $\|f_n'\|_{L^2}^2 = 2n \to \infty$ keeps them out of any bounded set of $H^1$, and in one dimension $H^1(\mathbb{R}) \hookrightarrow C^0$ restores bounded point evaluation. But one derivative buys this only up to dimension one. In two dimensions it does not: $\log\log(1/|x|)$, cut off away from the origin, lies in $H^1(\mathbb{R}^2)$ yet is unbounded there, so point evaluation is again unbounded and a point mass is no element of $(H^1(\mathbb{R}^2))'$. Theorem 1 is the precise threshold $p > d_a$ separating the two regimes, and the hierarchy at the end of this appendix works out what it costs.

**The dual, and the Dirac.** With the test space $V$ fixed, the distributions are exactly the rules that act on it, the elements of its dual. I need that dual roomy enough to hold three things at once: the steady-state distribution, the mass point the constraint piles at the limit, and the density part of every response; the motion of the mass points themselves lives one rung further down, and the paragraph on what the dual cannot hold, below, explains why no admissible dual reaches it.

**Definition 12** (Dual space and the pairing)**.** The dual space $V'$ is the space of bounded linear functionals $L : V \to \mathbb{R}$ [Brenner and Scott, 2008, (1.7.1)–(1.7.2)], acting through the pairing $\langle \mu, \varphi \rangle = \int \varphi \, d\mu(\theta, a)$. A density $g$ acts by $\varphi \mapsto \int g\,\varphi \, dx$, so functions embed in $V'$; the dual of a positive-order Sobolev space is a negative-order one, $V' = \ell^2(\Theta) \otimes W^{-1,p'}(\mathcal{A})$ with $1/p + 1/p' = 1$ [Brenner and Scott, 2008, (1.7.5)]. Throughout, this angle bracket is the $V'$–$V$ pairing and not a separate inner product: when $\mu$ has a density $g$ it reduces to the $L^2$ inner product $\int g\,\varphi \, dx$, and when $\mu$ is a genuine distribution such as the constraint's Dirac it is the extension $\int \varphi \, d\mu$ that survives where the inner product does not, the two agreeing wherever both are defined.

**Proposition 16** (The Dirac is admissible)**.** *When $p > d_a$, the Dirac $\delta_{(\theta,a)} : \varphi \mapsto \varphi(\theta, a)$ is a bounded linear functional on $V$, hence an element of $V'$, though it is represented by no $\int g\,\varphi \, dx$ [Brenner and Scott, 2008, (1.7.4)]. The same holds for a corner at which up to $d_a$ bounds bind at once: such a mass is concentrated at a single point $a \in \mathcal{A}$ of the predetermined coordinates and is admitted by the same embedding.*

The constraint's point mass is therefore a legitimate element of $V'$, the room the dual had to provide.

**What the dual cannot hold.** One object stays out of reach, and it is worth seeing precisely why. An atom of the distribution that moves differentiates into it: when a point mass sits at $q$ today and a perturbation of size $h$ shifts it to $q + h\,v$, the first-order response of the distribution is the limit

$$\lim_{h \to 0} \; \frac{\delta_{q + h\,v} \;-\; \delta_q}{h},$$

two spikes side by side, mass arriving next to mass leaving. Electrostatics calls this configuration a dipole, and I keep the name; Definition 3 in the main text records it. The pairing is exactly the two-spike limit

tested against $\psi$,

$$\frac{\psi(q + h\,v) \;-\; \psi(q)}{h} \;\longrightarrow\; v \cdot \nabla_a \psi(q) \qquad \text{as } h \to 0.$$

So reading a dipole asks for the *slope* of the test function at one point, where the Dirac asked only its value, and that is the whole difficulty. A $W^{1,p}$ test function with $p > d_a$ has a value at every point, its continuous representative (Theorem 1); its gradient sits one rung lower, an $L^p$ function defined only almost everywhere, and a single point is a null set, so the slope of $\psi$ at $q$ is not a meaningful question on $V$: the dipole is not in $V'$. Making the question meaningful would take a space whose members have continuous gradients, two derivatives in high integrability, $W^{2,p}$, one full derivative more than $V$ carries. And no admissible test space can pay that price: the kinked policy must itself be a test function, and the Koopman composition must preserve the space, so the smoothness is capped at one derivative whatever the space: the trilemma at the end of this appendix, run one rung up. Nothing beyond the dipoles will ever be needed: a distribution supported at a single point is a finite combination of the Dirac and its derivatives [Friedlander and Joshi, 1998, Thm. 3.2.1], and at first order only the mass and its dipoles survive.

So no admissible dual holds the moving atom, and this is not a defect of the choice of $V$ but a statement about the object: the first-order motion of an atom is not a measure. The resolution in Section 4.3 is bookkeeping on finitely many carriers: the carriers are finitely many, their motions fill the register $D$, finite-dimensional at $d_a = 1$ where every carrier is a point, and the responses live in $E = V' \oplus D$, the analysis carrying the density bulk and the register carrying the singular motions.

**Higher orders and the grades.** Section 5.2 runs the same accounting one order at a time. Each order of perturbation reads one more derivative of the test functions, the order-two transport pairing through Hessians, so the transport formulas are stated on the grades $V_k = \ell^2(\Theta) \otimes W^{k,p}(\mathcal{A})$, $V = V_1 \supset V_2 \supset \cdots$, one derivative gained per grade, the embedding condition $p > d_a$ unchanged. The trilemma below is not violated, because the grades never claim to be admissible test spaces in the sense of this appendix: the kinked policy $\bar{x}^*$ lies in $V_1$ and in no higher grade, and it is only the *transport* of the order-$k$ response that pairs on $V_k$. The objects that would breach the cap, the layers of Lemma 10 and the displaced dipoles, live in $D_k$, finitely many carriers each with a density-valued weight (Definition 7), and their pairings against the per-regime smooth policy are one-sided evaluations, defined off the binding boundary $\Gamma$ exactly as the first-order pairings below; the smooth-case analogue of the graded family is Tanaka [2022, Sect. 2.1, condition (I)]. The per-regime stability of the Koopman composition on the higher grades, and the push-forward algebra of $D_k$ beyond the dipole, are established where used, after Definition 7: the first regime by regime through the pinned-slot vanishing and the clause-2 change of variables, with a transversality clause on the iterated pre-images of $\Gamma$; the second outright, the per-carrier maps being bi-Lipschitz.

**Orbits, the binding boundary $\Gamma$, and the dipole pairings.** The dipoles of $D$ are read against the steady-state policy: each pairing is a number $v \cdot \nabla_a \bar{x}^*(\theta, u)$, the slope of $\bar{x}^*$ at the state $(\theta, u)$ the dipole sits at. Two definitions make clause 4 of Assumption 1, which guards these numbers, precise. The *forward orbit* of a state under the steady-state policy is the collection of states its households can visit: from $(\theta, a)$ the continuous coordinates move deterministically to $\bar{a}^*(\theta, a)$ while the Markov component jumps to any row $\theta'$ the transition allows, so the orbit collects, date by date, the finitely many states reached along positive-probability shock paths, countably many in all. The *binding boundary* $\Gamma$, a finite union because each bound contributes its own surface, collects the finitely many surfaces where two regimes meet: within a regime the policies $\bar{a}^*$ and $\bar{x}^*$ are continuously differentiable (Assumption 1), across a regime boundary their one-sided gradients disagree, so the gradients are undefined exactly there. Clause 4 asks that no point of any atom's forward orbit sit on $\Gamma$. Every state at which the transport ever places a dipole then lies where both the policy Jacobian $D_a \bar{a}^*$ that transports the directions and the gradient $\nabla_a \bar{x}^*$ that reads them out are defined, at every horizon. The clause is mild. Each atom contributes one condition per date, and a condition fails only when the calibration places one specific orbit point exactly on one of finitely many lower-dimensional surfaces, an event confined to a null set of calibrations; countably many dates contribute countably many such conditions, and a countable union of null sets is still null. Violating the clause therefore takes an exact coincidence, and a perturbation of the calibration removes it.

**The Gelfand triple.** Again, the model only ever sees the population through integrals. Every aggregate is a number $\int \varphi \, d\mu$, a mean, a moment, a marginal-propensity-weighted sum, got by integrating some function $\varphi$ against the distribution $\mu$. Two different kinds of object meet in that integral: the function $\varphi$ I integrate, and the population $\mu$ I integrate it against, which may pile up on atoms. For example, in the Krusell-Smith model, a lump of households lives at the borrowing limit. When perturbing the steady-state economy, I will move the population forward in time as the mirror image of the operation that takes expectations, precisely so that when I differentiate, the derivative falls on the smooth function $\varphi$ and never on the lumpy population. To set that up I keep the two kinds of object in separate spaces, $\varphi$ in a space $V$ and $\mu$ in its dual $V'$, with a third space in between for the objects that are of both kinds at once. Those three spaces are a Gelfand triple.

**Definition 13** (Gelfand triple)**.** A Gelfand triple, the evolution triple of Zeidler [1990, Def. 23.11], is a chain of three nested spaces $V \hookrightarrow L^2 \hookrightarrow V'$. At the top are the test functions $V = W^{1,p}(\mathcal{A})$, smooth enough that evaluating one at a single point is meaningful (Theorem 1); the moments and the steady-state policy are among them. At the bottom is the dual $V'$, holding everything that turns a test function into a number by integration, big enough to contain the constraint's lump, a Dirac mass that is no function at all. In between are the ordinary densities $L^2$, and the whole construction turns on what makes them special: a density belongs to both worlds at once. It is itself a function, so it could serve as a test function up in $V$; and it also integrates test functions, through $\varphi \mapsto \int g\,\varphi \, dx$, which is exactly what a distribution down in $V'$ does.

For the densities the two roles do more than coexist, they coincide. Riesz's theorem, stated just below, says that every bounded way of turning $L^2$ functions into numbers is already integration against some $L^2$ density: nothing acts on $L^2$ from outside, the densities are the whole of its dual. So on the middle space being a function and being a distribution are the same thing, and $L^2$ and its dual $(L^2)'$ are one and the same space. A density belongs to both worlds, and it is a special feature of $L^2$. The top space is not equal to its dual: $V'$ is strictly larger than $V$, which is precisely why it has room for the Dirac lumps that no test function could ever be. This single coincidence is what threads the three spaces onto one chain: a test function is in particular a density, so $V \hookrightarrow L^2$; a density is, by Riesz, also a distribution, so $L^2 \hookrightarrow V'$; and the middle space, being its own dual, is the one place where the world of functions and the world of distributions actually meet. For that reason $L^2$ is called the pivot of the triple.
One caveat, for later use. The middle space $L^2$ carries an inner product, but the top space $V = W^{1,p}$ does not once $p \neq 2$, so $V$ is a Banach and not a Hilbert space. This costs nothing here, because the only inner product I ever use is the one on $L^2$; in the standard terminology $V$ is the rigging of the triple, the smooth top space over the pivot, here Banach rather than Hilbert.

**Theorem 2** (Riesz representation)**.** *Every continuous linear functional $L$ on the Hilbert space $L^2$ is represented uniquely as*

$$L(v) \;=\; \langle u, v\rangle_{L^2}, \qquad \|L\| = \|u\|_{L^2},$$

*for some $u \in L^2$ [Brenner and Scott, 2008, (2.4.2), (2.4.6)]; the pivot is thereby identified with its own dual.*

**A caution for finite-dimensional intuition.** Three reflexes from $\mathbb{R}^n$ have to be set aside. Every linear functional over $\mathbb{R}^n$ is a dot product with a fixed vector, so a space and its dual are the same object; every functional is bounded; and all norms are equivalent, so one never has to choose a space at all. Each fails in infinite dimension, and the failures are exactly what the care above is for.

1. The dual $V'$ is genuinely larger than the space $V$: the Dirac $\delta_{(\theta,a)}$ is a bounded functional represented by *no* density, no $g$ with $\int g\,\varphi = \varphi(\theta, a)$, so $V'$ holds measures and distributions that are not functions at all.

2. Boundedness is no longer automatic but a real condition on the norm, as the spikes above show.

3. Inequivalent norms are genuinely different spaces with different duals, which is why the choice of $V$ is consequential and not cosmetic.

Riesz representation (Theorem 2) says the pivot $L^2$ is its own dual; it does not say $V$ is, and indeed $V' = \ell^2(\Theta) \otimes W^{-1,p'} \neq V$. The Gelfand triple is the device that keeps a function and integration-against-that-function distinct, the first in $V$ and the second in $V'$, while letting them coincide at the pivot. Identifying $V$ with $V'$ is the finite-dimensional habit to resist.

This is the stage on which the transfer operator of the next paragraph lives.

**Koopman and transfer operators.** The law of motion of the distribution is linear, so it is an operator, and I will have to differentiate it without the derivative ever reaching the kinked policy. The device is to build the forward map that pushes the population as the adjoint of the backward map that takes expectations: a derivative then falls on the smooth test function the expectation carries, never on the policy or the measure.

**Definition 14** (Koopman and transfer operators)**.** The Koopman operator $\mathcal{K}$ of Section 4.3 is the conditional-expectation operation Dynare 7 performs: it sends a function of tomorrow's state to its conditional expectation as a function of today's. Its adjoint $\Lambda = \mathcal{K}^*$, defined by $\langle \Lambda\mu, \varphi\rangle = \langle \mu, \mathcal{K}\varphi\rangle$ for all $\varphi \in V$, is the forward, or transfer (Perron–Frobenius), operator [Lasota and Mackey, 1994, Defs. 3.1.1, 3.2.3, 3.3.1].

Differentiating $\mathcal{K}$ asks only the outer test function to be smooth, never the policy, so no Dirac appears; $\Lambda$ pushes mass forward and reproduces the distribution's law of motion. The expectation step and the lottery push-forward are mirror images in that regard (Lemma 2).

Section 4.3 builds these two operators and proves they are an adjoint pair, with $\Lambda = \mathcal{K}^*$ the push-forward (Lemma 2); the Markov assumption on the shock is what makes $\mathcal{K}$ bounded on $V$, so that the adjoint exists at all. One entry of the Markov-operator dictionary is worth keeping here: the deterministic Koopman operator of Lasota and Mackey [1994, Def. 3.3.1], composition $\varphi \mapsto \varphi \circ S$ with a fixed map, is only the point-mass special case, since $\mathcal{K}$ averages in addition over the stochastic shock and so needs the general Markov-operator adjoint of Lasota and Mackey [1994, Def. 3.1.1, §3.3].

The push-forward $\Lambda$ of Lemma 2 is the Markov operator carried by the same kernel [Lasota and Mackey, 1994, (5.7.3)], and it is exactly the map Young [2010] discretizes. Young advances the histogram by sending each node's mass through the policy to $\bar{a}^*(\theta, a)$, which lands between two grid nodes, and splitting it onto the pair with the unique weights that conserve mass and mean. Those weights are the $P^1$ hat values at the landing point, so the lottery is the push-forward followed by the interpolation-dual projection onto the histogram (Lemma 11): not a convenient approximation but the Perron–Frobenius operator read on the nodal basis. On the grid the two are one transition matrix read two ways, the expectation contracting it along its rows and the push-forward along its columns, so the discrete transfer operator $\Lambda^d$ is the transpose of the discrete Koopman matrix. The expectation the economist computes and the distribution the economist simulates are one matrix transposed, the discrete face of $\langle \Lambda\mu, \varphi\rangle = \langle \mu, \mathcal{K}\varphi\rangle$. This is why the linearization lives on the Koopman side: its adjoint lands the derivative on the smooth $\varphi$ rather than dragging it across the atom and the kink, the exchange $\langle \Lambda^t(\cdots), \bar{x}^*\rangle = \langle \cdots, \mathcal{K}^t\bar{x}^*\rangle$ that identifies the expectation vectors of Auclert et al. [2021] in Section 4.4 and that Dynare 7's pipeline exploits in Section 4.5.

**Finite elements and the interpolant.** Everything so far is infinite-dimensional; to compute, Dynare 7 carries the distribution on a grid as a histogram. That histogram is a finite-element projection, so I borrow just enough finite-element vocabulary to name it precisely and, in the last step, to bound its error.

**Definition 15** (Finite element and nodal basis)**.** A finite element is Ciarlet's triple $(K, \mathcal{P}, \mathcal{N})$: an element domain $K$, a finite-dimensional space $\mathcal{P}$ of shape functions, and a basis $\mathcal{N} = \{N_i\}$ of nodal variables for the dual $\mathcal{P}'$ [Brenner and Scott, 2008, (3.1.1)]. The nodal basis $\{\phi_i\} \subset \mathcal{P}$ is the one dual to $\mathcal{N}$, $N_i(\phi_j) = \delta_{ij}$ [Brenner and Scott, 2008, (3.1.2)].

**Definition 16** (Interpolant)**.** For a function $v$ on which every nodal variable is defined, the interpolant is

$$\mathcal{I}_h v \;=\; \sum_i N_i(v)\,\phi_i$$

[Brenner and Scott, 2008, (3.3.1)], which reproduces nodal values.

The one-dimensional $P^1$ element has the hats $\phi_1(x) = 1 - x$, $\phi_2(x) = x$ [Brenner and Scott, 2008, (3.1.3)], and the tensor product gives the bilinear element on a rectangle [Brenner and Scott, 2008, (3.5.1)–(3.5.2)]. This hat is the kinked shape $1 - |x|$ of the weak-derivative example (Definition 10): its weak derivative is a bounded $L^\infty$ jump, never a Dirac, which is why the constraint atom never forces a derivative onto a delta. Young's lottery is exactly this interpolant acting on a point mass.

**Lemma 11** (Young's lottery is the discretized transfer operator)**.** *The dual interpolant $\mathcal{I}_h^*$, adjoint to $\mathcal{I}_h$, sends a point mass to its nodal split,*

$$\mathcal{I}_h^*\,\delta_x \;=\; \sum_i \phi_i(x)\,\delta_{a_i},$$

*the mass- and mean-conserving weights of Young's lottery, so the discrete forward map is $\Lambda^d := \mathcal{I}_h^*\Lambda$.*

*Proof.* The interpolant is $\mathcal{I}_h v = \sum_i v(a_i)\,\phi_i$, so for every test function $v$,

$$\langle \mathcal{I}_h^*\delta_x,\; v\rangle \;=\; \langle \delta_x,\; \mathcal{I}_h v\rangle \;=\; (\mathcal{I}_h v)(x) \;=\; \sum_i v(a_i)\,\phi_i(x) \;=\; \Big\langle \sum_i \phi_i(x)\,\delta_{a_i},\; v\Big\rangle,$$

which is the displayed identity. On a cell $[a_j, a_{j+1}]$ only $\phi_j, \phi_{j+1}$ survive, with $\phi_j(x) = (a_{j+1} - x)/(a_{j+1} - a_j)$ and $\phi_{j+1}(x) = (x - a_j)/(a_{j+1} - a_j)$, summing to one and reproducing the mean $\phi_j(x)\,a_j + \phi_{j+1}(x)\,a_{j+1} = x$: Young's split. Taking $x = \bar{a}^*(\theta, a)$ and composing with the push-forward $\Lambda$ (Lemma 2) gives $\Lambda^d = \mathcal{I}_h^*\Lambda$. □

**Collocation and prolongation.** Two further mesh operations, both built from the interpolant, carry the dynamics of Section 4.5. *Collocation* determines a member of the finite-element space by imposing an equation exactly at the mesh nodes: the unknowns are the nodal values, and an equation that acts pointwise in the state, node by node, decouples into one small system per node, solved independently. The time iteration of Section 3.3 and the per-node solves of the policy-response recursion are collocation in this sense. *Prolongation* carries a member of one mesh's finite-element space onto another, typically finer, mesh: evaluate the interpolant at the new mesh's nodes, that is, apply the new mesh's nodal variables to it, one rectangular matrix whose entries are the old mesh's hats evaluated at the new nodes. The evaluation is exact on the interpolant itself, so prolongation adds no error of its own; what the finer mesh receives carries, unchanged, the coarse mesh's interpolation error in representing the underlying function. Section 4.5 uses it to read the policy responses, solved on the coarse policy mesh, on the finer distribution mesh.

**Interpolation error.** One thing is left: that the histogram is faithful, returning the right aggregates as the grid is refined. This is an interpolation-error statement.

**Theorem 3** (Interpolation error)**.** *On a shape-regular mesh of width $h$, the $P^1$ interpolant satisfies*

$$\|\varphi - \mathcal{I}_h\varphi\|_{H^1} \;=\; O(h)\,|\varphi|_{H^2}$$

*for $\varphi \in H^2$ [Brenner and Scott, 2008, (4.4.4)–(4.4.5) per element, summed over the mesh as Thm. (4.4.20), estimate (4.4.21)].*

The consistency of the histogram is exactly this interpolation-error statement. For a kinked, merely Lipschitz integrand the nodal interpolant is not even defined, and the right tool is the averaging Scott–Zhang/Clément interpolant, defined for integrable functions and converging at the rate $h^{k-s}$ on $W_p^k$ [Brenner and Scott, 2008, (4.8.1) and Thm. (4.8.12), estimate (4.8.14)], which delivers the convergence rate honestly rather than pretending the smooth rate still holds.

**Which test space, and why not a tighter one.** The choice of $V$ is the price of admission for the whole construction: the dynamics never see the distribution $\bar{\Omega}^*$ as a function, only as a rule acting on $V$, so they live in the dual $V'$, and $V$ fixes at once which functions can be paired against the distribution and which distributions $V'$ can hold. Two things must fit inside it. The dual must hold the distribution's atoms, the mass the constraints pile at their limits, which forces point evaluation to be bounded on $V$, that is $V \hookrightarrow C^0$. And $V$ itself must hold the kinked steady-state policy and stay closed under the Koopman composition $\varphi \mapsto \varphi \circ \bar{a}^*$, which forbids asking more smoothness than the Lipschitz policy has. One would like to meet both demands inside a Hilbert space, where the inner product buys orthogonal projection and the self-dual $L^2$ pairing the primer leaned on. The question is whether the model allows it, and the answer comes in two parts: $W^{1,p}$ with $p > d_a$ always works, and a tighter, Hilbert space is available only when the steady-state policy is shaped gently enough.

**Lemma 12** (The frame is admissible)**.** *Under Assumption 1, $V \hookrightarrow C^0$, so point evaluation is bounded and every constraint atom, including a corner where up to $d_a$ bounds bind at once, lies in $V' = \ell^2(\Theta) \otimes W^{-1,p'}(\mathcal{A})$; the policy lies in $V$; and the Koopman operator maps $V$ to $V$, so its adjoint $\Lambda$ acts on $V'$. With several coordinates, on a regime that pins a coordinate the composition $\varphi \circ \bar{a}^*$ carries the regularity of the trace of $\varphi$ on the set $\{a^i = \underline{a}^i\}$, fractional rather than full; the identities of Section 4.3 are accordingly proved on Lipschitz test functions and reach all of $V$ by density.*

Lemma 12 asks two things of the test space, and $W^{1,p}$ with $p > d_a$ delivers both. I take them one at a time.

The first demand comes from the atoms. A point mass acts on a test function by evaluating it at one point, so it is a bounded functional exactly when point evaluation is bounded on $V$. This is what the Sobolev embedding buys (Theorem 1): for $p > d_a$ every test function has a continuous representative, and its value at any single point, a corner where all $d_a$ bounds bind included, is controlled by its norm (Proposition 16).

The second demand comes from the policy, and it comes twice. First, the steady-state policy $\bar{x}^*$ must itself be a test function: the aggregates of Section 3.1 pair the distribution against it, so it has to lie in $V$. It does, with nothing to spare. The policy is Lipschitz, kinked at the constraint (Assumption 1): it has exactly one bounded weak derivative, so it lies in $W^{1,\infty}$, hence in $W^{1,p}$ on the bounded $\mathcal{A}$, while a space asking for a second derivative would expel it. Second, the Koopman operator must map $V$ into $V$, because every derivative of Section 4.3 is taken through $\mathcal{K}$ and must land back in the space it started from. It does: $(\mathcal{K}\varphi)(\theta, a)$ averages the compositions $\varphi(\theta', \bar{a}^*(\theta, a))$ over tomorrow's row $\theta'$, and composing with the Lipschitz $\bar{a}^*$ multiplies the weak gradient of $\varphi$ by the bounded policy Jacobian $D_a\bar{a}^*$, leaving it in $L^p$, the chain rule that Corollary 1 makes precise below. One weak derivative, in integrability high enough to embed in $C^0$: enough continuity for the atoms, no more smoothness than the kinked policy has.

With several coordinates $d_a \geq 2$, the composition needs one more care, the refinement recorded in the lemma's last sentence. On a regime that pins coordinate $i$, every household of the regime lands on the set $\{a^i = \underline{a}^i\}$, so the composition $\varphi \circ \bar{a}^*$ reads only the values of $\varphi$ there, its trace. Restriction costs smoothness: the trace of a $W^{1,p}$ function is not a $W^{1,p}$ function of the carrier's variables but a member of the fractional class $W^{1-1/p,\,p}$ on it, defined and anchored below. Rather than manipulate fractional gradients, the identities of Section 4.3 are proved for Lipschitz test functions, whose restriction to such a set is again Lipschitz, and extend to all of $V$ by density, both sides of each identity being bounded functionals.

Two results of Section 4.3 were stated there and proved here, both refinements of the composition argument just given.

*Proof of Corollary 1.* The $\theta'$-average in $\mathcal{K}$ is a finite sum, so it suffices to differentiate one composition $\psi \circ \bar{a}^*$ at one shock node $\theta'$.

*State direction.* The test function $\psi$ is Lipschitz, hence differentiable off a Lebesgue-null set $N \subseteq \mathcal{A}$ by Rademacher's theorem [Evans and Gariepy, 2025, Thm. 3.2], and the landing map $\bar{a}^*$ is continuously differentiable on each regime (Assumption 1). Consider first a regime that leaves every coordinate free. Clause 2 of Assumption 1 bounds $|\det D_a\bar{a}^*|$ away from zero there, so the pre-image $(\bar{a}^*)^{-1}(N)$ of the null set $N$ under the landing map is itself Lebesgue-null, and $\nabla\psi$ is defined at the landing point $\bar{a}^*(\theta, a)$ for almost every state $(\theta, a)$ of the regime; at every such state both maps of the composition are differentiable, and the classical chain rule gives

$$\partial_{a^j}\big(\psi \circ \bar{a}^*\big)(\theta, a) \;=\; \sum_i \partial_{a^j}\bar{a}^{*\,i}(\theta, a)\;\big(\partial_{a'^i}\psi\big)\big(\theta', \bar{a}^*(\theta, a)\big).$$

On a regime that pins a coordinate, the composition is constant along the pinned directions, so the corresponding left-hand derivatives vanish, and the weights $\partial_{a^j}\bar{a}^{*\,i}$ of the pinned components vanish with them, which is the convention in the statement. The free coordinates of the same regime land within the set $\{a^i = \underline{a}^i\}$, and their derivatives are read as the tangential derivatives of the Lipschitz restriction of $\psi$ to it, which exist almost everywhere there by Rademacher's theorem applied within it; the free block of $D_a\bar{a}^*$ is non-degenerate by the same clause 2, so pre-images of face-null sets are null and the chain rule runs as before.

*Policy direction.* For a bounded direction $da$, the difference quotient

$$\frac{1}{h}\Big[\psi\big(\theta', \bar{a}^*(\theta, a) + h\,da(\theta, a)\big) - \psi\big(\theta', \bar{a}^*(\theta, a)\big)\Big]$$

is bounded by the Lipschitz constant of $\psi$ times $\|da\|_\infty$, and at almost every state $(\theta, a)$ at which $\nabla\psi$ exists at the landing point, the same set as above, it converges to $\sum_j da^j(\theta,a)\,(\partial_{a'^j}\psi)(\theta', \bar{a}^*(\theta,a))$; dominated convergence carries the limit through the $\theta'$-average. On each set $\{\bar{a}^{*i} = \underline{a}^i\}$ the pinned component of the directions of interest vanishes, $da^i = 0$, so the quotient displaces the landing point within the set $\{a^i = \underline{a}^i\}$ alone, and the tangential reading of the previous paragraph serves; with a single predetermined coordinate the whole direction vanishes there and the quotient with it. □

*Proof of Lemma 3.* Work at one pair of Markov nodes $(\theta, \theta')$; the finite $\pi$-average preserves every class in play. Split the state space into the regimes of Assumption 1 and treat each regime's contribution to $(\mathcal{L}^{(a)}\omega)^i = \Lambda\big(\sum_j \partial_{a^j}\bar{a}^{*i}\,\omega^j\big)$ separately.

*One predetermined coordinate.* Two regime types arise. On the binding regime the weight $\partial_a\bar{a}^*$ vanishes, so the weighted field carries no mass there: the states that a bare push-forward would collapse into the atom at the bound $\underline{a}$ are annihilated before the push. On an interior regime the landing map $\bar{a}^*(\theta,\cdot)$ is continuously differentiable with $\partial_a\bar{a}^*$ bounded away from zero (clause 2), hence strictly monotone with a continuously differentiable inverse $g(\theta,\cdot)$; the change of variables $a' = \bar{a}^*(\theta,a)$ gives, for every bounded test function $\varphi$,

$$\int \varphi\big(\bar{a}^*(\theta,a)\big)\,\partial_a\bar{a}^*(\theta,a)\,\omega(\theta,a)\,da \;=\; \int \varphi(a')\,\omega\big(\theta, g(\theta,a')\big)\,da',$$

the weight $\partial_a\bar{a}^*$ cancelling the Jacobian $\partial_{a'}g = 1/\partial_a\bar{a}^*$ of the substitution exactly. The regime therefore contributes the density $\omega(\theta, g(\theta,\cdot))$ on its image: a composition, so $\|\mathcal{L}^{(a)}\omega\|_{L^\infty} \le \|\omega\|_{L^\infty}$ after the $\pi$-average; in $L^{p'}$ whenever $\omega$ is, by the reverse change of variables and the boundedness of $\partial_a\bar{a}^*$; and atom-free. Summing the finitely many regimes and iterating preserves the class.

*Several predetermined coordinates.* On a regime that leaves every coordinate free, the same change of variables runs with the matrix $D_a\bar{a}^*$ in place of the scalar slope: the contributed density is bounded by $\sup\|D_a\bar{a}^*\| \,/\, \inf|\det D_a\bar{a}^*|$ times $\|\omega\|_{L^\infty}$, the numerator controlled by the Lipschitz bound and the denominator by clause 2, and no atom can form, pre-images of points being finite sets. On a regime that pins a coordinate, the pinned rows of the weight vanish but the free rows do not: the regime's weighted mass is pushed exactly onto the pinned coordinate's set $\{a^i = \underline{a}^i\}$, where it lands as a layer, its density obtained by the same change of variables run within the carrier through the non-degenerate free block of $D_a\bar{a}^*$. Full-binding corners contribute nothing, their whole weight vanishing. No point atom arises anywhere: an atom would require positive $\omega$-mass on the pre-image of a single point, and every such pre-image either meets a region where the weight vanishes or is a null subset of a non-degenerate regime. The created layers do not inherit the fractional face class of Assumption 1, and they do not need it: their densities are fiber integrals of the interior density, in $L^{p'}$ of their carriers by Minkowski's inequality, and an $L^{p'}$ weight is all a degree-zero layer's pairings consume; the main text records this next to Lemma 4. □

Whether one can do better, and test against a tighter Hilbert space, is not a matter of taste but of the shape the constraints draw in the stationary distribution. Three economic features decide it: which bounds bind, how many bind at once, and whether the boundary they trace runs along the predetermined-variable axes. The first sets where the atoms sit, the second how singular they are, and the third whether a Hilbert space can still hold the policy. Three notions, codimension, trace, and tensor product, turn these into a ladder of test spaces, every rung Hilbert and comfortable until the last.

The instinct is easiest to see by walking up the asset count. In the one-asset models the toolkit was built for, Krusell–Smith and one-asset HANK, only the borrowing constraint binds; the mass it leaves is a layer of constrained households along $\{a = 0\}$, a codimension-one carrier, and a Hilbert space holds it. The two-asset model is where the comfort runs out, and it runs out for an economic reason. Its liquidity-constrained households, the wealthy hand-to-mouth who hold illiquid wealth but no liquid buffer, pile on $\{b = 0\}$, still a codimension-one carrier a Hilbert space can hold; but the poorest of them hold neither asset and pile at the corner $(b,a) = (0,0)$, a codimension-two point mass, and evaluating a function at a point of the plane is unbounded on $H^1(\mathbb{R}^2)$. A more elaborate Hilbert space, the tensor product of one-dimensional factors, could still reach the corner if the policy's kink ran along an axis; but the threshold at which a household turns liquidity-constrained rises with its illiquid wealth, so the kink runs diagonally, and no Hilbert space holds the policy. The wealthy hand-to-mouth on a sloped frontier and the truly poor at the corner are exactly what no Hilbert space can carry, and that is why the safe choice is the Banach space $W^{1,p}$. The three notions below make each step of this instinct precise, and I take them in turn.

**Definition 17** (Codimension)**.** For a linear subspace $F$ of $\mathbb{R}^{d_a}$, the codimension is the familiar $\operatorname{codim} F = d_a - \dim F$, the number of independent directions transverse to $F$; a hyperplane has codimension one, the origin $\{0\}$ codimension $d_a$. The notion needed here is the manifold analogue of this, carrying the same count to curved sets: a set $S \subseteq \mathcal{A} \subseteq \mathbb{R}^{d_a}$ has codimension $k$ when, near each of its points, it is the common zero set of $k$ functions with linearly independent gradients, a surface of dimension $d_a - k$. In the flat case the $k$ functions are independent linear forms $\ell_1, \dots, \ell_k$,

$$\{\, x \in \mathbb{R}^{d_a} : \ell_1(x) = \cdots = \ell_k(x) = 0 \,\},$$

a plane of dimension $d_a - k$.

The atoms sit on these sets and on their date-by-date forward images, which the invariance $\Lambda\bar{\Omega}^* = \bar{\Omega}^*$ forces into the singular family (clause 1 of Assumption 1): the mass an unbinding row moves off a bound lands at interior orbit points and must carry singular mass there too. A binding bound itself plays two roles that Definition 2 keeps apart: the positive-measure subset $\{(\theta, a) \in \Theta \times \mathcal{A} : \bar{a}^{*i}(\theta, a) = \underline{a}^i\}$ over which the policy pins the coordinate flat, whose measure is the atom's mass; and the set $F_i = \{a \in \mathcal{A} : a^i = \underline{a}^i\}$, the subset of $\mathcal{A}$ alone on which the pinned households pile up. The set $\{\bar{a}^{*i} = \underline{a}^i\}$ says how much mass arrives; $F_i$ says where it sits; and it is the codimension of $F_i$, never of that set, that measures how singular the arriving mass is. In the two-asset model the liquid bound alone leaves a codimension-one layer on $\{b = 0\}$; both bounds binding at once leave a codimension-two corner at $(b, a) = (0, 0)$, where the poorest hold neither asset. The higher the codimension, the lower-dimensional the set the atom occupies and the harder it is as a functional, the precise threshold being what the trace theorem below supplies.

The reading to carry into that theorem is that a layer of mass spread over a codimension-one surface is still a bounded functional on $H^1$, where a point mass is not.

**Theorem 4** (Trace)**.** *For a domain $\mathcal{A}$ with piecewise-smooth boundary and a hypersurface $\Gamma$ of codimension one, the restriction $\varphi \mapsto \varphi|_\Gamma$, defined first on smooth functions, extends to a bounded linear map $H^1(\mathcal{A}) \to L^2(\Gamma)$: there is a constant $C$ with*

$$\|\varphi|_\Gamma\|_{L^2(\Gamma)} \;\le\; C\,\|\varphi\|_{H^1(\mathcal{A})}$$

*[Brenner and Scott, 2008, Thm. 1.6.6, printed for the boundary of a Lipschitz domain and applied to an interior $\Gamma$ through the two subdomains it bounds].*

This is what lets a measure living on $\Gamma$ act on $H^1$, and the step is worth taking slowly. Let $\nu$ be a finite measure on $\Gamma$ with an $L^2$ density, mass spread over a codimension-one carrier. It pairs with a test function by $\langle \nu, \varphi \rangle = \int_\Gamma \varphi|_\Gamma \, d\nu$, and by the Cauchy–Schwarz inequality and then the trace bound,

$$|\langle \nu, \varphi \rangle| \;\le\; \|\varphi|_\Gamma\|_{L^2(\Gamma)}\, \|\nu\|_{L^2(\Gamma)} \;\le\; C\,\|\nu\|_{L^2(\Gamma)}\, \|\varphi\|_{H^1},$$

so $\nu$ is a bounded functional on $H^1$, an element of its dual. A layer $\sigma^*$ of Definition 2 is exactly such a $\nu$, which is why one binding bound is harmless. A point mass is not: it sits on a single point, of codimension $d_a$, and pairing it with $\varphi$ would mean evaluating $\varphi$ there, which by the spikes above is unbounded on $H^1(\mathbb{R}^{d_a})$ once $d_a \ge 2$. The trace reaches codimension one, but not codimension $d_a$.

The mass just paired is the *layer* of Definition 2, the object Assumption 1 and Section 4.3 use throughout: a measure carried by a set $\{a : a^i = \underline{a}^i\}$ with a density on it, taken with respect to the carrier's own surface measure rather than the volume of $\mathcal{A}$. The wealthy hand-to-mouth households of the two-asset model are the running example: households holding illiquid wealth but no liquid buffer spread along $\{b = 0\}$, their density varying with the illiquid coordinate $a$. Clause 1 of Assumption 1 admits finitely many such layers in the steady state, and the velocity operator of Section 4.3 creates new ones, a regime that pins one coordinate pushing its mass onto that coordinate's set.

A moving layer can be displaced in two ways, and the distinction decides which tool serves the pairing. A *tangential* displacement slides the layer along its own carrier: the free coordinates move, the pinned one does not, the layer remains carried by the same set, and only its density shifts. A *normal* displacement would carry mass off the carrier. At the rows where the layer's own coordinate binds, that coordinate's response is zero (Proposition 2), so the tangential displacement is the only one open there, the wealthy hand-to-mouth drifting along $\{b = 0\}$ as their illiquid wealth moves; at an unbinding row the mass leaves the carrier for the interior, the layer analogue of the moving atoms that the finite dipole span $D$ of Section 4.3 tracks.

Pairing a tangentially displaced layer is the one place this document needs smoothness between zero and one derivative, and the notion deserves a plain definition. Between $L^p = W^{0,p}$, functions with a value in the mean, and $W^{1,p}$, functions with a full derivative in the mean, runs a scale of classes $W^{s,p}$ with $0 < s < 1$, carrying a fractional amount of differentiability: their increments $|u(x) - u(y)|$ are controlled, in an $L^p$-averaged sense, like the fractional power $|x - y|^s$,

$$|u|^p_{W^{s,p}(F)} \;=\; \int_F\!\int_F \frac{|u(x) - u(y)|^p}{|x - y|^{\dim F + s\,p}}\, dx\, dy \;<\; \infty.$$

Traces are where these classes arise, because restriction costs smoothness, and it costs exactly $1/p$ of a derivative: the trace of a $W^{1,p}(\mathcal{A})$ function on a codimension-one hypersurface $\Gamma$ lands in the fractional class $W^{1-1/p,\,p}(\Gamma)$ [Adams and Fournier, 2003, Thm. 7.39, carried from the flat trace to smooth hypersurfaces by Rem. 7.45(1)], the statement given there on the Besov scale $B^{1-1/p;\,p,p}$, the interpolation name of the same fractional class. The tangential gradient of the trace accordingly sits one full derivative lower, in $W^{-1/p,\,p}(\Gamma)$. A layer displaced along its own carrier pairs with exactly that tangential gradient, and the pairing is bounded, an honest integral, as soon as the layer's density lies in the dual fractional class $W^{1/p,\,p'}(\Gamma)$, the regularity clause 1 of Assumption 1 imposes on the steady-state layers; the created layers of Section 4.3 carry $L^{p'}$ densities and pair at degree zero, through plain traces, so they never call on the fractional class. I use the fractional classes only through this one pairing and do not develop them; the trace theorem in the $L^2$ form above is Brenner and Scott [2008, Thm. 1.6.6].

One remark extends the reach of the same tool. The trace bound is indifferent to whether the hypersurface $\Gamma$ is a carrier $\{a^i = \underline{a}^i\}$ on the boundary of $\mathcal{A}$ or an interior surface: the binding boundary $\Gamma_b$ of a bound, along which the layers of Section 5.2 sit, is a codimension-one hypersurface per Markov row (Lemma 7), so those layers pair with test functions through the same restriction map, with the same bound, weight against trace, the higher degrees reading traces of normal derivatives on the higher grades. The interface layers of the higher orders are, for the pairing, layers like any other; their carrier simply sits in the interior of $\mathcal{A}$.

**Definition 18** (Tensor-product Sobolev space)**.** The tensor-product space allows one weak derivative in each coordinate separately, the fully mixed $\partial_1 \cdots \partial_{d_a}$ included, but no second derivative in any single coordinate. Formally, on a product domain $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_{d_a}$, the tensor product $H^1(\mathcal{A}_1) \otimes \cdots \otimes H^1(\mathcal{A}_{d_a})$ is the completion of the finite sums of products $\varphi_1(a_1) \cdots \varphi_{d_a}(a_{d_a})$, $\varphi_i \in H^1(\mathcal{A}_i)$, in the norm controlling every mixed weak derivative $\partial_1^{\alpha_1} \cdots \partial_{d_a}^{\alpha_{d_a}}$ with each $\alpha_i \in \{0, 1\}$.

Why point evaluation is bounded on this space is the crux, and it is a one-line factorization. On a product the evaluation splits,

$$\big\langle \delta_{(a_1,\ldots,a_{d_a})},\; \varphi_1 \cdots \varphi_{d_a} \big\rangle \;=\; \varphi_1(a_1) \cdots \varphi_{d_a}(a_{d_a}) \;=\; \prod_{i=1}^{d_a} \big\langle \delta_{a_i},\; \varphi_i \big\rangle,$$

so the joint Dirac is the tensor product $\delta_{(a_1,\ldots,a_{d_a})} = \delta_{a_1} \otimes \cdots \otimes \delta_{a_{d_a}}$ of one-dimensional Diracs. Each factor $\delta_{a_i}$ is bounded on $H^1(\mathcal{A}_i)$, since in one dimension $H^1 \hookrightarrow C^0$ (one derivative beats half a dimension, $1 > 1/2$), and a tensor product of bounded functionals is bounded on the tensor product of the spaces, with norm the product of the norms. So the joint Dirac is bounded. The threshold is therefore half a derivative *per coordinate*, $s_i > 1/2$ for each $i$ separately, never $d_a/2$ spent jointly: the cost of admitting a corner does not grow with the number of assets. This is what isotropic $H^s$, charging $s > d_a/2$ for the same continuity, cannot match.

Whether a policy lies in this space is decided by its mixed weak derivative, and the word weak matters: $\partial_i \partial_j$ here is the distributional derivative, the only sense in which a Dirac can appear, a classical second derivative being an ordinary function wherever it exists. The question is whether this weak mixed derivative is an $L^2$ function, with no singular part. For the local model of a policy bending into a constraint, $f = \max(0, g)$ with $g$ smooth and the kink along $\Gamma = \{g = 0\}$, a distributional computation gives

$$\partial_i \partial_j f \;=\; \mathbf{1}\{g > 0\}\, \partial_i \partial_j g \;+\; (\partial_i g)(\partial_j g)\, \delta(g),$$

with $\delta(g)$ the surface Dirac on $\Gamma$, normalized as $\delta(g) = \delta_\Gamma / |\nabla g|$ by the coarea formula, the level-set remark of Section 5.2, which computes this same singular term by the surface-layer route. The first term is an ordinary $L^2$ function; the second is singular, a Dirac on the crease weighted by the product of the gradient's components. Whether that Dirac is present is best read through the moving kink. The

liquid policy bends at the liquidity frontier $b = b^*(a)$, flat at the bound below it and saving above, so its $b$-derivative jumps across the frontier. The mixed derivative $\partial_b \partial_a$ measures how that jump moves as illiquid wealth $a$ changes, that is, how the kink slides. If the frontier is flat, $b^*$ the same for every $a$, the kink sits still: stepping in $a$ runs along it, never across, and the mixed derivative is an ordinary function. If the frontier rises with wealth, $b^*(a)$ sloping, the kink slides as $a$ grows; stepping in $a$ then drags it across the point one sits at, and the derivative of a sliding step is a Dirac on the crease, of strength exactly the slope $b^{*\prime}(a)$, the $(\partial_i g)(\partial_j g)$ of the formula. So the singular term vanishes, and $f$ lies in the tensor-product space, exactly when the crease is axis-aligned, its normal having a single nonzero component. A flat frontier leaves nothing there; a rising one, the signature of the wealthy hand-to-mouth, leaves a Dirac the tensor space cannot hold, because that space buys the corner's continuity with one derivative in each coordinate, the fully mixed one included, and the Dirac lands exactly there. The last rung $W^{1,p}$ asks for no mixed derivative at all, so the slope of the frontier costs it nothing.

It is natural to object that a kink along a slanted line is no real obstruction, since a change of coordinates could rotate it onto an axis. It cannot, for three reasons. First, the geometry does not cooperate: the crease is in general curved, not straight, the liquidity threshold $b^*(a)$ being a nonlinear function of illiquid wealth, so no single affine map straightens it; and the sets $\{b = 0\}$ and $\{a = 0\}$ are already axis-aligned, so a map that straightens the diagonal crease tilts them and the codimension-one layers they carry. Two straight carriers and one curved crease cannot all meet the axes at once. Second, the tensor-product space is not preserved by coordinate changes that mix the variables: only separable maps, acting one coordinate at a time, leave the product structure and the mixed-derivative norm intact, whereas a rotation that straightens the crease turns $\partial_i \partial_j$ into a combination that includes the pure second derivatives $\partial_i^2$ and $\partial_j^2$, which the space does not control. Reparameterizing does not place the policy in the space; it replaces the space. Third, the histogram is itself built on a rectangular grid in the economic coordinates $(b, a)$, fixed by the axis-aligned bounds, with no rectangular grid available in the rotated frame. The first-order space of the last rung escapes all of this because it is coordinate-robust: it constrains only the size of the gradient, unchanged up to a constant under any bi-Lipschitz change of variables, and asks for no mixed derivative at all.

The ladder of admissible test spaces then runs as follows, from the tightest and most comfortable rung to the most general, each available only under an extra structural condition on the model.

1. *One continuous coordinate*, $d_a = 1$. Here $H^1(\mathbb{R}) \hookrightarrow C^0$ already, since $1 > 1/2$; the test space is Hilbert, point masses are admissible, and the Lipschitz policy lies in it. Krusell–Smith and the one-asset model live here, and nothing below is needed.

2. *Several coordinates, one bound binding*, so the atom's support is of codimension one (Definition 17). Isotropic $H^1(\mathbb{R}^{d_a})$ no longer embeds in $C^0$, but a codimension-one layer is still admissible by the trace theorem (Theorem 4). The Hilbert frame survives as long as no corner of positive mass forms, a codimension-$k$ corner requiring $s > k/2$.

3. *Several coordinates, bounds co-binding, creases aligned with the axes.* A codimension-two corner atom needs $V \hookrightarrow C^0$, which the tensor-product space of Definition 18 supplies through its factorized point evaluation, at the per-coordinate cost of more than half a derivative in each direction rather than $s > d_a/2$ overall. But the policy lies in it only if its kinks run along the axes: by the crease computation above, an off-axis crease leaves a Dirac in the mixed derivative, so the policy leaves the tensor space and the Koopman composition no longer preserves it.

4. *The general case*, bounds co-binding and creases not axis-aligned. Only a first-order space holds the policy, and the Sobolev embedding (Theorem 1), in its $p > d_a$ form, supplies one that still embeds in $C^0$: $W^{1,p}(\mathbb{R}^{d_a}) \hookrightarrow C^0$, a single weak derivative in high enough integrability buying continuity with no second derivative at all. This is Assumption 1.

The ladder is a trilemma. In $d_a \geq 2$ no space is at once first order, embedded in $C^0$, and Hilbert: a first-order Hilbert space is $H^1 = W^{1,2}$, and $W^{1,2}(\mathbb{R}^{d_a}) \hookrightarrow C^0$ would need $2 > d_a$. To reach $C^0$ one spends either a second derivative, by $H^s$ with $s > 1$ or by the tensor space, and gives up first order, or higher integrability, $W^{1,p}$ with $p > d_a$, and gives up the inner product. The kink forces first order and the corner atom forces $C^0$, so it is the Hilbert structure that must go (Table 1 sets the rungs side by side, and Figure 4 draws the trilemma). The dipole $\dot{\delta}_{(\theta,u)}[v]$ (Definition 3) sits one rung further still. Reading it takes the slope of the test function at a point, a second derivative on top of the $C^0$ embedding, which not even the tensor rung affords; and the two constraints that built the ladder, the kinked policy that must itself be a test function and the Koopman composition that must preserve the space, cap

| Setting | $V$ | $\hookrightarrow C^0$ via | Hilbert | Atom in $V'$ | Breaks when |
|---|---|---|---|---|---|
| $d_a = 1$ | $H^1(\mathbb{R})$ | $1 > 1/2$ | yes | point | (suffices) |
| $d_a \geq 2$, one bound | $H^1(\mathbb{R}^{d_a})$ | trace | yes | codim-1 layer | bounds co-bind |
| $d_a \geq 2$, axis-aligned | $\bigotimes_i H^1(\mathcal{A}_i)$ | factorized | yes | codim-$k$ corner | crease off-axis |
| general | $W^{1,p}$, $p > d_a$ | Sobolev | no | codim-$k$ corner | (safe choice) |

Table 1: The ladder of test spaces for the aggregate block. Every rung but the last is a Hilbert space; the kink forces first order and the corner atom forces the $C^0$ embedding, and in two or more asset dimensions no Hilbert space has both, so the two-asset model lands on the Banach space $W^{1,p}$. Figure 4 draws this trilemma.

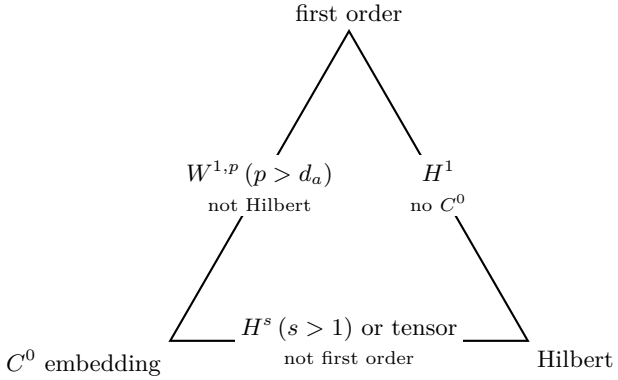


Figure 4: The test-space trilemma in $d_a \geq 2$. No space holds all three of first-order regularity, the $C^0$ embedding, and Hilbert structure; each edge carries the space with its two endpoint properties but not the third. The kink forces first order and the corner atom forces $C^0$, so the construction sits on the left edge and gives up the Hilbert corner, landing on $W^{1,p}$ with $p > d_a$.

the smoothness at one derivative on every rung. The ladder thus does its original job (the atoms need $C^0$, the policy caps the order at one, the density responses pair by Hölder) and settles a further point: the dipole is out of reach of every admissible test space at once, so the motion of the singular mass is necessarily bookkeeping, the finite dipole span $D$ of Section 4.3, and no bigger dual was there to buy.

The two-asset model sits on the last rung. Its stationary distribution carries a genuine codimension-two atom at $(b, a) = (0, 0)$, the poorest households pinned where both bounds bind (Figure 6); and its liquid-constraint boundary is not axis-aligned, since the threshold below which a household is liquidity-constrained rises with its illiquid wealth, so the policy's mixed derivative is singular on a sloped curve (Figure 5). Neither the trace rescue of the second rung nor the tensor space of the third applies, which is why Assumption 1 is pitched at the first-order Sobolev space $W^{1,p}$.

What the general choice gives up against the Hilbert rungs is the inner product and what it powers: the self-dual pivot $L^2$ is untouched, but orthogonal projection, the $L^2$-duality that would double the interpolation rate to $O(h^2)$, and the exact match between the tensor test space and the bilinear histogram basis are Hilbert luxuries that $W^{1,p}$ lacks. Little is lost in practice. The policy is only Lipschitz, so the consistency rate is the honest $O(h)$ of the Scott–Zhang interpolant whichever space one names; and $W^{1,p}$ has the clean dual $W^{-1,p'}$, so the Gelfand triple and the transfer-operator calculus of Section 4.3, which use only the duality pairing and never the inner product, carry over unchanged. The two pairings the distribution enters are met by this one first-order space: the atomic aggregation against $\bar{\Omega}^*$ through the continuity the Sobolev embedding gives, and the interior displacement of Section 4.3 through Hölder integrability. The motion of the singular mass itself sits outside every rung of the ladder, and Section 4.3 carries it as the finite dipole span $D$.

# D Differentiating the heterogeneous residual

This appendix derives the three partial derivatives of the residual $\mathcal{F}$ displayed in Section 4.2. Recall

$$\mathcal{F}\big(\bar{x}, \bar{x}^+, \bar{Y}\big)\langle \theta, a, \mathcal{Z}\rangle \;=\; F\big(\theta,\, a,\, \bar{x}(\theta, a, \mathcal{Z}),\, \mathbb{E}\big[\bar{x}^+ \,\big|\, \theta, a, \mathcal{Z}\big],\, \bar{Y}(\mathcal{Z})\big)\,,$$

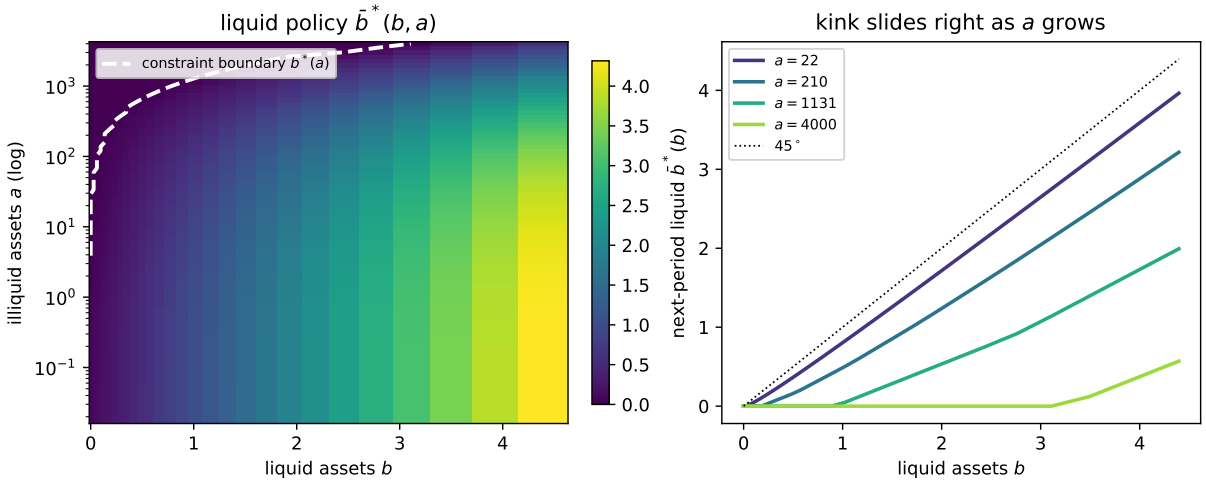


Figure 5: The liquid-saving policy of the two-asset model is kinked along a curve that is not axis-aligned. *Left:* the policy $\bar{b}^*(b,a)$ at the median productivity state, with the liquidity-constraint boundary $b^*(a)$ overlaid (dashed); the threshold rises with illiquid wealth, so the crease runs diagonally. *Right:* cross-sections $\bar{b}^*(b)$ at four illiquid levels, each flat at the bound and then rising with a bounded slope (the policy is Lipschitz), with the kink sliding right as $a$ grows. A diagonal crease puts a Dirac in the mixed derivative $\partial_b \partial_a \bar{b}^*$, so the policy is not in the tensor-product Sobolev space.

where $\mathbb{E}[\cdot \mid \theta, a, \mathcal{Z}]$ is the conditional expectation of Section 2, taken jointly over the next-period idiosyncratic state $\theta'$ and the aggregate innovation $\mathcal{E}'$, the next-period policy is evaluated at $\bar{x}^+\langle \theta, \theta', a, \mathcal{Z}, \mathcal{E}' \rangle = \bar{x}^+(\theta', \bar{a}(\theta, a, \mathcal{Z}), \mathcal{Z}')$, with $\theta'$ and $\mathcal{E}'$ the next-period draws the expectation integrates, and the predetermined output is $\bar{a} = p\,\bar{x}$. The aggregate path $(\mathcal{Z}, \mathcal{Z}')$ is an argument held fixed throughout, so the derivatives below are taken in the three function slots alone.

On each fixed regime $\mathcal{F}$ is continuously Fréchet differentiable (proof of Proposition 1), so by Proposition 14 its Fréchet derivative in any slot coincides with the directional (Gâteaux) derivative of Definition 8 and may be computed as the scalar limit

$$D_u \mathcal{F} \cdot du \;=\; \lim_{\alpha \to 0} \frac{\mathcal{F}(u + \alpha\, du) - \mathcal{F}(u)}{\alpha},$$

a one-dimensional differentiation in $\alpha$. I write $F_x, F_{x^+}, F_Y$ for the partial Jacobians of the pointwise residual $F$ at the unperturbed arguments, and use that the conditional expectation is linear and commutes with the $\alpha$-limit, the latter by dominated convergence since on the fixed regime the relevant policy derivative is a bounded $L^\infty$ object (the kink jump, no Dirac).

**The relevant-aggregates slot.** Perturbing $\bar{Y} \to \bar{Y} + \alpha\, d\bar{Y}$ moves only the fifth argument of $F$, and linearly, so

$$D_{\bar{Y}} \mathcal{F} \cdot d\bar{Y} \;=\; \lim_{\alpha \to 0} \frac{F\big(\ldots, \bar{Y}(\mathcal{Z}) + \alpha\, d\bar{Y}(\mathcal{Z})\big) - F\big(\ldots, \bar{Y}(\mathcal{Z})\big)}{\alpha} \;=\; F_Y\, d\bar{Y}.$$

**The next-period policy.** The slot $\bar{x}^+$ enters $\mathcal{F}$ only through the expectation, and linearly. By linearity of the expectation, $\mathbb{E}[\bar{x}^+ + \alpha\, d\bar{x}^+ \mid \theta, a, \mathcal{Z}] = \mathbb{E}[\bar{x}^+ \mid \theta, a, \mathcal{Z}] + \alpha\, \mathbb{E}[d\bar{x}^+ \mid \theta, a, \mathcal{Z}]$, and the evaluation point $\bar{a} = p\,\bar{x}$ is unchanged because $\bar{x}$ is held fixed. The difference quotient of $F$ in its fourth argument then gives

$$D_{\bar{x}^+} \mathcal{F} \cdot d\bar{x}^+ \;=\; F_{x^+}\, \mathbb{E}\big[d\bar{x}^+ \,\big|\, \theta, a, \mathcal{Z}\big].$$

**The current policy.** This is the only slot that enters $\mathcal{F}$ twice: directly as the third argument of $F$, and through the predetermined output $\bar{a} = p\,\bar{x}$ that locates the next-period policy inside the expectation. Perturbing $\bar{x} \to \bar{x} + \alpha\, d\bar{x}$ shifts the third argument by $\alpha\, d\bar{x}(\theta, a, \mathcal{Z})$ and the evaluation point of $\bar{x}^+$ by $\alpha\, p\, d\bar{x}(\theta, a, \mathcal{Z})$, a quantity that does not depend on the integration variables $(\theta', \mathcal{E}')$. Taking the difference

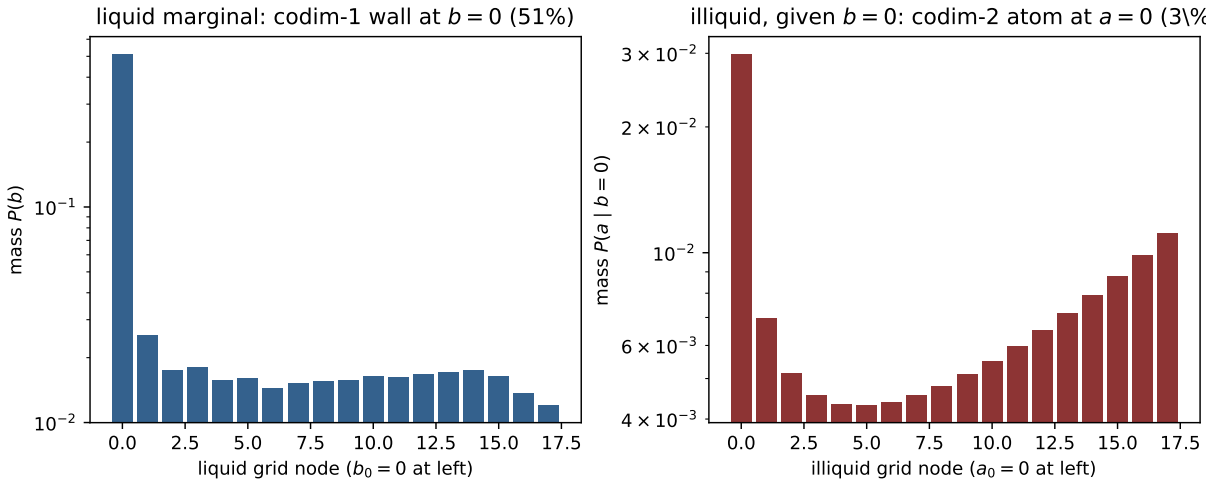


Figure 6: The two-asset stationary distribution carries atoms of two codimensions. *Left:* the liquid marginal, with a codimension-one atom at the borrowing constraint $b = 0$ holding about half the population (log scale). *Right:* the illiquid distribution conditional on $b = 0$, with a further atom at $a = 0$, the codimension-two corner where the poorest households hold neither asset. This corner atom is what the dual $V'$ must admit, and what a first-order test space below $C^0$ would miss.

quotient and letting $\alpha \to 0$, the chain rule across the third and fourth arguments of $F$ gives

$$D_{\bar{x}}\mathcal{F} \cdot d\bar{x} \;=\; F_x\, d\bar{x} \;+\; F_{x^+} \lim_{\alpha \to 0} \frac{\mathbb{E}[\bar{x}^+(\theta', \bar{a} + \alpha\, p\, d\bar{x}, \mathcal{Z}') \,|\, \theta, a, \mathcal{Z}] - \mathbb{E}[\bar{x}^+(\theta', \bar{a}, \mathcal{Z}') \,|\, \theta, a, \mathcal{Z}]}{\alpha}.$$

Pulling the limit inside the expectation and using that $p\, d\bar{x}$ is constant in $(\theta', \mathcal{E}')$, the inner limit is the directional derivative of $\bar{x}^+$ in its predetermined-state argument along $p\, d\bar{x}$,

$$\lim_{\alpha \to 0} \frac{\bar{x}^+(\theta', \bar{a} + \alpha\, p\, d\bar{x}, \mathcal{Z}') - \bar{x}^+(\theta', \bar{a}, \mathcal{Z}')}{\alpha} \;=\; D_a \bar{x}^+(\theta', \bar{a}, \mathcal{Z}')\; p\, d\bar{x},$$

so that

$$D_{\bar{x}}\mathcal{F} \cdot d\bar{x} \;=\; \big(F_x + F_{x^+}\, \mathbb{E}\big[D_a \bar{x}^+ \,\big|\, \theta, a, \mathcal{Z}\big]\, p\big)\, d\bar{x}.$$

This is the load-bearing identity: the current policy reaches the expectation only through the bound it sets on next period's state, and that single channel is the term $F_{x^+}\, \mathbb{E}[D_a \bar{x}^+ \mid \theta, a, \mathcal{Z}]\, p$. Assembling $J = D_{\bar{x}}\mathcal{F}(\bar{x}^*, \bar{x}^*, \bar{Y}^*)$ from it and inverting is the content of Proposition 1.

# E The order-$k$ decay classes: the derivation at $m = 2$

This appendix derives the order-$k$ decay hypothesis of Proposition 12 at $m = 2$, for the three families of terms the source $\Xi^{(2)}$ actually contains, and records exactly what the derivation consumes beyond the standing assumptions. Throughout, $\rho$ and $\lambda$ are the anticipation and propagation rates of Proposition 6, $r(u) = \lambda_m^u$ for $u \geq 0$ and $r(u) = \rho^{-u}$ for $u < 0$ the one-sided profile, read at the per-order propagation rate $\lambda_m = \lambda^{1/m}$, and $\mathcal{K}_m$ the kernel class of the hypothesis; the third family is what forces the enlargement $\lambda \to \lambda^{1/m}$, and the diagonal-profile paragraph below derives it and shows it exact.

**The class algebra.** One pointwise fact does the outer work: the corner profile is dominated by the interior one,

$$\lambda^t \rho^s \;\leq\; r(t - s) \qquad (t, s \geq 0),$$

since with $r$ read at the first order's own $\lambda$ the ratio is $(\lambda\rho)^s \leq 1$ for $t \geq s$ and $(\lambda\rho)^t \leq 1$ for $t < s$, and enlarging the propagation rate to $\lambda_m \geq \lambda$ only raises the right side, so the domination holds a fortiori. Three consequences follow at once: a product of $\mathcal{K}_1$-bounds lands in the $\mathcal{K}_2$-bound, the cross terms eaten

by the domination; a one-period shift costs a constant, $r(u \pm 1) \le \max(\lambda^{-1}, \rho^{-1})\, r(u)$; and the horizon convolutions close,

$$\sum_{\sigma \ge 0} \rho^{\sigma}\, r(t + \sigma - s) \;\le\; C\, r'(t - s),$$

with $r'$ built from any rates strictly between the originals and one, the polynomial factors of the middle ranges absorbed by reading the rates strictly inside their intervals; every statement below is read at such interior rates, exactly as Proposition 6 fixes $\rho$ and $\lambda$, a convention the per-order $\lambda_m$ inherits. The aggregate-equation family of $D^2\mathcal{G}$ is settled by this alone: each row reads the two kernel paths at finitely many dates within one period of the row date, through bounded steady-state coefficients.

**The two graded clauses.** The heterogeneous families consume two strengthenings of clause 2 of Assumption 2, stated here, their qualitative frame established at the grades of Section 5.2: (i) the state derivatives of the policy-response coefficients decay with the horizon, $\|D_a \times_\sigma\|_\infty \le C\rho^\sigma$, per regime; (ii) the mixing of clause 2 extends to the coefficient family and one grade up,

$$\big\|\mathcal{K}^m \times_\sigma - \langle \bar\Omega^*, \times_\sigma \rangle\big\|_\infty + \big\|\nabla_a\big(\mathcal{K}^m \times_\sigma\big)\big\|_\infty \;\le\; C\,\lambda^m \rho^\sigma, \qquad \big\|\nabla_a^2\big(\mathcal{K}^m \bar x^* - \langle \bar\Omega^*, \bar x^* \rangle\big)\big\|_\infty \;\le\; C\,\lambda^m.$$

Clause 2 as stated covers the first term of the first display at $\bar x^*$ only; the running scalar illustration is silent on both clauses, having no state and no distribution side, which is where they bite.

**The register mass at a horizon.** The carriers of the singular part $D_k$ multiply with the horizon, a $\Gamma$-layer splitting into its per-regime pre-images under $\Lambda$, so their surface area grows; their total mass does not. One lemma records this, the register-side companion of the stationarity count in Lemma 6.

**Lemma 13** (Total layer mass decays at the propagation rate)**.** *Fix the grade integrability $p$ of Definition 7 large enough, any $p > d_a$ admissible, the paired iterates being Lipschitz per regime. Then the total layer mass at horizon $t$ contracts,*

$$\sum_{c} \big\|w_c\big\|_{L^{p'}(\mathit{carrier}_c)} \;\le\; C\,(\lambda')^t, \qquad \lambda' \in (\lambda, 1),$$

*the sum over the carriers created by $t$ transfers, with $\lambda'_m$ in place of $\lambda'$ at order $m$; the growing carrier count does not defeat the mixing decay.*

*Proof.* A single-valued push $\bar a^*$ carries a layer $w\,\delta_\Gamma$ on $\Gamma$ to a layer $w'\,d\Gamma'$ on its image, with $w' \circ \bar a^* = w / J_\tau$ for $J_\tau$ the tangential Jacobian, so the total mass is invariant,

$$\int |w'|\, d\Gamma' \;=\; \int |w' \circ \bar a^*|\, J_\tau\, d\Gamma \;=\; \int |w|\, d\Gamma :$$

a larger image carries a proportionally smaller weight. The $R$ Markov branches split the mass with probabilities $\pi(\theta' \mid \theta)$ summing to one, so the branch count does not enter the summed mass either. The register prices weights in $L^{p'}$, whose per-step distortion is $(J_\tau)^{-1/p}$ (from $\|w'\|^{p'}_{L^{p'}} = \int |w|^{p'} J_\tau^{1-p'}\, d\Gamma$), and, the branch probabilities included,

$$\Theta \;=\; \sum_{\theta'} \pi(\theta' \mid \theta)\,(J_\tau)^{-1/p} \;\le\; \big(L^{d_a - 1}/\epsilon\big)^{1/p} \;\xrightarrow[p \to \infty]{}\; 1,$$

by the clause-1 and clause-2 bounds $J_\tau \ge \epsilon / L^{d_a - 1}$ of Lemma 7, Step 5; the nested grades admit any $p > d_a$, so $\Theta$ is brought as close to one as wanted. The weights are jumps of centered iterates, contracted by $\lambda$ per transfer through clause 2 of Assumption 2 and clause (ii) above, so one transfer multiplies the summed $L^{p'}$ mass by at most $\lambda\Theta$; choosing $p$ with $\Theta < \lambda^{-1/2}$ gives $\lambda\Theta \le \lambda^{1/2} =: \lambda' < 1$ and the stated bound. The orbit atoms and the first-order layers of Assumption 1 sit on fixed carriers and do not proliferate, so the $\Gamma$-layers are the only family the count concerns. □

**The three families.** The bilinear part of the second derivative of a cross-sectional row splits as

$$d^2 \int \bar x_t\, d\bar\Omega_t \;=\; \big\langle \bar\Omega^*,\; d^2 \bar x_t \big\rangle \;+\; 2\,\big\langle d\bar\Omega_t,\; d\bar x_t \big\rangle \;+\; \big\langle d^2 \bar\Omega_t,\; \bar x^* \big\rangle,$$

the parts linear in the unknown order-two path sitting inside $\tilde{J}_X$ on the left. For the first, unrolling the order-two horizon recursion of Section 5.2 backward bounds the composite coefficients by $\|{\times_{\sigma_1\sigma_2}}^{(2)}\| \le C\rho^{\sigma_1}\rho^{\sigma_2}$, clause (i) pricing the $D_a{\times_\sigma}$ factors in the source, and the horizon sums close by the convolution bound. For the second, the increments mirror Steps 1–2 of Proposition 6's proof, a source is mass-free so the pairing centers, with the coefficient $\times_{\sigma'}$ in the test-function slot, which is what clause (ii)'s first display prices; the date sums are the first-order Step-3 convolutions run twice. For the third, the four-term law of motion of Section 5.2 at $m = 2$ contributes displacements of displacements, whose pairings read $\nabla(\mathcal{K}^m\bar{x}^*)$ through their interior parts, clause 2 as stated, and, through the dipole entries of the displaced first-order response, the weights $\nabla_a\hat{a}$, clause (i), and second-derivative iterates, clause (ii)'s second display, the mixing running on the composite of a coefficient slice with a gradient iterate; and one Hessian pairing,

$$\big\langle \Lambda^m\, D^2_{\bar{a}^*}\Lambda(a,a)\,\bar{\Omega}^*,\ \bar{x}^*\big\rangle \;=\; \int \nabla^2\big(\mathcal{K}^m\bar{x}^*\big)\big[\,a,\ a\,\big]\; d\bar{\Omega}^*,$$

which is what clause (ii)'s second display prices; the ledger terms carry layer weights that are jumps of the same objects, bounded by twice the sup-norms already controlled, their traces priced by the trace clause of Assumption 3, and their uniformity near intersections of distinct boundaries $\Gamma_b$ is supplied by the uniformity-plus-angle clause at the interface-calculus paragraph of Section 5.2. Each family therefore lands in $\mathcal{K}_2$ read at the enlarged rate $\lambda_2 = \lambda^{1/2}$, the next paragraph deriving why the enlargement is forced and exact, and the general order runs the same three moves by induction, the clauses becoming the graded ladder one derivative per order and the rate $\lambda_m = \lambda^{1/m}$.

**The diagonal profile, and the exact rate enlargement.** The third family fixes the class's propagation rate, and it does so exactly. Its deposits live at dates $\tau \le \min_i s_i$, every coefficient slot being anticausal, and are read after a single relaxation chain $\Lambda^{t-1-\tau}$; substituting $w = \min_i s_i - \tau$, the $\tau$-sum of the all-at-once deposit, the order-$m$ sibling of the Hessian pairing above, evaluates on the orthant of positive gaps $u_i = t - s_i$ to a $(\lambda\rho^m)$-geometric series times the profile

$$\lambda^{\max_j u_j}\ \prod_{i=1}^{m}\rho^{(\max_j u_j)-u_i},$$

one mixing chain from the latest deposit and $\rho$-factors in the spreads between the gaps. On the all-equal diagonal $u_1 = \cdots = u_m = u$ the profile is $\lambda^u$, while the product bound of the class demands $\lambda^{mu}$: after the last shock date the order-$m$ response relaxes through one transfer-operator chain at the mixing rate, so it decays like the first-order response and not like its $m$-th power, and no fixed rate below one places it in the product class. The root enlargement does, and it is exact. The claim $\lambda^M\prod_i\rho^{M-u_i} \le \lambda^{\sum_i u_i}_m$ with $M = \max_j u_j$ reads, in units of $\log\lambda$,

$$M\ +\ \gamma\sum_{i=1}^{m}\big(M-u_i\big)\ \ge\ \frac{1}{m}\sum_{i=1}^{m}u_i,\qquad \gamma=\frac{\ln\rho}{\ln\lambda}>0,$$

and the maximum dominates the mean, the $\gamma$-term only helping; the diagonal shows no rate below $\lambda^{1/m}$ serves. Nested deposits, an order-$m'$ response displaced inside an order-$m$ one with $m' < m$, close by induction: the deposit's own chain covers every deposit-side gap, $(t-s_i)_+ \le t-\tau$ there, the inner response's rate $\lambda_{m'}$ covers the remaining gaps through $(t-s_i)_+ \le (t-\tau)+(\tau-s_i)_+$, and $1/m \le 1/m'$; the one chain $\lambda^{t-\tau}$ is shared by up to $m$ propagation gaps, which is where the $1/m$ originates. Products across orders compose: a $D^p\mathcal{G}$ contraction of kernels of orders $m_1+\cdots+m_p = m$, each in its own class, lands in $\mathcal{K}_m$, since $\sum_i\Sigma_i/m_i \ge \big(\sum_i\Sigma_i\big)/m$ with $\Sigma_i$ the sum of the $i$-th factor's positive gaps. The anticipation side, the corner parts, and the correction terms keep the first order's own rates, the domination $\lambda^t\rho^s \le r(t-s)$ holding a fortiori at the enlarged one.

**The edge kernel of the corrections.** Write ${\Xi^{(m)}}_t[\,\cdot,\dots,\cdot\,]$ for the source read as an $m$-linear functional of its first-order ingredient columns, so that ${\Xi^{(m)}}_t(s_1,\dots,s_m) = {\Xi^{(m)}}_t[\,K_{\cdot,s_1},\dots,K_{\cdot,s_m}\,]$ with $K_{\cdot,s}$ the first-order column at shock date $s$, and split each column into its diagonal limit and its correction, $K_{t,s} = H_{t-s} + c_{t,s}$ with $\|c_{t,s}\| \le C\lambda^t\rho^s$ (Proposition 6, item 3). Expanding by multilinearity gives one term per subset $A \subseteq \{1,\dots,m\}$ of slots fed the correction, $H$ in the rest, the all-diagonal slot read

against the whole-line gap kernel ${\Xi^{(m)}}_{\mathrm{gap}}$, the object Step 2 of Proposition 12's proof transforms,

$$
\begin{aligned}
{\Xi^{(m)}}_t\big(s_1,\dots,s_m\big) - {\Xi^{(m)}}_{\mathrm{gap}}\big(t-s_1,\dots,t-s_m\big) \;=\; & \sum_{\emptyset\neq A\subseteq\{1,\dots,m\}} {\Xi^{(m)}}_t\big[\,g_1^A,\dots,g_m^A\,\big] \\
& + \Big({\Xi^{(m)}}_t\big[\,H_{\cdot-s_1},\dots,H_{\cdot-s_m}\,\big] \;-\; {\Xi^{(m)}}_{\mathrm{gap}}\big(t-s_1,\dots,t-s_m\big)\Big), \\
g_i^A \;=\; c_{\cdot,s_i}\ (i\in A), \qquad & H_{\cdot-s_i}\ (i\notin A),
\end{aligned}
$$

the first sum the forms with the correction $c$ in the slots of $A$, the parenthesized term the whole-line reads at the negative dates the half-line form pins at the steady state, each $A$-term bounded by $C\,\lambda^{|A|\,t}\prod_{i\in A}\rho^{s_i}\prod_{i\notin A} r(t-s_i)$, one factor $\lambda^t$ per correction slot, in class by applying the domination slot by slot over $A$; the per-slot factors matter beyond $m=2$, where a single $\lambda^t$ across a mixed $A$ with $2\le|A|<m$ would no longer be dominated by the class bound. Pushed through the four correction terms of Proposition 12, item 3: the pre-date-0 tail sums $\sum_{v<0}\|j_{X,t-v}\|\prod_i\rho^{\,s_i-v}\le C\,\lambda^t\prod_i\rho^{\,s_i}$, the full corner profile; the Hankel product, the source's own edge, and the resolvent term run sums of the form $\sum_v\rho^v\prod_i r(v-s_i)$, which evaluate to $C\,\rho^{\max_i s_i}$ at interior rates, so those three terms carry $C\,\lambda^t\rho^{\max_i s_i}$, geometric in the row and in the latest shock date. That weaker column profile is all the downstream summability uses.

# F A notation and code dictionary

This appendix is for the reader who wants to dig into the three codebases. The first-order dynamics are one object computed three ways, and the three differ less in mathematics than in where they sit on the ladder from continuous operator to running code.

Bhandari et al. [2023] work at the top of the ladder, with continuous linear operators on functions, the most general level because nothing depends on the discretization basis, at the cost of eventually building large sparse matrices. Auclert et al. [2021] work at the bottom, with finite vectors and matrices on the grid throughout, gaining the speed of a specialized $P^1$ code that never forms those matrices explicitly. Dynare 7 keeps the functionals of the top level for the mathematics and discretizes only at the end, taking the heterogeneous block of Bhandari et al. [2023] and the aggregate assembly of Auclert et al. [2021].

The one trap for a code reader is that the same symbol can name an object at two rungs. Bhandari et al. [2023]'s operators $\mathcal{M}$, $\mathcal{L}^{(a)}$, $\mathcal{I}^{(a)}$ are continuous (their §3.3, Lemma $4^{HA}$); the variables called `M` and `L` in their Julia code are discrete matrices that approximate them, and converge to them only as $h\to 0$ (Section 4.5). I split the dictionary into the two rungs to keep them apart: Table 2 collects the continuous objects, Table 3 their discrete incarnations.

| This paper | Bhandari et al. [2023] | Auclert et al. [2021] | Object |
|---|---|---|---|
| $\mathfrak{x}$ | x (Lemma 3) | user backward step | implicit-function backward map |
| $\times_s$ | $\times_s$ | $dy^T_{T-s}$ | policy-response coefficient ($\times_s$ adopted from Bhandari et al.), $-J^{-1}F_{x_+}\mathbb{E}[\,\cdot\,]$ |
| $\mathcal{K}$ | $\mathbb{E}_\epsilon[\,\cdot\mid\theta,a]$ | – | Koopman (conditional expectation) |
| $\Lambda$ | $\bar\Lambda$ | – | transfer (Perron–Frobenius) |
| M | $\mathcal{M}$ | – | displacement, $\Lambda(\,\cdot\,\bar\Omega^*)$ (Prop. 3) |
| – | $\mathcal{L}^{(a)}$ | – | velocity transport, $\Lambda\big(\sum_j \partial_{aj}\bar a^{*i}\ \ \cdot^{\,j}\ \big)$ (Lem. 4) |
| – | $\mathcal{I}^{(a)}$ | – | aggregation pairing $\sum_j\langle\cdot^{\,j},\partial_{aj}\bar x^*\rangle$ (Lem. 5) |
| $\delta_{(\theta,u)}[v]$, $E=V'\oplus D$ | atoms inside $\hat\omega_t$ (Assumption 2) | – | dipole, the moving singular mass |
| $\mathcal{K}^t\bar x^*$ | – | $\mathcal{E}_t$ (Def. 1, eq. (23)) | expectation vectors, Koopman iterates of the policy |

Table 2: Continuous objects (the mathematics). Bhandari et al. [2023]'s velocity-field operators $\mathcal{L}^{(a)}$, $\mathcal{I}^{(a)}$ carry no symbol of this paper's own: they are cited for comparison, inside the identification passages of Sections 4.3 and 4.4, and the identifications are Lemma 4 and Lemma 5.

Two notes for the code reader. First, Bhandari et al. [2023]'s histogram law of motion (their distribution-update routine, `FirstOrderApproximation.jl:536`) is written $\hat\omega_{t+1}=L\,\hat\omega_t+M\,\hat a_t$ with the *discrete* matrices `L` and `M`: it is their Appendix C, not the continuous Lemma 4 operator law, and a comment in the code flags that it tracks the change in the histogram, not in the cumulative distribution function. Second, the identity $\mathsf{M}=\mathcal{M}$ of Bhandari et al. [2023] is literal in their code, `FO.M =`

| This paper | Bhandari et al. [2023] (Julia) | Auclert et al. [2021] (ecta toolkit) | Dynare 7 | Discretizes |
|---|---|---|---|---|
| $J$ | `f= -J^-1 (FOA:171)` | finite difference | `f3 (compute_impulse_responses)` | $D_{\bar{x}}\mathcal{F}$ |
| $\times_s$ recursion | `compute_Lemma3! (:183)` | `backward_iteration_fakenews (:562)` | `compute_impulse_responses` | the backward recursion |
| $\mathsf{M}^h$ (curlyD) | `FO.M=Λ*(Φ'.*ω̄) (:228)` | `forward_step_shock (:740)` | `curlyDs (solve)` | $-\nabla\cdot(\mathsf{M}\,\cdot)$ |
| – | `FO.L (:219)` | in forward step | in transition | $\mathcal{L}^{(a)}$, the slope-weighted transition matrix |
| curlyY | `compute_Proposition1! (:269)` | `jac (:301)` | `curlyYs (solve)` | $\langle\bar{\Omega}^*, \times_s\rangle$ |
| curlyP | `IL (:250)` | `curlyPs (:596)` | `expectations (compute_expected_pol)` | $\mathcal{K}^t\bar{x}^*$; Lem. 5 reads it as $\mathcal{I}^{(a)}(\mathcal{L}^{(a)})^t$ |
| fake-news $F$ | `Ft=Et*Dt (ABRS_KS:174)` | `build_F (:611)` | `transition_matrices (compute_transition_matrices.f08)` | diagonal increments $\mathcal{F}_{t,s}$ |
| $J_{t,s}$ | `Jt (ABRS_KS:186)` | `J_from_F (:620)` | `heterogeneous_jacobian (compute_heterogeneous_jacobian)` | block quasi-Toeplitz $J$ (Props. 5 and 6) |
| `dr.G` | `solve_Xt! (:315)` | host model | `G (solve)` | $\partial\widehat{X}/\partial\mathcal{E}$ |
| lottery | `find_stationarydistribution_λ! (AiyagariModel.jl:270)` | `forward_step_1d (utils.py:208)` | `mat.d.ind, mat.d.w (load_steady_state)` | interpolation push-forward |

Table 3: Discrete incarnations (the code). Unqualified line numbers are into `FirstOrderApproximation.jl` (Bhandari et al. [2023], repo https://github.com/apb-umn/bbeg, commit `015c363fdb6f9ac4ae2524717c22528c4622411d`) and `het_block.py` (Auclert et al. [2021], the `ecta200326-sup-0002` replication package); the other files are named in place. All pointers re-checked against the shipped sources.

`Λ*(Φ'.*ω̄)` being the discretization of $\int \bar{\Lambda}\,\bar{\Omega}^*(\cdot)$; the continuous identity is proved in Section 4.3, and the rate at which the discrete object approaches it is recorded in Section 4.5.

The second-order objects of Section 5.2 have their column in the same repository, in `SecondOrderApproximation.jl`. The state-slot Dirac of the first second-derivative display after Corollary 3, in its $d_a = 1$ reduction, is their $\bar{x}_{aa}$ decomposition, computed in the routine of that name; the order-two policy recursion after Lemma 8 is their Lemma-3 second-order assembly, the pre-factored first-order operator `f` applied to a bilinear source; the boundary sensitivities of Corollary 3 and the $C_{1,1}$ weight are computed jump by jump in their kink routine; the four-term law of motion of Section 5.2 is their Lemma-4 second-order block; and the same-operator closures of Propositions 12 and 13 reuse, in their code as here, the first-order factorizations, the aggregate one solving with the stored first-order decomposition and the risk block solving one more linear system in the same matrices. Line numbers shift with their repository and are omitted; the routine names above are stable anchors.

# G Multilinear interpolation in Dynare 7

Every discrete operator of Section 4.5 is one of two things read on the $P^1$ basis: the interpolant $\mathcal{I}_h$ of Definition 16, or the transfer operator $\Lambda^d = \mathcal{I}_h^*\Lambda$ of Lemma 11. Both, and the derivative stencil $\mathsf{M}^h$ that carries the displacement, reduce to one computational primitive: evaluating a tensor-product hat at a landing point, and in most call sites its gradient too. This appendix records how the toolkit's Fortran kernel evaluates that primitive, since the recipe is not the naive one and the same routine reappears seven times across the code. Nothing here is new mathematics; it is a description of the implementation for the reader of the sources under `mex/sources/heterogeneity/`.

**The object.** On the $d_a$-dimensional box of predetermined coordinates $a \in \mathcal{A}$, the $P^1$ interpolant is the tensor product of the one-dimensional hats of Definition 16, the bilinear $Q^1$ element of Section 4.5 at $d_a$ coordinates. Bracketing the landing point in each coordinate $k$ leaves a lower node and the hat weight $w_k \in [0,1]$ carried on it, so $\mathcal{P}^{i_k}$ takes the value $w_k$ at the landing point and its upper neighbour $\mathcal{P}^{i_k+1}$ the value $1-w_k$. Indexing a corner of the bracketing cell by the set $U \subseteq \{1,\dots,d_a\}$ of coordinates taken at their *upper* node, the interpolant of an array $v$ carried on the grid is

$$\mathcal{I}_h v \;=\; \sum_{U\subseteq\{1,\dots,d_a\}} \beta_U\, v_{a_U}, \qquad \beta_U \;=\; \prod_{k\notin U} w_k \prod_{k\in U}(1-w_k),$$

with $a_U$ the corner node and $\beta_U$ its tensor-product hat weight, the $\prod_i \mathcal{P}^i$ of Section 4.5 written corner by corner. The sum has $2^{d_a}$ terms.

**Bracketing one coordinate.** The one-dimensional bracket is a binary search on the sorted grid (`bracket_linear_weight` in `utils/interpolation.f08`): for a query $a_k^q$ it returns the lower index $i_k$

**Algorithm 1** Multilinear interpolation and its gradient at one landing point

**given** weights $w_k$, reciprocal widths $1/h_k$, ratios $r_k^{\uparrow}, r_k^{\downarrow}$, lower-corner index $m$, strides
$\beta \leftarrow \prod_k w_k$; $\quad d\beta_k \leftarrow -(1/h_k)\,\beta/w_k$ for each $k$ ▷ value and gradient at the lower corner
$U \leftarrow \emptyset$; $\quad$ accumulate $\beta\, v_m$ into the value and $d\beta_k\, v_m$ into the $k$-th gradient
**for** $t = 1, \ldots, 2^{d_a} - 1$ **do**
  $k_f \leftarrow k_f(t)$
  **if** $k_f \notin U$ **then** ▷ lower → upper
    $m \leftarrow m + \text{stride}_{k_f}$; $\quad U \leftarrow U \cup \{k_f\}$
    $\beta \leftarrow \beta\, r_{k_f}^{\uparrow}$; $\quad d\beta_k \leftarrow d\beta_k\, r_{k_f}^{\uparrow}$ for $k \neq k_f$
  **else** ▷ upper → lower
    $m \leftarrow m - \text{stride}_{k_f}$; $\quad U \leftarrow U \setminus \{k_f\}$
    $\beta \leftarrow \beta\, r_{k_f}^{\downarrow}$; $\quad d\beta_k \leftarrow d\beta_k\, r_{k_f}^{\downarrow}$ for $k \neq k_f$
  **end if**
  $d\beta_{k_f} \leftarrow -\,d\beta_{k_f}$ ▷ the gradient in the flipped coordinate changes sign, never scales
  accumulate $\beta\, v_m$ into the value and $d\beta_k\, v_m$ into the $k$-th gradient
**end for**

with $\boldsymbol{a}_{[i_k]} \le a_k^q < \boldsymbol{a}_{[i_k+1]}$ and the weight

$$w_k \;=\; \frac{\boldsymbol{a}_{[i_k+1]} - a_k^q}{\boldsymbol{a}_{[i_k+1]} - \boldsymbol{a}_{[i_k]}},$$

clamped to the end cell when the query falls outside the grid. The reciprocal cell width $1/h_k$, with $h_k = \boldsymbol{a}_{[i_k+1]} - \boldsymbol{a}_{[i_k]}$, is kept for the gradient below.

**The corner walk.** Forming the $2^{d_a}$ products $\beta_U$ independently costs $d_a$ multiplications each. The kernel instead walks the corners so that consecutive corners differ in one coordinate, and updates the weight by a single ratio at each step. The neighbouring weights satisfy

$$\frac{\beta_{U\cup\{k\}}}{\beta_U} \;=\; \frac{1-w_k}{w_k} \;=:\; r_k^{\uparrow}, \qquad \frac{\beta_{U\setminus\{k\}}}{\beta_U} \;=\; \frac{w_k}{1-w_k} \;=:\; r_k^{\downarrow},$$

so moving one coordinate from its lower node to its upper node multiplies the running weight by $r_k^{\uparrow}$, and moving it back multiplies by $r_k^{\downarrow}$; the two ratios are tabulated once (`compute_coefficient_updates` in `utils/gray_code.f08`). The order that visits every corner changing one coordinate per step is the binary-reflected Gray code $g(t) = t \oplus (t \gg 1)$, whose successive codes differ in the single bit $k_f(t) = \text{trailz}\big(g(t) \oplus g(t-1)\big) + 1$, the coordinate to flip at step $t$ (`generate_flip_indices`). The corner node's position in the flattened array moves by the stride of that coordinate, $\pm\,\text{stride}_{k_f}$, where $\text{stride}_1 = 1$ and $\text{stride}_k = \text{stride}_{k-1}\,\text{dim}_{k-1}$ (`compute_strides`); the idiosyncratic shock $\theta$ is the innermost axis, so on the joint shock-and-state array every stride is scaled by the number of shock nodes and the walk moves only along the continuous coordinates. Algorithm 1 is the traversal, in its value-and-gradient form.

Each of the $2^{d_a} - 1$ steps costs one multiplication and one integer shift, against the $d_a$ multiplications a fresh product would take, and it reuses the corner node it just visited rather than recomputing an index.

**Boundaries where a weight is zero or one.** When a landing point falls on a node face, some $w_k$ is 0 or 1 and the ratios $r_k^{\uparrow}, r_k^{\downarrow}$ divide by zero. Those coordinates are pulled out of the multiplicative walk and tracked by a counter. A coordinate with $w_k = 1$ contributes only through its lower node, one with $w_k = 0$ only through its upper node (`is_hard_one`, `is_hard_zero`); the low-corner weight $\beta = \prod_{k:\,0<w_k<1} w_k$ skips them. The counter $z$ holds the number of these coordinates currently on the wrong node, initialized to the number with $w_k = 0$ (all wrong at the all-lower corner) and moved by $\pm 1$ whenever the walk flips such a coordinate; a corner is accumulated only when $z = 0$. A point mass landing on a corner node where several bounds bind is then represented exactly, with no division, matching the exact representation claimed in Section 4.5.

**The gradient in the same pass.** The distribution stencil $\mathsf{M}^h$ of Section 4.5 needs the hat gradient at the landing point, the $\nabla\mathcal{P} = \pm 1/h$ that stands in for the weak derivative. Because $\beta_U$ depends on

Table 4: The multilinear-interpolation kernel across the toolkit

| Routine | Realizes | Section |
|---|---|---|
| *Value walk* (one running weight $\beta$) | | |
| `interpolate` | array evaluation at landing points | Def. 16 |
| `aggregate_residuals_tensor` | steady-state push-forward $\Lambda^d$ (scatter) | Lem. 11, H.3 |
| `compute_expected_y` | conditional expectation $\mathcal{I}_h\mathcal{K}$ (gather) | 4.5 |
| `compute_Phi_tilde_e` | sparse tabulation of $\mathcal{I}_h\mathcal{K}$ with $\boldsymbol{\pi}$ | 4.5 |
| *Value-and-gradient walk* ($\beta$ and the $d_a$ accumulators $d\beta_k$) | | |
| `set_het_input` | backward step: expected policy and its slopes | H.1 |
| `compute_expected_dx` | expected policy slope $\mathbb{E}[\partial_a \bar{x}^* \mid \theta, a]$ | 4.5 |
| `compute_curlyDs` | displacement stencil $\mathsf{M}^h$ | H.2 |

the landing coordinate $a_k$ only through the single factor $w_k$ or $1 - w_k$, and $\partial_{a_k} w_k = -1/h_k$,

$$\partial_{a_k}\beta_U \;=\; \mp\,\frac{1}{h_k}\,\beta_U^{\setminus k}, \qquad \beta_U^{\setminus k} \;=\; \prod_{\ell\notin U,\,\ell\neq k} w_\ell \prod_{\ell\in U,\,\ell\neq k} (1-w_\ell),$$

the sign being $-$ when coordinate $k$ sits at its lower node and $+$ at its upper node. The derivative accumulator therefore runs the same walk with coordinate $k$'s hat factor replaced by $\mp 1/h_k$: it starts at $-(1/h_k)\,\beta/w_k$, scales by $r^{\uparrow}, r^{\downarrow}$ under flips of the other coordinates exactly as the value does, and only changes sign when the walk flips coordinate $k$ itself (lines 13 and 2 of Algorithm 1). Carrying the $d_a$ accumulators alongside $\beta$ returns the interpolated value and its full gradient from one traversal, each gradient direction keeping its own hard-boundary counter. This is the array the code assembles as the hat-derivative stencil of Section 4.5, a bounded jump and never an impulse.

**Factoring the shock expectation.** The idiosyncratic shock is discrete and carries no interpolation, so wherever the primitive computes a conditional expectation the Markov kernel $\boldsymbol{\pi}$ is applied once, by a dense matrix product, outside the corner walk. The gather form (evaluating an array at the landing points, the discrete Koopman reading $\mathcal{I}_h\mathcal{K}$) premultiplies the data by $\boldsymbol{\pi}$ before the walk; the scatter form (splitting mass onto the bracketing nodes, the push-forward $\Lambda^d$) applies $\boldsymbol{\pi}^\top$ after. The two are the one transition matrix of Section 4.5 read along its rows or its columns, the same walk transposed, and neither lets the shock kernel into the inner $2^{d_a}$ loop.

**Where the kernel is used.** The primitive appears in seven places, in three families. Table 4 lists them against the continuous object each realizes; all live under `mex/sources/heterogeneity/`, and all draw on the shared primitives of `utils/gray_code.f08` through the `gray_code_cache` that stores the flip sequence, strides, ratios, and boundary flags.

The scatter and gather rows are the same walk read two ways, and the tabulation row emits the walk's per-corner weights as the sparse triplets of a matrix assembled once rather than applying them to data. The gradient family adds the $d_a$ accumulators of Algorithm 1 to the value walk and nothing else.

# H Backing out the responses of policy functions and the histogram

The three MATLAB helpers below reconstruct the cross-sectional responses of the "Backing out the heterogeneous responses" paragraph of Section 4.5 from the outputs of `heterogeneity_solve`. They are post-processing functions: none solves the model, each unpacks and recombines what the solver already computed. All three assume Dynare is on the path (`addpath dynare/matlab;`), so that the MEX `compute_expected_dx` and the package internals are reachable, and all three read subfields the interface warning of Section 4.5 marks as internal and liable to move, chiefly `oo_.heterogeneity.dr.G`, `oo_.heterogeneity.dr.curlyDs`, and the steady-state lottery data `oo_.heterogeneity.mat.d.{ind,w,dims}` and `mat.Mu`. The listings are the versions shipped for the Ottawa Dynare workshop.

## H.1 The policy response

The function `het_policy_responses` returns the cross-section $\hat{x}_t(\theta, a)$ on the joint shock-state grid. It recomputes the policy-response coefficients $\times_s$ (the array `impulse_responses` the solver forms at step 1 but does not persist) through a local copy of the private `compute_impulse_responses`, reads the equilibrium aggregate path off `dr.G`, and contracts the two by the anti-causal convolution of Section 4.2. If `solve.m` changes upstream, the local copy must be re-synchronized.

```
function [xhat, info] = het_policy_responses(M_, oo_, options_, shock_name, var_name)
% het_policy_responses  Reconstruct the response of a heterogeneous policy
% function to a one-time unit deviation of one aggregate shock at period 0.
%
% Requires: addpath dynare/matlab; so that the MEX `compute_expected_dx` and
% the heterogeneity package internals are reachable.
%
% Inputs:
%   M_, oo_, options_  Standard Dynare workspace (after a successful
%                      heterogeneity_compute_steady_state + heterogeneity_solve).
%   shock_name         Name of the aggregate shock (string), in M_.exo_names.
%   var_name           Name of the heterogeneous variable (string), in
%                      M_.heterogeneity(1).endo_names.
%
% Outputs:
%   xhat  [N_sp x T] cross-section of the variable's response on the
%         joint shock-state grid. xhat(:, t) is the response at period t-1
%         to a unit deviation of shock_name at period 0.
%   info  struct with fields:
%           .grids        oo_.heterogeneity(1).steady_state.pol.grids
%           .shocks_grid  oo_.heterogeneity(1).steady_state.shocks.grids
%           .x_bar        steady-state policy values for var_name
%           .var_idx      index of var_name in H_.endo_names
%           .shock_name
%           .T

oo_het = oo_.heterogeneity(1);
H_     = M_.heterogeneity(1);
T      = double(options_.heterogeneity.solve.truncation_horizon);

if ~ismember(shock_name, M_.exo_names)
    error('het_policy_responses:bad_shock', ...
          'Shock "%s" not found in M_.exo_names.', shock_name);
end
var_idx = find(strcmp(H_.endo_names, var_name), 1);
if isempty(var_idx)
    error('het_policy_responses:bad_var', ...
          'Variable "%s" not found in M_.heterogeneity(1).endo_names.', var_name);
end

sizes   = oo_het.sizes;
indices = oo_het.indices;
N_sp    = sizes.N_sp;
N_Y_all = sizes.N_Y.all;

impulse_responses = local_compute_impulse_responses(int32(T), H_, oo_het);
% [H_.endo_nbr x N_Y x N_sp x T]

Yhat_pad = build_Yhat_with_timing(oo_het.dr, indices, sizes.N_Y, shock_name, T);
% [N_Y_all x 2T]

IR_var = squeeze(impulse_responses(var_idx, :, :, :));   % [N_Y x N_sp x T]
IR_mat = reshape(permute(IR_var, [2 1 3]), N_sp, N_Y_all * T);

xhat = zeros(N_sp, T);
for t = 1:T
    win = Yhat_pad(:, t:t+T-1);          % [N_Y x T]
    xhat(:, t) = IR_mat * win(:);
end

info = struct();
info.grids       = oo_het.steady_state.pol.grids;
info.shocks_grid = oo_het.steady_state.shocks.grids;
info.x_bar       = oo_het.steady_state.pol.values.(var_name);
info.var_idx     = var_idx;
```

```
info.shock_name  = shock_name;
info.T           = T;

end


function Yhat_pad = build_Yhat_with_timing(dr, indices, N_Y, shock_name, T)
% Build the (N_Y_all x 2T) padded matrix whose iY-th row is the equilibrium
% path of the iY-th aggregate input slot, accounting for its timing (lag,
% current, lead, exo).
%
% Padding extends T periods of zeros on the right so the convolution
% xhat(om, t) = sum_iY sum_lag IR(iY, om, lag) * Yhat_pad(iY, t+lag-1)
% can safely index up to t+lag-1 = 2T-1.

N_Y_all = N_Y.all;
Yhat_pad = zeros(N_Y_all, 2*T);
in_Y = 0;

% Convention. path(k) is the equilibrium response at math-period (k-2) of
% the underlying variable to a unit shock at math-period 0, so:
%   path(1) = period -1 = 0
%   path(2) = period  0 (impact)
%   ...
%   path(T+1) = period T-1
%   path(T+2..2T+1) = 0     (past the truncation horizon)
% Then "value at convolution position t of input iY (with timing tau)" is
%   path(t + 1 + tau): tau = -1 (lag) -> path(t)
%                      tau =  0 (cur) -> path(t+1)
%                      tau = +1 (lead)-> path(t+2)

for i = 1:N_Y.lag
    in_Y = in_Y + 1;
    path = irf_path(dr, indices.Y.names.lag{i}, shock_name, T);
    Yhat_pad(in_Y, :) = path(1:2*T);
end
for i = 1:N_Y.current
    in_Y = in_Y + 1;
    path = irf_path(dr, indices.Y.names.current{i}, shock_name, T);
    Yhat_pad(in_Y, :) = path(2:2*T+1);
end
for i = 1:N_Y.lead
    in_Y = in_Y + 1;
    path = irf_path(dr, indices.Y.names.lead{i}, shock_name, T);
    Yhat_pad(in_Y, 1:2*T-1) = path(3:2*T+1);
end
for i = 1:N_Y.exo
    in_Y = in_Y + 1;
    name = indices.Y.names.exo{i};
    path = irf_path_exo(name, shock_name, T);
    Yhat_pad(in_Y, :) = path(2:2*T+1);
end

end


function path = irf_path(dr, name, shock_name, T)
% Padded equilibrium response of an endogenous aggregate `name` to a unit
% deviation of `shock_name` at math-period 0. path(k) is the period-(k-2)
% value; path(1) and path(T+2..2T+1) are zero.
g = dr.G.(name).(shock_name);
path = zeros(1, 2*T+1);
path(2:T+1) = g(:, 1).';
end


function path = irf_path_exo(name, shock_name, T)
% Padded path for an exogenous input. For a one-time unit shock at math
% period 0 to `shock_name`, the exogenous input's value is 1 at period 0
% if name == shock_name, else 0; in either case 0 at all other periods
% (varexo without an explicit process equation is treated as an iid
% innovation with the deterministic single-period perturbation).
path = zeros(1, 2*T+1);
```

```
if strcmp(name, shock_name)
    path(2) = 1;
end
end


function impulse_responses = local_compute_impulse_responses(T, H_, oo_het)
% Local copy of heterogeneity.solve > compute_impulse_responses, which is a
% private function inside solve.m and therefore not callable from outside.
% Kept verbatim except for the function name. If solve.m changes upstream,
% sync this body.
mat     = oo_het.mat;
sizes   = oo_het.sizes;
indices = oo_het.indices;

N_sp = sizes.N_sp;
N_Y  = sizes.N_Y.all;

impulse_responses = zeros(H_.endo_nbr*N_Y, N_sp, T);

ind_x  = H_.endo_nbr + (1:H_.endo_nbr);
ind_xe = H_.endo_nbr + ind_x;

Ex  = compute_expected_dx(mat.pol.x_bar, mat.pol.ind, mat.pol.w, mat.pol.inv_h, mat.Mu, mat.
    pol.dims);
Ae3 = pagemtimes(mat.dF(:, ind_xe, :), Ex);
f3  = mat.dF(:, ind_x, :);
col = H_.state_var;
f3(:, col, :) = f3(:, col, :) + Ae3;

RHS3 = cat(2, mat.dF(:, ind_xe, :), mat.dF(:, indices.Y.ind_in_dF, :));
Z3   = -pagemldivide(f3, RHS3);
cFx_next = Z3(:, 1:H_.endo_nbr, :);
x        = Z3(:, H_.endo_nbr+1:end, :);

x_perm = reshape(x, [], N_sp);
impulse_responses(:, mat.pol.P, 1) = (x_perm(:, mat.pol.Q) / mat.pol.U) / mat.pol.L;

for s = 2:T
    Ex_dash = impulse_responses(:,:,s-1) * mat.pol.Phi_e;
    x       = pagemtimes(cFx_next, reshape(Ex_dash, H_.endo_nbr, N_Y, N_sp));
    x_perm  = reshape(x, [], N_sp);
    impulse_responses(:, mat.pol.P, s) = (x_perm(:, mat.pol.Q) / mat.pol.U) / mat.pol.L;
end

impulse_responses = reshape(impulse_responses, H_.endo_nbr, N_Y, N_sp, T);

end
```

## H.2 The distribution response

The function `het_distribution_responses` returns the cross-section $\hat{\Omega}_t$ of the joint distribution. It reads the stored displacement arrays `dr.curlyDs` $= \mathsf{M}^h p \times_s$, rebuilds the steady-state push-forward $\Lambda^d$ through `build_steady_state_pushforward`, and iterates the forward law of motion of Section 4.5 from the predetermined start $\hat{\Omega}_0 = 0$.

```
function [omhat, info] = het_distribution_responses(M_, oo_, options_, shock_name)
% het_distribution_responses  Reconstruct the response of the joint
% distribution to a one-time unit deviation of one aggregate shock at
% period 0.
%
% Implements the BBEG-Appendix-C recursion
%   omhat_0 = 0
%   omhat_{t+1} = Lambda_ss * omhat_t + M^h * abar_t
% where Lambda_ss is the steady-state push-forward (policy lottery
% composed with the shock chain), abar_t is the equilibrium policy shift
% at math-period t, and M^h is the divergence-form operator that maps
% policy shifts to distribution drifts.
%
% By linearity, the period-t drift M^h * abar_t equals
%   sum_iY sum_{tau=0..T-1} curlyDs.(iY)(tau+1, :)' * Yhat^iY(t+tau)
```

```
% where curlyDs.(iY)(s+1, :)' is M^h applied to BBEG's unit-shock policy
% response at anticipation lag s (read directly from oo_.dr.curlyDs).
%
% Requires: addpath dynare/matlab;
%
% Inputs:
%   M_, oo_, options_  Standard Dynare workspace (after a successful
%                      heterogeneity_compute_steady_state + heterogeneity_solve).
%   shock_name         Name of the aggregate shock (string), in M_.exo_names.
%
% Outputs:
%   omhat  [N_om x T] flat distribution responses. omhat(:, k) is the
%          response at math-period (k-1). The (e, a)-tensor view is in
%          info.omhat_tensor.
%   info   struct with fields:
%            .Lambda        sparse N_om x N_om steady-state push-forward
%            .omhat_tensor  reshape into [n_shock_1 x ... x n_state_1 x ... x T]
%            .grids         oo_.heterogeneity(1).steady_state.d.grids
%            .shocks_grid   oo_.heterogeneity(1).steady_state.shocks.grids
%            .shock_name
%            .T
%            .inputs_used   cell array of curlyDs input names actually used

oo_het = oo_.heterogeneity(1);
T      = double(options_.heterogeneity.solve.truncation_horizon);

if ~ismember(shock_name, M_.exo_names)
    error('het_distribution_responses:bad_shock', ...
          'Shock "%s" not found in M_.exo_names.', shock_name);
end

sizes   = oo_het.sizes;
indices = oo_het.indices;
N_e     = sizes.N_e;
N_a_d   = sizes.d.N_a;
N_om    = N_e * N_a_d;
N_Y     = sizes.N_Y;

Lambda   = build_steady_state_pushforward(oo_het);
Yhat_pad = build_Yhat_with_timing_d(oo_het.dr, indices, N_Y, shock_name, T);

% Pre-collect, for each aggregate-input slot iY, the curlyDs slice of shape
% [N_om x T] in om-major layout. We store transposed so the per-iteration
% contraction is one matrix-vector multiply.
[curlyDs_T, input_names] = collect_curlyDs(oo_het.dr.curlyDs, indices, N_Y, T, N_om);
inputs_used = input_names(~cellfun(@isempty, input_names));

omhat = zeros(N_om, T);
for k = 1:T-1
    drift = zeros(N_om, 1);
    for iY = 1:N_Y.all
        if isempty(curlyDs_T{iY})
            continue
        end
        win = Yhat_pad(iY, k:k+T-1).';            % [T x 1]
        drift = drift + curlyDs_T{iY} * win;       % [N_om x T] * [T x 1]
    end
    omhat(:, k+1) = Lambda * omhat(:, k) + drift;
end

% Tensor view: [n_shock_1 x ... x n_state_1 x ... x T]
shock_names = indices.shocks;
state_names = indices.states;
tensor_dims = zeros(1, numel(shock_names) + numel(state_names) + 1);
for i = 1:numel(shock_names)
    tensor_dims(i) = sizes.shocks.(shock_names{i});
end
for i = 1:numel(state_names)
    tensor_dims(numel(shock_names) + i) = sizes.d.states.(state_names{i});
end
tensor_dims(end) = T;
omhat_tensor = reshape(omhat, tensor_dims);
```

```
info = struct();
info.Lambda       = Lambda;
info.omhat_tensor = omhat_tensor;
info.grids        = oo_het.steady_state.d.grids;
info.shocks_grid  = oo_het.steady_state.shocks.grids;
info.shock_name   = shock_name;
info.T            = T;
info.inputs_used  = inputs_used;

end


function [curlyDs_T, input_names] = collect_curlyDs(curlyDs_struct, indices, N_Y, T, N_om)
% Returns:
%   curlyDs_T   {1 x N_Y.all} cell, curlyDs_T{iY} is [N_om x T] (transposed
%               from the [T x N_om] storage so we can right-multiply by a
%               [T x 1] window). Empty if the field is not present.
%   input_names {1 x N_Y.all} cell, the corresponding dr.curlyDs field
%               name (empty if missing).

curlyDs_T = cell(1, N_Y.all);
input_names = cell(1, N_Y.all);

in_Y = 0;
[in_Y, curlyDs_T, input_names] = pull(in_Y, curlyDs_T, input_names, curlyDs_struct, ...
    indices.Y.names.lag,     '_lag',   T, N_om);
[in_Y, curlyDs_T, input_names] = pull(in_Y, curlyDs_T, input_names, curlyDs_struct, ...
    indices.Y.names.current, '',       T, N_om);
[in_Y, curlyDs_T, input_names] = pull(in_Y, curlyDs_T, input_names, curlyDs_struct, ...
    indices.Y.names.lead,    '_lead',  T, N_om);
[in_Y, curlyDs_T, input_names] = pull(in_Y, curlyDs_T, input_names, curlyDs_struct, ...
    indices.Y.names.exo,     '',       T, N_om);              %#ok<ASGLU>

end


function [in_Y, curlyDs_T, input_names] = pull(in_Y, curlyDs_T, input_names, ...
    curlyDs_struct, names, suffix, T, N_om)
for i = 1:numel(names)
    in_Y = in_Y + 1;
    fld = [names{i} suffix];
    if isfield(curlyDs_struct, fld)
        slice = reshape(curlyDs_struct.(fld), T, N_om);   % [T x N_om]
        curlyDs_T{in_Y}   = slice.';                       % [N_om x T]
        input_names{in_Y} = fld;
    end
end
end


function Yhat_pad = build_Yhat_with_timing_d(dr, indices, N_Y, shock_name, T)
% Same convention as build_Yhat_with_timing inside het_policy_responses:
% Yhat_pad is [N_Y.all x 2T], indexed by aggregate-input slot in solve.m
% order (lag block, current block, lead block, exo block). Column k holds
% the equilibrium value of the iY-th input slot at math-period (k-1),
% with the appropriate timing shift baked in.
N_Y_all = N_Y.all;
Yhat_pad = zeros(N_Y_all, 2*T);
in_Y = 0;

for i = 1:N_Y.lag
    in_Y = in_Y + 1;
    p = irf_path_d(dr, indices.Y.names.lag{i}, shock_name, T);
    Yhat_pad(in_Y, :) = p(1:2*T);
end
for i = 1:N_Y.current
    in_Y = in_Y + 1;
    p = irf_path_d(dr, indices.Y.names.current{i}, shock_name, T);
    Yhat_pad(in_Y, :) = p(2:2*T+1);
end
for i = 1:N_Y.lead
    in_Y = in_Y + 1;
    p = irf_path_d(dr, indices.Y.names.lead{i}, shock_name, T);
```

```
    Yhat_pad(in_Y, 1:2*T-1) = p(3:2*T+1);
end
for i = 1:N_Y.exo
    in_Y = in_Y + 1;
    name = indices.Y.names.exo{i};
    p = irf_path_d_exo(name, shock_name, T);
    Yhat_pad(in_Y, :) = p(2:2*T+1);
end

end


function p = irf_path_d(dr, name, shock_name, T)
g = dr.G.(name).(shock_name);
p = zeros(1, 2*T+1);
p(2:T+1) = g(:, 1).';
end


function p = irf_path_d_exo(name, shock_name, T)
p = zeros(1, 2*T+1);
if strcmp(name, shock_name)
    p(2) = 1;
end
end
```

## H.3 The steady-state push-forward

The function `build_steady_state_pushforward` assembles the sparse $N_\Omega \times N_\Omega$ steady-state transfer operator $\Lambda^d$ of Section 4.5 (the policy lottery on the predetermined-state sub-grid composed with the shock Markov chain), in the flat index $\texttt{om} = (j-1)N_e + i$ used throughout `dynare/matlab/+heterogeneity`.

```
function Lambda = build_steady_state_pushforward(oo_het)
% build_steady_state_pushforward  Sparse N_om x N_om matrix that maps a
% distribution column-vector at time t to the same vector at time t+1
% under the steady-state policy lottery (on the pre-determined state
% sub-grid) composed with the steady-state shock Markov chain.
%
% Convention. The joint flat index is om = (j-1)*N_e + i, where i runs
% over the Kronecker-flattened shock index (1..N_e) and j runs over the
% Kronecker-flattened a-state index (1..N_a). This is the layout used
% throughout `dynare/matlab/+heterogeneity` and the matching MEX files.
%
% Math. For source (e=i, a=j_src), the policy lottery on the a-grid
% produces destination corner k with low-side weight
%   beta(om, k) = prod over dims of (w(om, l) if bit_l of k is 0 else 1-w(om, l)),
% and destination a-state index
%   a_dst(om, k) = 1 + sum_l state_stride(l) * (ind(om, l) - 1 + bit_l(k)).
% Then the shock chain transitions to e' = e_dst with probability
% mu(e, e_dst). Composing:
%   Lambda(om_dst, om_src) = sum_corner beta(om_src, k) * mu(e_src, e_dst)
% where om_dst = (a_dst(om_src, k) - 1)*N_e + e_dst.
%
% Inputs:
%   oo_het  oo_.heterogeneity(1) (must contain mat.d.ind, mat.d.w,
%           mat.Mu, mat.d.dims).
%
% Output:
%   Lambda  sparse N_om x N_om matrix.

ind   = oo_het.mat.d.ind;       % [N_om x n] int32, low-corner state idx
w     = oo_het.mat.d.w;         % [N_om x n] double, low-corner weight
Mu    = oo_het.mat.Mu;          % [N_e x N_e] double, shock chain
dims  = double(oo_het.mat.d.dims);  % [1 x n] int32, N_a per dim

[N_om, n] = size(w);
N_e = size(Mu, 1);
N_a = prod(dims);
if N_om ~= N_e * N_a
    error('build_steady_state_pushforward:size_mismatch', ...
          'N_om = %d does not match N_e * N_a = %d * %d = %d.', ...
```

```
          N_om, N_e, N_a, N_e*N_a);
end

state_stride = [1, cumprod(dims(1:end-1))];   % 1 x n
Kcorn = bitshift(1, n);

om_src_all = (1:N_om).';
e_src_all  = mod(om_src_all - 1, N_e) + 1;

ii_cells = cell(Kcorn, 1);
jj_cells = cell(Kcorn, 1);
vv_cells = cell(Kcorn, 1);

for k = 0:Kcorn-1
    bits = bitand(k, bitshift(1, 0:n-1)) ~= 0;   % 1 x n logical

    W_eff = w;
    for l = 1:n
        if bits(l)
            W_eff(:, l) = 1 - W_eff(:, l);
        end
    end
    beta_k = prod(W_eff, 2);

    a_dst_idx = ones(N_om, 1);
    for l = 1:n
        a_dst_idx = a_dst_idx + state_stride(l) * (double(ind(:, l)) - 1 + bits(l));
    end

    om_src_rep = repmat(om_src_all, N_e, 1);
    e_dst_rep  = repelem((1:N_e).', N_om, 1);
    a_dst_rep  = repmat(a_dst_idx, N_e, 1);
    e_src_rep  = repmat(e_src_all, N_e, 1);
    beta_rep   = repmat(beta_k, N_e, 1);

    om_dst_rep = (a_dst_rep - 1)*N_e + e_dst_rep;
    mu_rep = Mu(sub2ind([N_e, N_e], e_src_rep, e_dst_rep));

    v_rep = beta_rep .* mu_rep;
    nz = v_rep ~= 0;

    ii_cells{k+1} = om_dst_rep(nz);
    jj_cells{k+1} = om_src_rep(nz);
    vv_cells{k+1} = v_rep(nz);
end

ii = cat(1, ii_cells{:});
jj = cat(1, jj_cells{:});
vv = cat(1, vv_cells{:});
Lambda = sparse(ii, jj, vv, N_om, N_om);

end
```

# I An executable determinacy check, and what it returns

Section 5.1 argues that the winding number of $\det j_X$ is the sound determinacy diagnostic, that it reads off arrays the solver already forms, and that its verdict is the continuous model's, independent of the mesh and the horizon (Proposition 10). This appendix makes the check executable and records what it returns on the two shipped models, the scalar Krusell–Smith economy and the two-asset HANK model. One code serves both: determinacy is a property of the stacked operator, not of the model's size.

The cleanest object to read is the stacked state Jacobian $\tilde{J}_X$ itself, the operator `compute_state_jacobian` assembles inside `solve.m` and whose $T \times T$ corner Dynare inverts to return `dr.G`. Reading $\tilde{J}_X$ directly, rather than reassembling its symbol from the aggregate blocks and the heterogeneous Jacobian by hand, is what keeps the check short and model-agnostic: the block $(a, b)$ of $\tilde{J}_X$, at rows and columns $(a-1)T + (1{:}T)$ in the order of `M_.endo_names`, is the sequence-space operator relating equation $a$ to variable $b$, and its symbol is the limit of its diagonals, read off the last row (propagation, $k \geq 0$) and the last column (anticipation, $k < 0$), the corner where the truncation lives sitting far from both.

## I.1 Exposing the operator

A two-line hook exposes the operators the solver already forms. It goes in `+heterogeneity/solve.m`, immediately after the line that inverts the corner,

```
G = -full(state_jacobian) \ full(shock_jacobian);
```

and reads

```
% (experiments) expose the stacked operators for the determinacy check of
% Appendix ax:determinacy-check; gate behind an option in production, since
% state_jacobian is (endo_nbr*T) square and large at long horizons.
dr.state_jacobian = state_jacobian;
dr.shock_jacobian = shock_jacobian;
```

The check then runs in the same session, right after `heterogeneity_solve`, on the operators the hook exposed. There is no round-trip: the whole diagnostic stays in MATLAB, alongside the sequence-space steady state the toolkit already loaded.

```
dr   = oo_.heterogeneity(1).dr;
endo = cellstr(M_.endo_names);   exo = cellstr(M_.exo_names);
T    = size(dr.state_jacobian, 1) / numel(endo);
determinacy_check(dr.state_jacobian, dr.shock_jacobian, endo, exo, T, {'K', 'Z'});
```

## I.2 The check

The diagnostic reads the symbol blocks $j_k$ off $\tilde{J}_X$, assembles the matrix symbol $j_X(z) = \sum_k j_k z^k$, and returns four numbers. The winding of $\det j_X$ around the unit circle is the verdict of Proposition 7, $w = 0$ for determinacy. The minimum $\min_{|z|=1} |\det j_X|$ is the determinant floor the finite-sample certificate of Proposition 10 compares the truncation error against. The winding of the reflected symbol $z \mapsto j_X(1/z)$ is the finite-section condition of the truncation license (it must be $-w = 0$ with the operator invertible). And the dimension of the kernel of the pure block-Toeplitz operator $\mathsf{T}(j_X)$, read from a moderate corner, is zero exactly when the factorization is canonical, every partial index $\kappa_i = 0$ (Proposition 9, item 2), the condition the two-pass inverse and the cheap higher-order cascade need beyond $w = 0$. The check closes with a self-validation: it reconstructs $-j_X(z)^{-1} j_{\mathcal{E}}(z)$ from the symbols and compares it, on the circle, against the symbol of the exact response $G = -\tilde{J}_X^{-1}\tilde{J}_{\mathcal{E}}$ the solver returns, so that agreement is the empirical form of the central theorem, the state-space and sequence-space readings meeting in one Fréchet derivative.

```
function out = determinacy_check(SJ, SE, endo, exo, T, validate)
% Determinacy diagnostic: reads the stacked state_jacobian symbol and applies the
% winding test of Section 5.1. One code for any model (KS, two-asset HANK, ...).
% Inputs: SJ, SE the stacked state/shock Jacobians (from the solve.m hook), the
% variable-name cells endo, exo, the horizon T, and validate = {var, exo} for the
% self-check. Reports w, min|det|, the reflected condition, and canonicity.
    nV = numel(endo); nX = numel(exo);
    jX = symbol_blocks(SJ, nV, nV, T);
    jE = symbol_blocks(SE, nV, nX, T);
    JX = @(z) evalsym(jX, z, T);
    JE = @(z) evalsym(jE, z, T);
    [w, mn]  = winding(JX, 8192);
    [wr, mr] = winding(@(z) JX(1/z), 8192);
    % canonicity: dim ker of the pure block-Toeplitz Toep(j_X)_M (0 => all kappa_i = 0)
    M = 80; Tp = zeros(nV*M);
    for i = 1:M
        for j = 1:M
            idx = (i-j) + T;
            if idx >= 1 && idx <= 2*T-1
                Tp((i-1)*nV+(1:nV), (j-1)*nV+(1:nV)) = jX(:,:,idx);
            end
        end
    end
    sv = svd(Tp); kerdim = sum(sv < 1e-7*max(sv));
    % validation: -j_X^{-1} j_E vs the exact -SJ\SE on one (var<-exo) response
    relerr = NaN;
    if nargin > 5 && ~isempty(validate)
        ai = find(strcmp(endo, validate{1})); xi = find(strcmp(exo, validate{2}));
        Gc = -SJ \ SE(:, (xi-1)*T+(1:T));            % sparse solve, T columns
        Gblk = full(Gc((ai-1)*T+(1:T), :));
```

```
        t0 = floor(3*T/4); kk = (-floor(T/4)):(floor(T/4));
        kk = kk(t0-kk >= 1 & t0-kk <= T);
        zz = exp(1i*2*pi*(0:255)/256); rel = zeros(size(zz));
        for m = 1:numel(zz)
            z = zz(m);
            gp = -(JX(z) \ JE(z)); gp = gp(ai, xi);
            gt = 0; for k = kk, gt = gt + Gblk(t0, t0-k)*z^k; end
            rel(m) = abs(gp - gt) / (abs(gt) + 1e-9);
        end
        relerr = median(rel);
        fprintf('  validation median rel-err = %.2e  (%s<-%s)\n', relerr, validate{1},
            validate{2});
    end
    fprintf('  winding(det j_X) = %d   min|det j_X| = %.4g   determinate = %d\n', w, mn, w==0)
        ;
    fprintf('  reflected winding = %d, invertible = %d\n', wr, mr > 1e-6);
    fprintf('  canonical (dim ker Toep = 0) = %d\n', kerdim == 0);
    out = struct('winding', w, 'mindet', mn, 'reflected', wr, ...
                 'canonical', kerdim==0, 'validation', relerr);
end

function jk = symbol_blocks(M, nrow, ncol, T)
% converged symbol blocks j_k: diagonal limit from the last row (k>=0) and the last
% column (k<0); the truncation correction lives in the top-left corner, far from both.
    jk = zeros(nrow, ncol, 2*T-1);
    for a = 1:nrow
        for b = 1:ncol
            lr = full(M((a-1)*T+T, (b-1)*T+(1:T)));
            lc = full(M((a-1)*T+(1:T), (b-1)*T+T));
            for k = 0:T-1,      jk(a,b,k+T) = lr(T-k);  end
            for k = -(T-1):-1,  jk(a,b,k+T) = lc(T+k);  end
        end
    end
end

function J = evalsym(jk, z, T)
    zp = reshape(z.^(-(T-1):(T-1)), 1, 1, []);
    J = sum(jk .* zp, 3);
end

function [w, mn] = winding(f, N)
    zc = exp(1i*2*pi*(0:N-1)/N); d = zeros(1, N);
    for m = 1:N, d(m) = det(f(zc(m))); end
    ph = unwrap(angle([d, d(1)]));
    w = round((ph(end) - ph(1)) / (2*pi)); mn = min(abs(d));
end
```

## I.3 What it returns on the two shipped models

The check is run on the Krusell–Smith model of Section 1 and on the shipped two-asset HANK model, the latter re-solved at the longer horizon $T = 800$ its slow mixing needs (Table 5). Both are determinate, $w = 0$, and both factorize canonically, so the two-pass inverse and the cheap cascade apply to each; the two-asset model, with $n > 1$ aggregate unknowns, is the only one of the two that *could* have shown the benign non-canonical index-zero case of Proposition 9, and it does not. The self-validation lands at $8 \times 10^{-5}$ for the scalar model and $2 \times 10^{-2}$ for the two-asset model at $T = 800$: the symbol read off $\tilde{J}_X$ reproduces the solver's own `dr.G`, the two readings of the same operator.

| Model | $n$ | $T$ | validation | $w$ | $\min \lvert \det j_X \rvert$ | canonical |
|---|---|---|---|---|---|---|
| Krusell–Smith | 5 | 400 | $8 \times 10^{-5}$ | 0 | 1.00 | yes |
| Two-asset HANK | 22 | 800 | $2 \times 10^{-2}$ | 0 | 0.0022 | yes |

Table 5: The determinacy check on the two shipped models. Both determinate ($w = 0$), both canonical (all partial indices zero), the reflected condition holding in each. The validation column is the median relative error between the symbol reconstruction of $-\tilde{J}_X^{-1}\tilde{J}_\mathcal{E}$ and the solver's `dr.G`.

Three readings of the table are worth recording. First, the verdict is horizon-independent, as Proposition 10 promises: for the scalar model the winding stays 0 when the horizon is cut from 400 to 50, the

determinant floor staying near 1 throughout, so the diagnostic returns the continuous model's verdict rather than an artifact of the truncation. Second, the determinant floor is the quantity that sizes the finite-sample certificate, and it is model-specific: 1.0 for the scalar model but 0.0022 for the two-asset model, roughly five hundred times smaller and stable across $T = 300$ and $T = 800$, so it is the genuine floor and not a truncation artifact. Because the certificate asks the truncation error to fall below that floor, and because the error carries the constant of the geometric decay rather than the bare tail (for the scalar model the mixing rate $\lambda \approx 0.97$, the symbol constant $C \approx 2.5 \times 10^3$, and the determinant's Lipschitz constant $L \approx 6 \times 10^2$ put the sharp certificate's crossing at $T \gtrsim 434$, past the default horizon even though the verdict is already robust), the horizon a sharp certificate needs grows both as the model mixes more slowly and as its floor shrinks: the two-asset model, slow-mixing and low-floored at once, is the demanding case, and its default truncation horizon should be set accordingly. Third, the canonicity read is the payoff of working with $n > 1$: the scalar model cannot separate a canonical factorization from a determinate one, $w = 0$ forcing $\kappa = 0$ when there is a single partial index, whereas the two-asset model can, and the empty kernel of $\mathsf{T}(j_X)$ is what certifies that its factorization is canonical rather than the benign non-canonical alternative.