Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

When Finite Learning Looks Reliable: Equilibrium Consistency and Iteration Dynamics in Learning from Prices

Mateo Bodon

mateo.bodon@yale.edu

Yale University

Research Article

Keywords: learning from prices, fixed-point iteration, computational reproducibility, numerical stability, market clearing

Posted Date: September 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10743822/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

When Finite Learning Looks Reliable: Equilibrium Consistency and Iteration Dynamics in Learning from Prices

Mateo Bodon
Yale University, New Haven, Connecticut, USA
Corresponding author: mateo.bodon@yale.edu
ORCID: 0009-0004-5012-835X

August 2026

Abstract

Finite learning paths can look numerically reliable even when the iteration that generates them does not reliably approximate equilibrium. I study this distinction in a learning-from-prices model by separating three objects: the equilibrium conditions in the article, the fixed-point residual in deposited code, and the finite path generated by the updating rule. At the benchmark calibration, exact algebra shows that one denominator mismatch in the deposited private fixed-point routine produces a distinct nonzero root. The article/updater root satisfies independently reconstructed market clearing; the deposited-routine root leaves a nonzero mean-supply residual. Their price functions differ by 3.53% at the fixed point and 16.64% after the first prescribed update. Yet the plotting driver never uses the discrepant fixed-point output, so I do not infer a change in any displayed figure or in the article's reported economic mechanism. The broader result concerns finite iteration. The partially revealing map has a Jacobian equal to a scalar times a rank-two oblique projector: the factorization yields a sharp local-stability condition, although the first step can still amplify error. Exact enumeration classifies 39 of 100 cells in the deposited sweep and 1,833 of 8,100 prespecified parameter cells as unstable. A scalar reduction identifies the singular boundary and every finite preimage of the pole, separating convergent, two-cycle, and explosive regimes. Together, the results show why a behavioral horizon, equilibrium validity, and a numerical approximation horizon must be evaluated separately.

Keywords: learning from prices; fixed-point iteration; computational reproducibility; numerical stability; market clearing.
JEL classification: C63, D82, D83, G14.

1 Introduction

Finite rounds of learning can look numerically reliable even when the iteration that generates them does not reliably approximate equilibrium. The reason is that three closely related objects can diverge: the equilibrium conditions stated in an article, the residual equations solved by deposited code, and the finite sequence produced by an updating rule. Agreement among them is often assumed rather than established. In a learning model, that assumption is especially consequential because a finite number of reasoning rounds may be the economic object of interest even when the same recurrence is unstable or close to a pole.

I examine this joint problem using the replication package for Bastianello and Fontanier [2025]. The article develops a theory of partial-equilibrium thinking (PET) and learning from prices; its stated equilibrium conditions, behavioral mechanism, and broad convergence results are contributions of the original authors. In particular, Sections D.1–D.2 of the online appendix derive the K-level recurrence and its convergence properties, Section F treats partially revealing prices, and Section G states the private-signal equations [Bastianello and Fontanier, 2024a]. The deposited private fixed-point routine, however, differs in one denominator from the article's online appendix and the package's own private updating routine. I analyze both specifications without altering either. I call the deposited residual the *deposited-D system* and the article/companion-updater residual the *article/updater-H system*, using the denominator labels defined below. These labels describe the observed source forms without assigning a version history and prevent a solver's root from being treated as an equilibrium before market clearing is checked. The deposited driver assigns the discrepant fixed-point output but never reads it again; the plotted private path instead comes from the article/updater-H routine. The mismatch is therefore confined to a deposited fixed-point routine, not evidence that the article's reported figures or economic mechanism are wrong.

The first result is exact. At the benchmark calibration, the two residuals have different unique denominator-safe roots. Three price coefficients coincide; the mean-supply loading differs. Independent Gaussian projection and market-clearing conditions give zero residuals at the article/updater-H root, coefficient by coefficient, but a mean-supply residual of $5000/116781$ at the deposited-D root. The difference is not confined to the limiting roots. From the prescribed common start, the normalized price-function gap is 16.639312% after one update, compared with 3.526856% at the fixed point. Thus an early iterate can separate the two systems more sharply than the eventual comparison does.

The second result concerns the partially revealing source map. At its exact nonzero root the four-dimensional Jacobian factors as $J = -\varepsilon\Pi$, with $\Pi^2 = \Pi$ and rank two. Thus two directions are eliminated in one update and two share the eigenvalue $-\varepsilon$. The projector identity rules out Jordan-chain amplification after the first linearized update. Because Π is generally oblique, however, its operator norm can exceed one, so the spectrum alone does not bound first-step amplification. Exact enumeration finds 39 unstable cells in the deposited package's 100-point sweep and 1,833 unstable cells in a prespecified 8,100-cell variance grid. By contrast, the displayed $h = 0.60$ benchmark is stable and extremely well converged. Local stability therefore governs what persists after the first projection; it does not by itself bound the size of that first step.

Third, I derive a parameter-uniform scalar reduction for the core shared by the two private maps. It identifies the pole and all of its finite preimages and distinguishes three regimes. When $t > 1$, every common forward-defined orbit converges in the core. At $t = 1$, the reduced coordinate alternates unless it begins at the fixed point; for $t < 1$, it expands unless it begins there. For $t \leq 1$, convergence of the remaining coefficients requires additional restrictions, and the article/updater-H convergence set is strictly contained in the deposited-D convergence set. The rational recurrence and its convergence boundary already appear in the authors' appendix; my contribution is applying them to the two implemented coefficient fibers and deriving the explicit pole and preimage diagnostics. Related explosive PET dynamics are studied in Bastianello and Fontanier [2026].

Finally, I evaluate finite-horizon behavior under a design fixed before the outcomes were computed. Undefined iterations and pole cases remain visible rather than being filtered out. At the displayed benchmark, a large first-step discrepancy coexists with rapid eventual convergence. Across the grid, some unstable paths pass a finite tolerance before their limiting classification is visible; all 1,833 unstable paths approach a singular boundary limit under the prespecified-start closed orbit. Because the limit lies on the source map's

special-branch boundary, it is neither an equilibrium nor an attractor. Early closeness is therefore evidence about a particular horizon, not proof of convergence.

The most directly relevant prior results are in Appendices D and G of the original article. Finite rounds of inference from prices also appear in Zhou [2022] and Carvajal and Zhou [2025a]; adaptive learning from prices and convergence to rational-expectations equilibrium are studied directly by Bray [1982], Blume et al. [1982], and Carvajal and Zhou [2025b]. Broad learning stability has a longer history including Grandmont [1998]. The contribution is narrower: an exact, source-preserving comparison of two implemented fixed-point systems, their market-clearing implications, and prespecified finite-iteration consequences. This focus complements work on comparing numerical solution methods [Aruoba et al., 2006] and the growing evidence on computational reproducibility in finance and economics [Christensen and Miguel, 2018, Pérignon et al., 2024]. In *Computational Economics*, Kirkby [2023] separates reproducibility from independent replication across fourteen quantitative-macroeconomic papers. That precedent motivates the independent-code check here, but the conclusion remains specific to the available evidence: the mismatched fixed-point output is not read by the plotting driver. Online Resource 1 provides the complete claim ledger, source boundary, proofs, and prespecified protocol. Online Resource 2 contains the reviewer-safe reproduction archive. The Statements and Declarations section reports the use of generative AI tools in preparing and checking this work.

The paper proceeds from the bounded source comparison to the exact market-clearing result, then to local stability, global pole geometry, and finite-horizon evidence. The final sections synthesize the economic interpretation and state the contribution and claim boundaries.

2 What the deposited routines solve

2.1 Source forms and evidentiary scope

The analysis begins with the official article, online appendix, and Zenodo v1 package [Bastianello and Fontanier, 2024b]. The source-faithful evidence comes from one recorded, unmodified MATLAB R2023b Update 11 execution; the current reproduction does not rerun it. Independent code agrees with the finite paths stored in MATLAB figure objects 10 and 11 to 4.44×10^{-16} and with the deposited-D coefficients to 4.94×10^{-13}. I do not claim an independent reconstruction of figure objects 1–9. Online Resource 1 records the archive hashes, execution diagnostics, tolerance, and exact source-to-output custody.

The deposited driver assigns the output of `fixedpoint_ree` once, but no later line reads it; the private-price figure instead uses rows from `klevelupdating_private`, whose denominator is article/updater-H. The mismatch therefore does not feed a displayed package result. Online Resource 1 provides the complete source-object table and static call-graph boundary.

2.2 Private-map notation

Let $a, b, c, A > 0$ denote the prior variance, signal-noise variance, supply-noise variance, and risk aversion in the private-signal routine. Write $x = (p_0, v, q, z)$ for the constant, value-signal, mean-supply, and realized-supply price coefficients. Define

$$D(x) = abv^2 + c(a+b)z^2, \qquad H(x) = D(x) - abv. \tag{1}$$

The article/updater-H residual is the fixed-point residual of

$$p_0' = \frac{bcz^2 - abvp_0}{H}, \qquad v' = \frac{acz^2}{H},$$
$$q' = -\frac{abv(q+z)}{H}, \qquad z' = -\frac{Aabcz^2}{H}. \tag{2}$$

The deposited-D map differs only by using D rather than H in the denominator of the q' equation. The displayed algebra cancels a factor of v present in the source, so the source-faithful common domain is $vDH \neq 0$; I do not extend either map through $v = 0$ by cancellation. Undefined values are never filled by continuity without labeling the boundary.

3 One denominator, two roots, and a market-clearing test

Set

$$t = A^2bc, \qquad \kappa = a + t(a+b). \tag{3}$$

The denominator difference enters only the mean-supply loading at the fixed point. That narrow algebraic difference nevertheless shifts the entire conditional price schedule when mean supply is held fixed.

Proposition 1 (Exact private roots)**.** *For positive primitives, every denominator-safe nonzero fixed point of the common core of* (2) *satisfies*

$$p_0^\star = \frac{tb}{\kappa}, \qquad v^\star = \frac{a(1+t)}{\kappa}, \qquad z^\star = -Abv^\star. \tag{4}$$

The article/updater-H and deposited-D mean-supply loadings are uniquely

$$q_H^\star = \frac{Aab}{\kappa}, \qquad q_D^\star = \frac{Aab(1+t)}{\kappa(2+t)}. \tag{5}$$

At the benchmark calibration $(a, b, c, A) = (4/25, 8/25, 19/20, 2)$, Proposition 1 gives

$$x_H^\star = (304/581, 277/581, 80/581, -4432/14525),$$
$$x_D^\star = (304/581, 277/581, 11080/116781, -4432/14525). \tag{6}$$

The two routines therefore agree on the three core coefficients and disagree on one loading. The remaining question is economic rather than numerical: which root satisfies the independently reconstructed clearing conditions? A nonlinear solver returning a root cannot answer that question because a root is defined relative to the equations supplied to the solver.

For a candidate (p_0, v, q, z), define the Gaussian-projection coefficients and posterior variance

$$\alpha = \frac{acz^2}{D}, \qquad \ell = \frac{abv}{D}, \qquad \sigma_P^2 = \frac{abcz^2}{D}. \tag{7}$$

The two posterior normal equations and the coefficient-by-coefficient clearing residual used below are, respectively,

$$N(x) = \begin{pmatrix} \alpha(a+b) + \ell a v - a \\ \alpha a v + \ell(a v^2 + c z^2) - a v \end{pmatrix}, \tag{8}$$

$$R(x) = \begin{pmatrix} 1 - \alpha - \ell v - p_0 \\ \alpha + \ell v - v \\ -\ell z - q \\ \ell z - z - A\sigma_P^2 \end{pmatrix}. \tag{9}$$

These equations reconstruct the private Gaussian projection and aggregate constant-absolute-risk-aversion (CARA) normal market clearing independently of either fixed-point residual.

Proposition 2 (Independent market-clearing check)**.** *At the benchmark calibration, the article/updater-H root satisfies* $N(x_H^\star) = 0$ *and* $R(x_H^\star) = 0$ *exactly. The deposited-D root has residual vector*

$$R_D = (0, 0, 5000/116781, 0) \tag{10}$$

in coefficient order (p_0, v, q, z) *and unit-mean-supply excess demand* $15625/61104$.

The proposition supports precise terminology. The deposited-D object is the unique denominator-safe fixed point of the deposited residual, but it is not an equilibrium of the independently reconstructed market-clearing system. After mean supply is set to one, the excess demand is expressed in the reconstruction's demand units. It is an exact coefficient-level failure of market clearing, not a percentage discrepancy or welfare loss.

Proposition 3 (Fixed and first-update gaps)**.** *On the original prespecified source normalization—the realization grid* $\xi = .20, .21, \ldots, 1.80$*, with prior mean, mean supply, and realized supply each fixed at one—the deposited-D versus article/updater-H relative sup-norm gap is* $125000/3544233 = 3.526856\%$ *at the fixed point and* $125000/751233 = 16.639312\%$ *after one update from the prescribed common start. The denominators are horizon-specific article/updater-H price sup norms on that grid:* $2519/2075$ *at the fixed point and* $3879/2575$ *at the first update. The ratio of the first-update normalized gap to the fixed-point normalized gap is* $168773/35773$*. The ratio of the corresponding unnormalized constant gaps is instead* $603/103$.

Both normalized gaps exceed the prespecified one-percent numerical-discrepancy threshold. This is an outcome-independent reporting rule, not a welfare or return benchmark. Economic importance depends on units and on the joint distribution of realizations and supply, which this exercise does not estimate. Table 1 separates the exact identities from their decimal reporting forms and from the interpretation each number can support.

Table 1: Exact private-map identities and prespecified numerical discrepancy.

Object	Exact value	Decimal
Article/updater-H root	$(304/581, 277/581, 80/581, -4432/14525)$	–
Deposited-D mean-supply loading	11080/116781	0.094878
Mean-supply residual	5000/116781	0.042815
Fixed gap (fixed-H norm)	125000/3544233	3.526856%
First-update gap (time-1 H norm)	125000/751233	16.639312%
Normalized amplification	168773/35773	4.717888
Raw coefficient-gap amplification	603/103	5.854369

The discrepancy is largest after the first prescribed update, not in the limiting comparison. A finite iterate can therefore look like an innocuous step toward convergence while temporarily moving the two price functions farther apart. The separate 101-point-grid results and their distinct normalization are retained as claims C14–C15 in Online Resource 1; they do not replace the original-grid comparison in Proposition 3.

4 Stability after the first learning update

The fixed-point mismatch is one part of the reliability problem. A separate question is whether a finite path generated by a source updater remains a good approximation to its own nonzero fixed point. The partially revealing map makes that distinction exact.

Let $h \in (0, 1]$ denote the partial-revelation weight and retain positive a, b, c, A. Define

$$E = h^2 + A^2bc, \qquad \bar{H} = h + A^2bc, \qquad P = bE + ah\bar{H}. \tag{11}$$

Write $x_h = (p_0, p_s, q, p_Z)$, where p_s is the signal-price coefficient and p_Z is the realized-supply coefficient.

Proposition 4 (Exact nonzero partial root)**.** *On the positive domain, the partially revealing source residual has the unique nonzero root*

$$x_h^\star = \left(\frac{bE}{P}, \frac{ah\bar{H}}{P}, \frac{Aabh(1-h)}{P}, -\frac{Aab\bar{H}}{P} \right). \tag{12}$$

All four original rational residuals vanish exactly.

At $(a, b, c, A, h) = (4/25, 3/25, 1/10, 2, 3/5)$, this is $(85/193, 108/193, 16/193, -216/965)$.

Proposition 5 (Projector stability)**.** *At (12), the exact Jacobian of the source update is*

$$J = -\varepsilon\Pi, \qquad \varepsilon = \frac{h(1-h)}{h^2 + A^2bc}, \tag{13}$$

where $\Pi^2 = \Pi$ *and* $\operatorname{rank}(\Pi) = 2$*. Hence* $J^2 = -\varepsilon J$ *and*

$$\operatorname{spec}(J) = \{0, 0, -\varepsilon, -\varepsilon\}. \tag{14}$$

For $0 < h < 1$*, the minimal polynomial is* $\lambda(\lambda + \varepsilon)$*. At* $h = 1$*,* $J = 0$ *and the minimal polynomial is* λ*. Local asymptotic stability holds exactly when* $|\varepsilon| < 1$*;* $|\varepsilon| = 1$ *is a boundary.*

The factorization separates two stages of local behavior. The first update projects a perturbation onto a two-dimensional image and can amplify it because Π is generally oblique. Every later linearized update scales what remains by $|\varepsilon|$. The eigenvalues therefore characterize persistence after the first step, not the size of that step. Online Resource 1 gives the full projector matrix and exact symbolic certificate. Its boundary experiment isolates persistence within $\mathrm{im}(\Pi)$; it does not estimate general off-image first-step amplification or a grid of operator norms that was not prespecified.

4.1 Exact parameter censuses

The census asks where this persistence factor is smaller than one. For the deposited package's variance sweep, exact enumeration gives instability at $h = 0.06, 0.07, \ldots, 0.44$, or 39 of 100 cells. At the displayed $h = 0.60$, $\varepsilon = 10/17$; the $K = 99$ error expressions contain the factor $(10/17)^{98}$ and satisfy the recorded verification bounds. On the prespecified grid $b/a, c/a \in \{2^{-4}, \ldots, 2^4\}$ and $h \in \{.01, \ldots, 1\}$, 1,833 cells are unstable and 6,267 stable. There are no exact boundary cells. The closest cell has $\varepsilon = 625/626$, an exact margin $1/626$ from one. Figure 1 reports the number of unstable h values within every variance-ratio cell; the printed cell counts make the classification readable without relying on color.

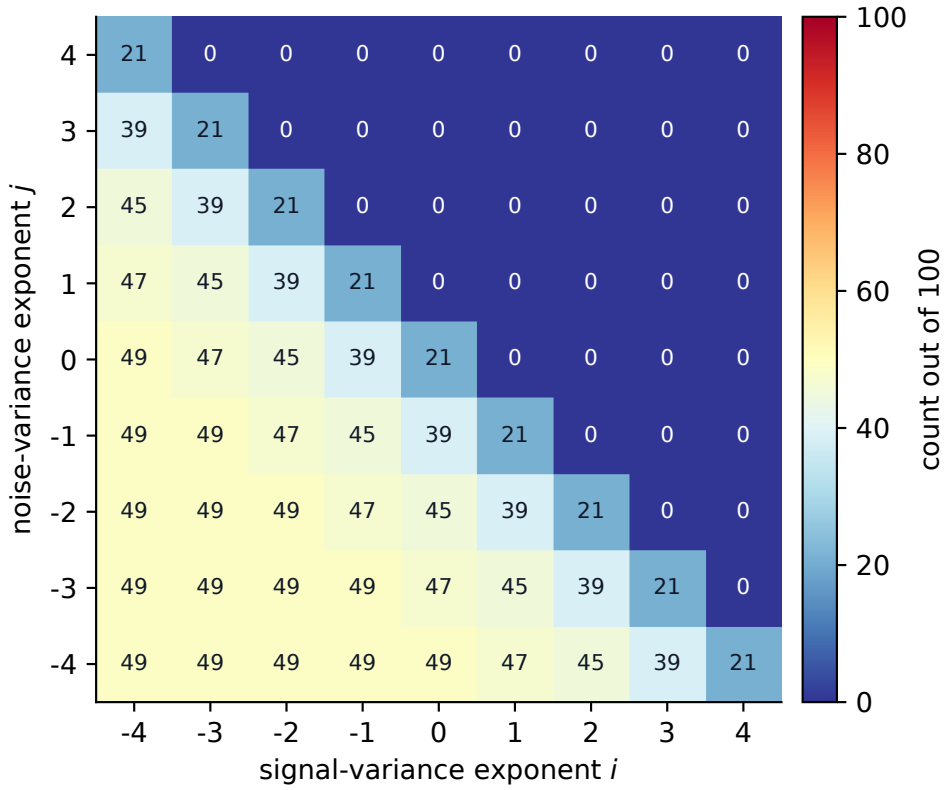


Fig. 1 Number of unstable h values in each variance-ratio cell. Every cell prints its count so the result does not depend on color perception. Exact rational comparisons give 1,833 unstable and 6,267 stable grid cells, with no boundary cells

5 Global geometry: convergence, cycling, and the pole

After one defined update, the shared four-dimensional core collapses to a one-dimensional Möbius recurrence. In that coordinate, convergence, two-cycling, expansion, and finite pole hits are governed by one scalar multiplier. The theorem is stated on the common forward-defined domain because an initially finite point can still be a finite preimage of a later pole.

Theorem 1 (Common core and convergence sets)**.** *Let $a, b, c, A > 0$, define t, κ by (3), and consider the article/updater-H and deposited-D maps on their common source domain $vDH \neq 0$. Restrict attention to finite orbits for which every subsequent iterate stays in this domain. After one defined update, both maps satisfy $z = -Abv$. For every subsequent defined iterate, set*

$$y = \frac{\kappa v}{a}, \qquad w = 1 - \frac{1+t}{y}. \tag{15}$$

Then

$$y' = \frac{ty}{y-1}, \qquad w' = -\frac{w}{t}. \tag{16}$$

The pole is $y = 1$, equivalently $w = -t$. An invariant-line initial coordinate w_1 reaches the pole at reduced index $n \geq 1$ if and only if

$$w_1 = (-t)^n. \tag{17}$$

If $t > 1$, every finite common forward-defined orbit converges under each map to that map's nonzero fixed point, so the two core convergence domains coincide. If $t = 1$, the reduced coordinate alternates unless $w_1 = 0$; if $0 < t < 1$, its magnitude grows unless $w_1 = 0$. For $0 < t \leq 1$, convergence to the nonzero root also requires $p_0 + v = 1$. On the invariant line, the article/updater-H convergence set further requires $q = q^{\star}_H$, whereas the deposited-D q fiber converges from every finite q. Thus the article/updater-H convergence set is a strict subset of the deposited-D convergence set for $0 < t \leq 1$.

The reduction makes the regimes transparent. When $t > 1$, the distance-like coordinate w contracts while alternating in sign. At $t = 1$, its magnitude is constant; below one, its magnitude expands. The qualifications in the theorem matter: convergence is asserted only for finite common forward-defined orbits, and a path that reaches $w = -t$ terminates at the pole.

Equation (16) has the same Möbius form as equation (D.33) in the original appendix: under $y = \widetilde{\gamma}/\beta$ and $t = \alpha/\beta$, that recurrence becomes $y' = ty/(y-1)$ [Bastianello and Fontanier, 2024a]. Here the same form arises from the private coefficient core. The contribution is its explicit application to the two implemented coefficient fibers and the resulting pole and preimage diagnostics. I use "convergence set" rather than "basin of attraction" in the nonattracting $t \leq 1$ regimes.

6 Finite-horizon evidence: when early accuracy misleads

The prespecified design fixes the horizons, accuracy bands, starts, pole neighborhoods, and price grid before outcomes are computed. It follows paths through $K = 1000$, reports an approximation horizon only when a path remains inside the stated band, and retains undefined and pole cases rather than filtering them. Online Resource 1 gives the complete protocol, typed-state rules, error definitions, and exact start list.

A finite-K bounded-reasoning prediction is a distinct economic object, not a numerical failure merely because it differs from equilibrium. Here, "approximation horizon" means only approximation to the relevant

nonzero fixed point. The boundary experiment separately isolates persistence after the first projection, while the pole experiment stops every exact path at its first undefined update.

6.1 Benchmark approximation horizons

At the displayed partial benchmark, the coefficient path first enters and thereafter remains within the 10, 5, 1, 0.1, and 0.01 percent bands at $K = 3, 5, 10, 20, 20$, respectively. From the private certainty-equivalent start, both systems first enter and remain within those bands at $K = 10, 20, 20, 50, 50$. Thus a large first-step error can coexist with rapid eventual convergence at the benchmark. These numerical errors do not by themselves establish a welfare, return, demand, or empirical effect.

The full start ledger retains the zero-vector domain distinction, fixed-point controls, and every prespecified scaling. No start or undefined classification is dropped; the detailed start-by-start results are in Online Resource 1 and Online Resource 2.

Figures 2 and 3 report, respectively, the finite-horizon errors across the deposited sweep and the exact scalar geometry behind the limiting regimes; the accompanying diagnostics are developed in the following two subsections.

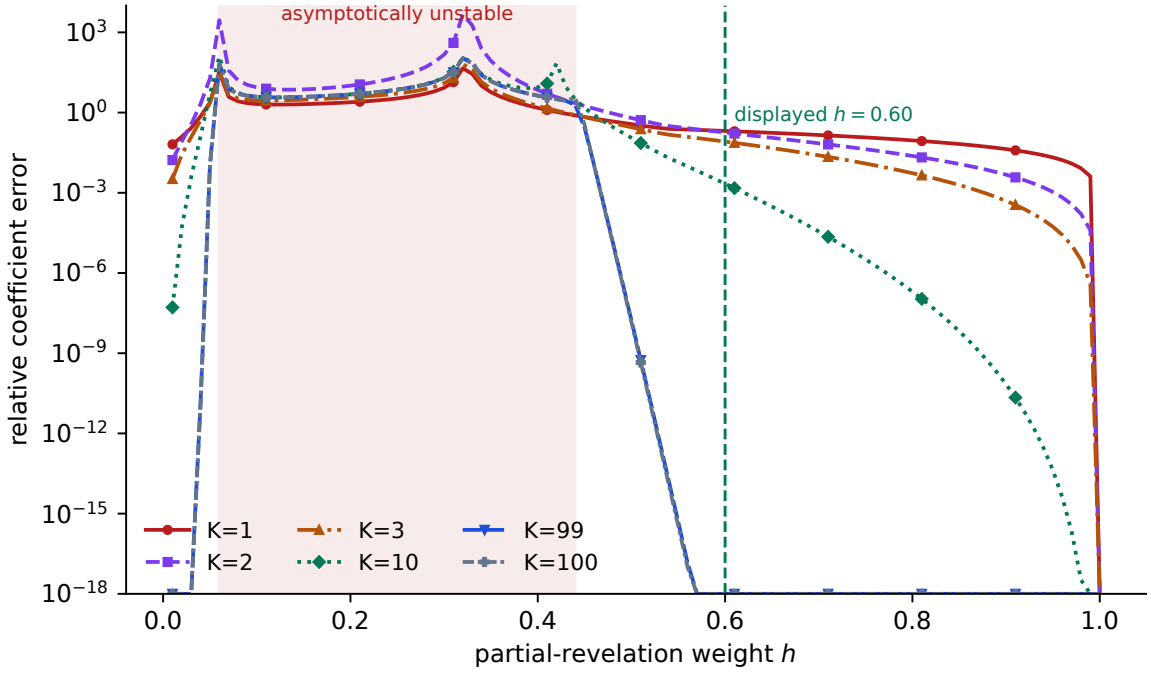


Fig. 2 Finite coefficient errors in the deposited 100-cell sweep. Early closeness and asymptotic stability are distinct statements. All six horizons used in the full census ($K = 1, 2, 3, 10, 99, 100$) are shown. Values below plotting resolution are displayed at the plotting floor; the machine-readable results retain explicit status labels rather than reporting false exact zeros

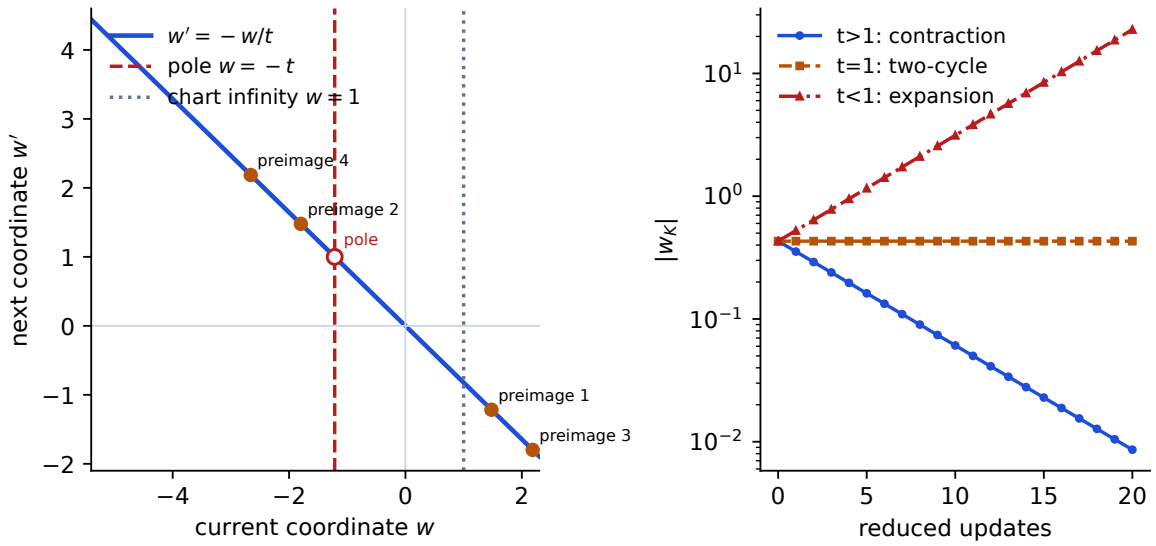


Fig. 3 Scalar geometry. The left panel plots the benchmark map, the true finite-state pole, chart infinity ($w = 1$, corresponding to $y = \infty$ under the coordinate change), and first finite pole preimages. The right panel uses distinct line styles and markers to show contraction for $t > 1$, a two-cycle magnitude for $t = 1$, and expansion for $t < 1$

6.2 Near-boundary and pole behavior

Within the projector image, persistence changes exactly at $|\varepsilon| = 1$: paths above one do not enter a strict accuracy band, the boundary preserves the perturbation, and paths below one enter tighter bands at rates governed by $|\varepsilon|^K$. All five zero-offset pole cases terminate at their analytically prescribed indices; nearby starts can pass arbitrarily close to the denominator singularity and generate large intermediate coefficients. An exact forward-orbit check finds no finite pole at any integer update $K \geq 1$ for the prespecified-start orbit in any of the 100 deposited-sweep cells or 8,100 grid cells. This cell-specific result does not establish orbit safety for other starts or parameter values.

6.3 Finite versus asymptotic approximation

For the full 100-cell source sweep and 8,100-cell grid, exact classifications are paired with finite-horizon errors and persistent-band checks. All 6,267 stable grid paths are within five percent of the nonzero fixed root at $K = 1000$. All 1,833 unstable paths are within five percent of the analytically derived singular boundary limit by $K = 1000$. The latter object is not called an attractor: the source switches branches at zero signal loading, so the limiting boundary is not an ordinary fixed point of the map. The unstable paths make the distinction concrete: a finite iterate can pass a tolerance before the path's limiting classification becomes visible. Early numerical closeness is evidence about a horizon, not proof of convergence. Table 2 collects the headline stability, horizon, and denominator diagnostics.

Table 2: Stability, horizon, and denominator census.

Diagnostic	Exact/result	Interpretation
Deposited sweep	39/100 unstable	$h = 0.06, \ldots, 0.44$
Prespecified variance grid	1,833/8,100 unstable	6,267 stable
Closest grid cell	$\epsilon = 625/626$	margin $1/626$
Displayed $h = 0.60$	$\epsilon = 10/17$	stable
Benchmark 1% persistent horizon	K=10	partial coefficient error
Private 1% persistent horizon	K=20	both coefficient systems
Exact all-horizon partial-orbit check	0 finite poles	8,200 prespecified cells
Exact boundary cells	0	no boundary value assigned to a class

7 What the results mean

The denominator mismatch changes only one private price coefficient at the fixed point, but that coefficient multiplies mean supply and therefore enters the entire prespecified price grid as a level shift. Figure 4 shows three facts. First, the two routines share the same core coefficients. Second, their q coordinates oscillate toward different limits. Third, the deposited-D versus article/updater-H price gap does not decline monotonically with iteration count: the first update produces a much larger gap than either the tenth update or the limiting comparison. Reporting only a converged root or only a selected finite iterate would miss this interaction.

Because mean supply is fixed at one and only q differs, the prespecified-grid sup norm, root-mean-square (RMS) gap, and $\xi = 1$ gap are three measures of the same conditional level shift. The shift scales linearly with mean supply. No empirical distribution over value or supply is imposed, so this is a coefficient-level sensitivity exercise rather than a distributional economic effect. On the separate 101-point analysis grid, the first-update normalized gap is exactly $750000/3204433$; that grid's 23.405077% result is distinct from the headline 16.639312% result under the original-source normalization. The separate-grid fixed-point RMS gap and later path metrics are retained in Online Resource 1 and the machine-readable results, but they determine no classification.

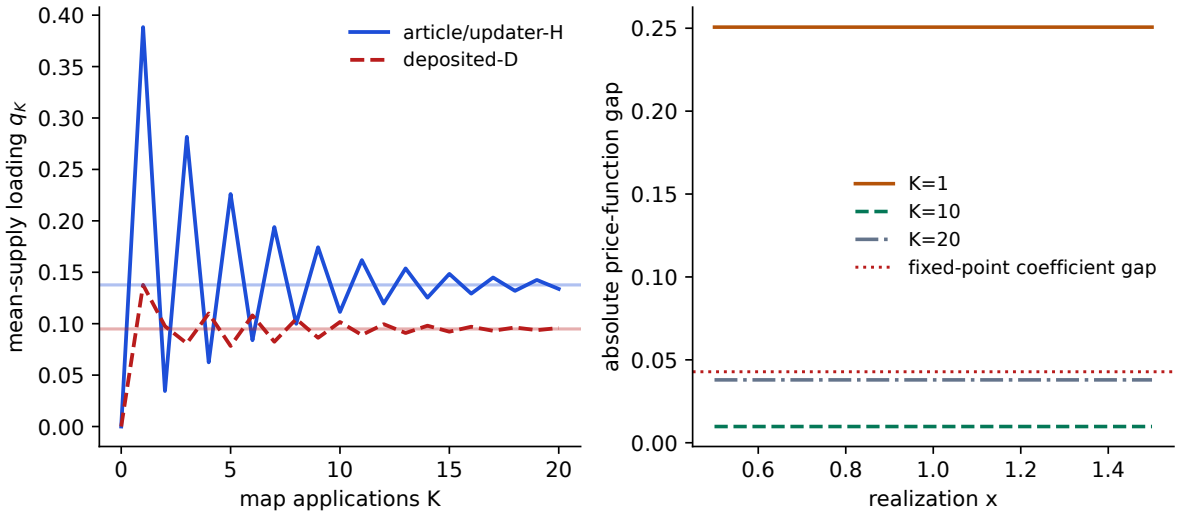


Fig. 4 Deposited-D and article/updater-H private paths from the common prescribed start, distinguished by line style. The price-function gap is constant across ξ at a fixed K because the two paths differ only in the mean-supply loading

The partially revealing results refine the interpretation. Instability is not evidence that the model's mechanism is absent; it is evidence that a particular finite source iteration is not a uniform approximation to its nonzero fixed point. Conversely, stability does not guarantee that one or three iterations approximate the fixed point closely. A responsible finite-reasoning claim should therefore distinguish the behavioral horizon from the numerical approximation horizon and report the map's stability geometry.

8 Contribution boundaries and limitations

Table 3 distinguishes the contribution claimed here from previously established results, excluded conclusions, and source-availability boundaries.

Table 3: Contribution and limitation boundaries.

Category	Included	Excluded
Previously established results	equilibrium equations; recurrence; authors' convergence boundary	no novelty claim
Implementation consistency	deposited-D denominator; exact roots; market-clearing discrepancy	not a general correction of the paper
Finite iteration	prespecified counts, horizons, and pole margins	no tuned thresholds or starts
Bounded source agreement	MATLAB objects 10–11 and coefficients	no independent objects 1–9 reproduction
Source distribution	public citation and custody hashes; independent code	no third-party source files redistributed

The original authors derive the equilibrium equations, rational recurrence, convergence condition in the theoretical model, PET mechanism, displayed benchmark, and broad instability interpretation. I make no

novelty claim for established numerical tools such as fixed-point iteration, spectral-radius checks, projector algebra, exact rational arithmetic, or sensitivity grids. The contribution of this paper is limited to the denominator mismatch in the deposited private fixed-point routine, the exact root and market-clearing check, implementation-specific parameter censuses, the explicit pole and preimage application, and the prespecified finite-horizon normalized-discrepancy results.

Several limitations are substantive. I do not independently reproduce MATLAB figure objects 1–9. The source-faithful evidence comes from one recorded MATLAB R2023b Update 11 execution. I did not rerun that execution, and its log contains non-success diagnostics from seven nonlinear solves. The prespecified grids are broad relative to the benchmark but are not a probability distribution over empirical calibrations. Local projector stability is exact, whereas economic relevance outside the prespecified grid requires a calibration argument. The scalar theorem applies only on the common forward-defined, pole-avoiding domain. Singular boundary limits are not equilibria merely because a stored finite path approaches them.

The exercise is deterministic algebra and a designed parameter census. It has no sampling population, standard error, or statistical estimand; the grid fractions must not be read as empirical frequencies.

9 Conclusion

A numerical equilibrium claim is strongest when the article, a deposited fixed-point residual, and the finite updater path are treated as separately testable objects. Here, the deposited private fixed-point routine solves a system that differs from the article and companion updater. Its output is not read by the plotting driver, so the mismatch is not evidence of an error in the original paper's reported economics. It nevertheless demonstrates why three questions must be answered separately: whether a numerical root satisfies the intended equilibrium conditions, whether the updater is stable, and whether a finite iterate is close to the relevant fixed point.

The analysis also shows why those questions cannot be collapsed into a single convergence label. An oblique first projection can amplify error even when the remaining local dynamics contract; an early iterate can pass a tolerance even when the path is unstable; and a pole can limit the domain before any limiting claim is meaningful. A finite-round economic interpretation should therefore identify both its behavioral horizon and its numerical approximation horizon. The accompanying package reproduces every derived result, table, and figure from independent deterministic code while leaving the recorded deposited-source execution untouched.

A Proofs

Proof of Proposition 1. At a defined fixed point, the v and z equations in (2) imply $z = -Abv$. Substituting into (1) gives $D = b\kappa v^2$ and $H = bv(\kappa v - a)$. The v fixed-point equation becomes $v = atv/(\kappa v - a)$, so the nonzero solution is $v = a(1 + t)/\kappa$. The p_0 equation then gives $p_0 = 1 - v = tb/\kappa$. Substitution into the two respective q equations gives (5). Positivity of primitives and the source-domain denominator conditions exclude additional denominator-safe nonzero roots. □

Proof of Proposition 2. Substituting $x_H^\star$ into (7) gives $\alpha = 152/581$, $\ell = 125/277$, and $\sigma_P^2 = 1216/14525$. Direct substitution into (8) and (9) cancels every numerator exactly. Because the two roots differ only in q, substitution of $x_D^\star$ leaves the first, second, and fourth entries unchanged at zero and the third entry equal to

$$-\ell z^\star - q_D^\star = \frac{5000}{116781}.$$

Dividing by the exact demand normalization $A\sigma_P^2 = 2432/14525$ gives $15625/61104$. Thus both the posterior and market-clearing checks can be verified from the displayed equations without using a deposited residual. □

Proof of Proposition 3. The difference between systems is a constant shift in the prespecified affine price function. On the original $\xi = .20, .21, \ldots, 1.80$ grid, with prior mean, mean supply, and realized supply fixed at one, the fixed-point unnormalized constant-coefficient gap is $5000/116781$ and the article/updater-H fixed-point price sup norm is $2519/2075$; their quotient is $125000/3544233$. At the first update the unnormalized constant-coefficient gap is $15000/59843$ and the contemporaneous article/updater-H price sup norm is $3879/2575$; their quotient is $125000/751233$. Cross multiplication gives $(125000/751233)/(125000/3544233) = 168773/35773$. Applying the same operation to the unnormalized differences $(15000/59843)/(5000/116781)$ gives $603/103$. □

Proof of Proposition 4. The source residual can be reduced to four affine identities after clearing its positive denominators:

$$\begin{aligned} Pp_0 &= bE, & Pp_s &= ah\bar{H}, \\ Pq &= Aabh(1-h), & Pp_Z &= -Aab\bar{H}. \end{aligned}$$

Their coefficient matrix has nonzero determinant on the positive domain, so the displayed root is unique. Substitution into the original rational residuals gives four zero numerators. No point on the source's $p_s = 0$ special branch satisfies this nonzero system. □

Proof of Proposition 5. Differentiate the original source update at (12). Exact factorization gives $J = -\varepsilon\Pi$. Online Resource 1 displays Π in terms of L, U, V and records the executable symbolic certificate. Direct multiplication gives $\Pi^2 = \Pi$ and $\operatorname{tr}\Pi = 2$. A projector's rank equals its trace, hence rank two. Its spectrum is two zeros and two ones, which yields the spectrum of J and $J^2 = -\varepsilon J$. When $0 < h < 1$, $\varepsilon > 0$ and both distinct factors occur, giving minimal polynomial $\lambda(\lambda + \varepsilon)$. At $h = 1$, $\varepsilon = 0$ and $J = 0$, so the minimal polynomial reduces to λ. To settle the boundary rather than infer it from linearization, hold (p_0, p_s, p_Z) at their exact-root values and write $u = q - q^\star$. The original rational update leaves those three coordinates fixed and gives the exact fiber recurrence $u' = -\varepsilon u$. Every nonzero fiber perturbation therefore contracts when $|\varepsilon| < 1$, forms a nonconvergent two-cycle when $\varepsilon = 1$, and grows when $\varepsilon > 1$. Together with the strict spectral condition for the smooth source map, this proves the stated if-and-only-if local-stability classification. □

Proof of Theorem 1. One defined application of either map gives $z' = -Abv'$. On that invariant line, $D = b\kappa v^2$ and $H = bv(\kappa v - a)$, hence

$$v' = \frac{atv}{\kappa v - a}.$$

The definitions of y and w then give (16). Iteration gives $w_n = (-1/t)^{n-1} w_1$. Setting $w_n = -t$ yields $w_1 = (-t)^n$, proving (17).

For the remaining fibers, let $e = p_0 + v - 1$. Direct substitution gives

$$e' = -\frac{e}{y-1},$$
$$q'_H = \frac{-q_H + Abv}{y-1},$$
$$q'_D = -\frac{q_D}{y} + \frac{Aab}{\kappa}.$$

Write $q^\star_H = Aab/\kappa$ and let $u_H = q_H - q^\star_H$. The article/updater recurrence then gives the exact identity

$$u'_H = -\frac{u_H}{y-1}.$$

If $t > 1$, $w_n \to 0$ and hence $y_n \to 1 + t$. Choose any $\rho \in (1, t)$. After a finite, forward-defined prefix, $|y_n - 1| > \rho$; therefore both e_n and $u_{H,n}$ contract at rate at most $\rho^{-1} < 1$. For the deposited-D fiber, put $q^\star_D = Aab(1+t)/[\kappa(2+t)]$ and $u_D = q_D - q^\star_D$. Its recurrence has the form

$$u'_D = a_n u_D + b_n, \qquad a_n = -\frac{1}{y_n} \longrightarrow -\frac{1}{1+t}, \qquad b_n \longrightarrow 0.$$

Thus $|a_n| \le r < 1$ eventually, and a standard geometric-convolution bound for this affine recurrence gives $u_{D,n} \to 0$. Together with $v_n \to a(1+t)/\kappa$, $z_n = -Abv_n$, and $p_{0,n} = 1 - v_n + e_n$, this proves convergence of every finite common forward-defined orbit when $t > 1$.

On the invariant line, e and u_H have multiplier $-1/t$, while u_D has multiplier $-1/(1+t)$. If $t = 1$, the reduced coordinate alternates unless $w_1 = 0$; if $t < 1$, it expands unless $w_1 = 0$. On the fixed core, convergence of e requires $e = 0$, and convergence of the article/updater-H q coordinate also requires $q = q^\star_H$. By contrast, the deposited-D q coordinate converges from every finite initial q. This proves the stated inclusion. The forward-defined restriction excludes precisely the finite pole preimages and other source-domain denominator failures. □

B Closed partial source orbit

For the deposited zero initialization, let

$$\delta = \frac{bh(1-h)}{P}, \qquad x_K = (-\varepsilon)^{K-1}, \quad K \ge 1.$$

After the special first update the stored orbit has

$$p_{s,K} = \frac{p^\star_s}{1 + \delta x_K}, \quad p_{0,K} = 1 - p_{s,K}, \quad p_{Z,K} = -\frac{Ab}{h} p_{s,K}, \quad q_K = \frac{q^\star + r x_K}{1 + \delta x_K},$$

where $r = q_1(1+\delta) - q^\star$ and q_1 is the mean-supply coefficient after the source's special first update. This form produces the exact stability census and the singular-boundary limit. If $\varepsilon > 1$, $|x_K|$ diverges; the orbit

approaches $(1, 0, r/\delta, 0)$ along parity subsequences when defined. Because the source uses a special branch at $p_s = 0$, this limit is not labeled a fixed point or attractor.

C Reproducibility

The reported-number ledger documents two corrections to the analysis specification: the normalized-gap ratio is kept distinct from the raw coefficient-gap ratio, and the unstable partial path is classified by its branch-dependent singular boundary rather than by the zero vector. Neither correction changes a parameter grid, horizon, accuracy band, or initial condition.

Running `make rebuild` checks the exact symbolic identities, regenerates results, renders exactly four figures and four table assets with Matplotlib (three tables appear in the main text, while the source-custody table remains supporting detail), builds the identified and contingency anonymous article and supplement PDFs, and checks the resulting manifests. JSON represents undefined states with typed strings; it rejects NaN and infinity. Absolute local paths, credentials, secrets, and private account or control identifiers are excluded from the reviewer archive.

References

S. Borağan Aruoba, Jesús Fernández-Villaverde, and Juan F. Rubio-Ramírez. Comparing solution methods for dynamic equilibrium economies. *Journal of Economic Dynamics and Control*, 30(12):2477–2508, 2006. doi: 10.1016/j.jedc.2005.07.008. URL https://doi.org/10.1016/j.jedc.2005.07.008.

Francesca Bastianello and Paul Fontanier. Online appendix for: Expectations and learning from prices, 2024a. URL https://paulfontanier.github.io/papers/onlineappendix_PETinGE_BF_Restud.pdf. Sections D, F, and G.

Francesca Bastianello and Paul Fontanier. Replication package for: Expectations and learning from prices. Zenodo, version v1, 2024b. URL https://doi.org/10.5281/zenodo.10780393.

Francesca Bastianello and Paul Fontanier. Expectations and learning from prices. *The Review of Economic Studies*, 92(3):1341–1374, 2025. doi: 10.1093/restud/rdae059. URL https://doi.org/10.1093/restud/rdae059.

Francesca Bastianello and Paul Fontanier. Partial equilibrium thinking, extrapolation, and bubbles. *The Review of Financial Studies*, 2026. doi: 10.1093/rfs/hhag053. URL https://doi.org/10.1093/rfs/hhag053. Published online 3 June 2026.

Lawrence E. Blume, Margaret M. Bray, and David Easley. Introduction to the stability of rational expectations equilibrium. *Journal of Economic Theory*, 26(2):313–317, 1982. doi: 10.1016/0022-0531(82)90006-0. URL https://doi.org/10.1016/0022-0531(82)90006-0.

Mateo Bodon. Reproduction materials for when finite learning looks reliable: Equilibrium consistency and iteration dynamics in learning from prices. GitHub release, version 1.0.2, 2026. URL https://github.com/MateoBodon/learning-from-prices-reproduction/releases/tag/v1.0.2. Release archive SHA-256: 724ee05465804ce11d77c9e897aebb04fd9742abf8901c414731ed0d16eb773c.

Margaret M. Bray. Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory*, 26(2):318–339, 1982. doi: 10.1016/0022-0531(82)90007-2. URL https://doi.org/10.1016/0022-0531(82)90007-2.

Andrés Carvajal and Hang Zhou. Learning to bet (rationally) with logs. *Journal of Economic Dynamics and Control*, 180:105181, 2025a. doi: 10.1016/j.jedc.2025.105181. URL https://doi.org/10.1016/j.jedc.2025.105181.

Andrés Carvajal and Hang Zhou. On the sophistication of financial investors and the information revealed by prices. *Journal of Dynamics and Games*, 12(4):406–414, 2025b. doi: 10.3934/jdg.2024024. URL https://doi.org/10.3934/jdg.2024024.

Garret Christensen and Edward Miguel. Transparency, reproducibility, and the credibility of economics research. *Journal of Economic Literature*, 56(3):920–980, 2018. doi: 10.1257/jel.20171350. URL https://doi.org/10.1257/jel.20171350.

Jean-Michel Grandmont. Expectations formation and stability of large socioeconomic systems. *Econometrica*, 66(4):741–781, 1998. doi: 10.2307/2999573. URL https://doi.org/10.2307/2999573.

Robert Kirkby. Quantitative macroeconomics: Lessons learned from fourteen replications. *Computational Economics*, 61:875–896, 2023. doi: 10.1007/s10614-022-10234-w. URL https://doi.org/10.1007/s10614-022-10234-w.

Christophe Pérignon, Olivier Akmansoy, Christophe Hurlin, Anna Dreber, Felix Holzmeister, Jürgen Huber, Magnus Johannesson, Michael Kirchler, Albert J. Menkveld, Michael Razen, and Utz Weitzel. Computational reproducibility in finance: Evidence from 1,000 tests. *The Review of Financial Studies*, 37 (11):3558–3593, 2024. doi: 10.1093/rfs/hhae029. URL https://doi.org/10.1093/rfs/hhae029.

Hang Zhou. Informed speculation with k-level reasoning. *Journal of Economic Theory*, 200:105384, 2022. doi: 10.1016/j.jet.2021.105384. URL https://doi.org/10.1016/j.jet.2021.105384.

Statements and declarations

Funding. The author received no funds, grants, or other support for this work.

Competing interests. The author declares no competing interests.

Author contribution. Mateo Bodon completed the conceptualization, methodology, formal analysis, software, validation, investigation, visualization, and writing.

Use of generative AI tools. Generative AI tools served as supplementary research and technical aids during literature review, for assistance with programming syntax and package use during code development, and for secondary and adversarial checks of mathematical derivations, code, and reproducibility materials. Mateo Bodon determined the research questions, methods, mathematical arguments, implementation, interpretation, and wording; independently reviewed the cited sources and every reported equation, result, figure, table, code file, and passage in the prepared package; and accepts full responsibility for the work.

Data availability. No new empirical data were collected or analyzed. The examined public replication package is available on Zenodo at https://doi.org/10.5281/zenodo.10780393. Derived numerical results are included in the accompanying reviewer-safe reproduction archive.

Code availability. Independent deterministic code and the derived results, figures, and tables are available in the public version 1.0.2 release of Bodon [2026] and in the accompanying Online Resource 2 reviewer reproduction archive. The release does not redistribute or execute the original third-party source code; it instead gives the source's public citation and cryptographic hashes.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM1SupplementaryAppendix.pdf
- ESM2ReviewerReproduction.zip