

SUPPLEMENTARY MATERIALS: Computational Validation of a Kinetic PDE for Reinforcement Learning in Market-Entry Games

Esther Bou Dagher *

SM1. Supplementary model and analytical details. This section records the probability-map checks and analytical details omitted from Section 2 for brevity.

SM1.1. Probability-map geometry and analytical assumptions. The analytical results require an increasing map $p : \mathbb{R} \rightarrow (0, 1]$ satisfying

$$(SM1.1) \quad p \in C^2(\mathbb{R}), \quad 0 \leq p'(x) \leq c_p(1 - p(x)), \quad |p''(x)| \leq c_p(1 - p(x)).$$

These are the probability-map assumptions used in the well-posedness analysis of [1], used in the a priori estimates and strong-solution theory leading to Theorem 4.4. The sorting result, Theorem 5.1(i), additionally assumes

$$(SM1.2) \quad |p''(x)| \leq c_p p'(x),$$

which is the additional curvature condition used in the long-time sorting analysis of [1]. For Branch A, (2.7),

$$\inf_{x \in \mathbb{R}} p_A(x) = 0.05,$$

and direct differentiation gives

$$(SM1.3) \quad p'_A(x) \geq 0, \quad p'_A(x) \leq \lambda(1 - p_A(x)), \quad |p''_A(x)| \leq \lambda^2(1 - p_A(x)), \quad |p''_A(x)| \leq \lambda p'_A(x).$$

Thus $p_A(x)$ satisfies the probability-map regularity and curvature conditions used in Theorems 4.4 and 5.1(i), motivating its role as the *theorem-facing* branch. For aggregate learning, Theorem 5.1(ii) further assumes $0 < p_{\min} < \kappa$ and the transport-dominance condition

$$(SM1.4) \quad \hat{c}_p \sqrt{\tau}(M - 1) < p_{\min}^4,$$

which is the transport-dominance condition used to control the long-time capacity error in [1]. These conditions bound the long-time capacity error; a further smallness condition in the theorem then places the limiting aggregate entry probability inside the Nash-equilibrium interval. Branch B, (2.8), is smooth and increasing, but

$$(SM1.5) \quad \inf_{x \in \mathbb{R}} p_B(x) = 0.$$

It therefore falls outside the positive-floor assumption of Theorem 5.1(ii), so the analytical guarantee of aggregate learning does not apply directly. This does not exclude the sorting conclusion in Theorem 5.1(i); Branch B is retained to test numerically whether PDE agreement, aggregate learning, and sorting persist when the positive-floor assumption is relaxed.

*CEREMADE (CNRS UMR 7534), PSL University, Université Paris-Dauphine, Place du Maréchal de Lattre de Tassigny, 75775 Paris Cedex 16, France (esther.bou-dagher@dauphine.psl.eu).




Figure SM1. Probability-map geometry for Branches A and B, showing $p(x)$, $p'(x)$, and $p''(x)$. Branch A has a positive whole-line floor, whereas Branch B approaches zero as $x \rightarrow -\infty$. The derivative panels show how the two maps convert the same propensity change into different local changes in entry probability.

SM2. Supplementary computational design and calibration. This section records the design details omitted from Section 3.1. The main text keeps the capacity rule, the baseline template, the moment-matching condition, and the capacity-matched set because these objects are used throughout the paper. Here we give the component interpretation, the pushforward calibration plot, and the full computational design table.

SM2.1. Initial-law component interpretation. The baseline law in (3.2) is a smooth, mildly multimodal test density. The coefficients 0.80, 0.10, 0.10 are relative component amplitudes rather than exact probability masses, because the three Gaussian components have different widths. The central component has width 0.85, while the two side components have width 0.45 and centres displaced by ± 1.20 from c . All of these quantities were frozen before particle-PDE discrepancies were inspected.

The role of the template is to give a non-Gaussian but controlled initial law. A single Gaussian would mainly test aggregate motion of one smooth bump; the side components make the comparison sensitive to the internal placement of probability mass. After normalisation, the centre $c = c_*$ is calibrated separately for each branch and population size using (3.3).

The calibration audit records total mass, $|a(0)|$, $\beta(0)$, and the pushforward $\nu_0 = p_{\#}f_0$. Across the 2×100 baseline cases, the largest residual in the moment-matching condition was

$$8.31 \times 10^{-14}.$$

This is at floating-point/root-solver precision and is negligible relative to the operational tolerances used later.

SM2.2. Frozen computational design registry. Table SM1 summarises the principal computational layers. A case is a distinct combination of the non-seed factors shown. Each reduced design was fixed in its case registry before the corresponding diagnostic was evaluated.




Figure SM2. Frozen initial-law calibration. The left panels show the translated propensity densities and the right panels show their probability-space pushforwards. Integer-capacity variation causes only small translations in propensity space, whereas the two probability maps generate different probability laws despite the common calibration $\int p(x)f_0(x) dx = \kappa(M)$.

Table SM1

Principal frozen computational designs.

Layer	Branches	M values	ρ values	Seeds	Reports	Purpose
PDE verification	2	7	fixed	deterministic	4	four-level space-time refinement
Particle verification	2	100	9	4	5	1800 cases; 7200 trajectories
Finite-time comparison	2	21	9	2	31	378-case observable atlas
Oracle and closure	2	13	4	8	31	104 paired mechanism cases
Direct jumps	2	13	4	6	31	624 case-seed simulations
Off-capacity bridge	2	selected	4	up to 20	common	attraction, noise, and censoring
Exact generator	2	$M \leq 300$	selected	exact + MC	4	six weak observables; 384 rows

SM3. Supplementary details for the deterministic implementation.

SM3.1. Finite-volume flux and refinement settings. This subsection gives the implementation details behind (3.5). The advective flux uses MUSCL reconstruction [3] with the minmod limiter [2]. The limited slope is

$$\sigma_i = \minmod(g_i - g_{i-1}, g_{i+1} - g_i).$$

It uses neighbouring cell values to reconstruct the left and right states at an interface,

$$g_{i+1/2}^- = g_i + \frac{1}{2}\sigma_i, \quad g_{i+1/2}^+ = g_{i+1} - \frac{1}{2}\sigma_{i+1}.$$

The upwind state is then chosen according to the sign of $c(t)$, so the advective flux uses information from the side from which the flow arrives. The limiter suppresses artificial oscillations

near steep gradients while retaining second-order behaviour in smooth regions. The diffusive flux is centred,

$$F_{i+1/2}^{\text{diff}} = -d(t) \frac{g_{i+1} - g_i}{\Delta x},$$

which is the symmetric interface approximation to $\partial_x g$. Since diffusion has no preferred direction, no upwind choice is used. No positivity projection is applied, so conservation error, negative mass, and boundary mass remain visible diagnostics rather than quantities corrected after the solve. The method-of-lines system is integrated with SciPy’s variable-order BDF method [4]. The local error is controlled against

$$\text{atol} + \text{rtol}|f_i|,$$

where atol controls small values and rtol controls error relative to the size of f_i . The fixed refinement settings are listed in Table SM2. The maximum step is restricted in scaled time over $\theta = 2$; this is not an explicit-scheme stability condition, but a safeguard so that time error is reduced together with space error.

Table SM2

Deterministic PDE refinement settings used in the independent solver audit.

Grid level	N_x	atol	rtol	N_t in $\Delta\theta_{\max} = 2/N_t$
1	101	10^{-5}	10^{-7}	80
2	201	2×10^{-6}	2×10^{-8}	120
3	401	5×10^{-7}	5×10^{-9}	160
4	801	10^{-7}	10^{-9}	240

The 801-cell grid is the reference grid after the audit. Coarser solutions are interpolated to its cell centres before common-grid L^1 comparisons, so that densities are compared at the same spatial locations. Bounded-window masses are also computed on their native grids, keeping threshold–cell alignment effects visible. For a diagnostic Q , let U_h be the numerical PDE solution computed with mesh width h . Define

$$E_h = \|U_h - U_{h/2}\|_Q, \quad E_{h/2} = \|U_{h/2} - U_{h/4}\|_Q, \quad q = \frac{\log(E_h/E_{h/2})}{\log 2}.$$

The value $q \simeq 2$ means that halving the mesh width reduces the observed difference by about a factor of four. The finest adjacent-grid differences are used as deterministic resolution scales in the operational thresholds (3.16).

SM3.2. Supplementary details for the weak-generator audit. This subsection gives the construction behind the weak-generator hierarchy in (3.21). The audit is local: it freezes one stored microscopic state and measures how each approximation changes the expected next-round evolution of a smooth observable. This is different from $D_{W_1,p}$, which is a trajectory-level law discrepancy accumulated over many updates. Let $\mathbf{X}_n = (X_n^1, \dots, X_n^M)$ be fixed. Conditional on this state, the probabilities $p(X_n^1), \dots, p(X_n^M)$ are fixed and only the next Bernoulli actions are random. For focal agent i , the one-step weak generator is the conditional expectation of the increment of $\psi(X_n^i)$, divided by τ . The implementation evaluates this for

every focal agent and averages uniformly over $i = 1, \dots, M$. For the focal agent, $K_{-i} = \sum_{j \neq i} \delta_n^j$, $\bar{p}_{-i} = \frac{1}{M-1} \sum_{j \neq i} p(X_n^j)$. Here LOO means “leave one out”: the focal agent is removed from the attendance count because its own action is handled separately. Since the other agents usually have unequal entry probabilities, K_{-i} is Poisson–binomial. This exact distribution is computed by dynamic programming for $M \leq 300$, so the primary generator values contain no Monte Carlo sampling noise. The five operators are

- G_{micro} : exact finite-player conditional update,
- G_{LOO} : replace K_{-i} by $\text{Binomial}(M-1, \bar{p}_{-i})$,
- G_{nonlocal} : replace $\bar{p}_{-i}$ by the global mean $\alpha = \frac{1}{M} \sum_j p(X_n^j)$,
- G_{KM2} : replace the exact finite shift by its second-order Taylor form,
- G_{PDE} : evaluate the coded PDE weak form.

Thus $G_{\text{micro}} - G_{\text{LOO}}$ measures the effect of unequal other-agent probabilities beyond their leave-one-out mean. The next difference, $G_{\text{LOO}} - G_{\text{nonlocal}}$, measures the local effect of replacing focal leave-one-out information by the global closed-model moment. Indeed, $\bar{p}_{-i} = \frac{M\alpha - p(X_n^i)}{M-1}$, so different focal agents can see slightly different environments before this replacement. The term “nonlocal” means that the exact finite shift $\psi(x + \Delta x) - \psi(x)$ is still retained. The Kramers–Moyal step is introduced only in G_{KM2} , where this finite shift is replaced by $\psi'(x)\Delta x + \frac{1}{2}\psi''(x)(\Delta x)^2$. Therefore $G_{\text{nonlocal}} - G_{\text{KM2}}$ is the local finite-jump truncation component. Finally, $G_{\text{KM2}} - G_{\text{PDE}}$ checks that the second-order weak form and the coded PDE operator agree up to floating-point error. Independent one-step Monte Carlo simulations are used only as validation of the exact quadrature: they redraw the next Bernoulli actions from the same frozen state and estimate the same conditional expectation. The telescoping identity (3.22) is exact for the signed local generator components, but these components should not be interpreted as additive percentages of the global Wasserstein error.

SM3.3. Independent numerical verification. The deterministic and microscopic implementations were audited before any matched particle–PDE discrepancy was interpreted. Table SM3 reports the branchwise deterministic audit, while Table SM4 reports the microscopic implementation and Bernoulli-calibration checks. In Table SM4, the overall entries are sums for the case and trajectory counts and maxima for the numerical residuals. The zero negative-mass result in Table SM3 is not produced by clipping, because no positivity projection is applied.

Table SM3

Independent deterministic PDE verification by probability branch. The audit used $N_x \in \{101, 201, 401, 801\}$ over seven population cases per branch.

Diagnostic	Branch A	Branch B
Cases passed	7/7	7/7
Observed q range	1.985658–2.000602	1.978483–2.013186
Maximum mass defect	1.887379×10^{-15}	1.998401×10^{-15}
Maximum negative mass	0	0
Maximum outer-domain mass	7.038239×10^{-28}	5.684786×10^{-26}

Table SM4

Independent microscopic-simulator verification. Each branch contains nine ρ levels, 100 population sizes, and four seeds per case. The predeclared Bernoulli screen is $|Z_{\text{global}}| \leq 4$.

Diagnostic	Branch A	Branch B	Overall
Cases passed	900/900	900/900	1800/1800
Seeded trajectories	3600	3600	7200
Maximum update residual	4.440892×10^{-16}	2.220446×10^{-16}	4.440892×10^{-16}
Maximum attendance residual	0	0	0
Maximum $ Z_{\text{global}} $	2.672861	3.059714	$3.059714 < 4$

SM3.4. Complete finite-time diagnostics. The verification tables above support the independent audits reported in Section 4.1. The following figures supplement the matched particle–PDE comparisons in Sections 4.2–4.3. Figure SM3 shows the complete (ρ, M) atlas, while Figure SM4 isolates the dependence on ρ at three representative population sizes.




Figure SM3. Complete finite-time diagnostics. Panel (a) shows the raw discrepancy atlases for Branches A and B. Aggregate errors are concentrated at small M and large ρ and broadly decrease with population size. The Wasserstein discrepancy also exhibits a small-population excess and a branch-dependent level at large M . Panel (b) shows discrepancies divided by their operational thresholds; values below one pass the declared criterion. The α and β errors pass throughout, as does the Branch A W_1 error, whereas the Branch B W_1 discrepancy becomes resolvable as the finite-sampling allowance contracts. Panel (c) shows the 40- and 100-bin half- L^1 diagnostics. These bin-dependent quantities provide stringent shape checks but are not interpreted as invariant continuum total-variation distances.




Figure SM4. Dependence on ρ at representative population sizes $M = 30, 915$, and 30000 . Rows correspond to Branches A and B, and columns to D_α , D_β , and $D_{W_{1,p}}$. The aggregate errors increase with ρ but remain markedly smaller at larger M . The Wasserstein discrepancy is strongly ρ -dependent at $M = 30$ and nearly flat at $M = 30000$. Because $\theta_{\max} = 1000\rho^2$, varying ρ changes both jump stress and observation time; the curves therefore describe their combined effect, not a fixed-horizon jump-size intervention.

SM4. Supplementary mechanism diagnostics. This supplement expands the robustness and scope checks supporting Section 5. The corrected covariance statistic is finite-dimensional and resolution-dependent; the aggregate-path calculation is an in-sample diagnostic intervention; unconditional and conditional-positive jump quantiles answer different questions; and the exact generator components are local one-step differences rather than additive shares of the trajectory-level Wasserstein distance.

Table SM5

Corrected closure regression and finite-time oracle calibration. HC3 uncertainty is reported for the regression.

Response	Term	Coef.	HC3 SE	95% low	95% high	p
$\log \chi_{\text{ex}}^+$	$\log M$	-1.001017	0.003876	-1.008708	-0.993326	4.65×10^{-143}
$\log \chi_{\text{ex}}^+$	$\log \rho$	0.129224	0.010808	0.107782	0.150667	5.45×10^{-21}
$\log \chi_{\text{ex}}^+$	Branch B	0.049603	0.015826	0.018205	0.081002	0.00226
Oracle relative improvement: $q_{0.10}$		-0.043261				
Oracle relative improvement: median		-0.002314				
Oracle relative improvement: $q_{0.75}$		-0.000195				
Oracle relative improvement: $q_{0.90}$		0.000252				




Figure SM5. Resolution and reliability checks for corrected finite-player dependence. The score χ_{ex}^+ is computed at $K = 10, 20, 40$, with each resolution using its own independent-multinomial baseline. Although its numerical level depends on K , the nearly parallel log-log trends preserve the approximate M^{-1} decay. Marked points fail $M \geq 5K$ and are omitted only from the reliability-restricted fit.

Table SM6

Primary unconditional and conditional-positive jump regressions.

Response	Term	Coef.	HC3 SE	95% low	95% high	p
unconditional $\log J_{90}$	$\log \rho$	0.992954	0.004304	0.984416	1.001493	3.57×10^{-138}
unconditional $\log J_{90}$	$\log M$	0.000604	0.001853	-0.003072	0.004280	0.745
unconditional $\log J_{90}$	Branch B	0.133054	0.005964	0.121222	0.144886	1.30×10^{-40}
conditional-positive $\log J_{90}$	$\log \rho$	0.994065	0.003271	0.987576	1.000554	$< 10^{-100}$
conditional-positive $\log J_{90}$	$\log M$	-0.004653	0.001395	-0.007420	-0.001887	0.00119
conditional-positive $\log J_{90}$	Branch B	0.119307	0.004537	0.110307	0.128308	$< 10^{-40}$




Figure SM6. (a) Fraction of particle-round observations with $\Delta p = 0$ at the smallest and largest tested ρ values. (b) Unconditional versus conditional-positive case-level 90th percentiles; points above the equality line show the larger move size after zero updates are removed. (c) J_{90}/ρ against M ; approximate flatness is the visual counterpart of the negligible residual M exponent in the unconditional fit. (d) Relative seed-to-seed interquartile range versus source-stream size. The absence of systematic growth checks that the fixed reservoir remains adequate as the full stream becomes larger.

Table SM7

Associative law-error regression and accompanying diagnostic checks.

(a) Primary associative regression

Predictor	Coef.	HC3 SE	95% CI		p
			Lower	Upper	
J_{90}	1.221759	0.884402	−0.533307	2.976825	0.170
χ_{ex}^+	0.236905	0.053209	0.131313	0.342496	2.25×10^{-5}
$\log M$	0.000089	0.000113	−0.000135	0.000313	0.436
$\log \rho$	0.000128	0.000520	−0.000904	0.001160	0.806
Branch B indicator	0.004413	0.000309	0.003800	0.005027	1.14×10^{-25}

(b) Nested explanatory power

Model	R^2	ΔR^2
Branch + scales	0.697708	—
+ χ_{ex}^+	0.885224	0.187516
+ J_{90}	0.887941	0.002717
+ J_{90}^+	0.888873	0.000932

(c) Two-jump collinearity audit

Predictor	R_j^2	VIF
J_{90}	0.998981	981.063
J_{90}^+	0.998998	997.965
χ_{ex}^+	0.582435	2.395
$\log M$	0.625613	2.671
$\log \rho$	0.915821	11.879
Branch B indicator	0.135393	1.157

Notes. Panel (a) reports the primary HC3 regression. Panel (b) adds predictor blocks sequentially. Panel (c) deliberately includes both J_{90} and J_{90}^+ to diagnose their near-redundancy; this two-jump specification is not used as the primary model. Here R_j^2 is obtained by regressing predictor j on the remaining predictors, and $\text{VIF}_j = (1 - R_j^2)^{-1}$.




Figure SM7. (a) Pearson correlations among unconditional jump size, conditional-positive jump size, corrected dependence, and the closed-PDE W_1 error. (b) Nested-model R^2 . Adding corrected dependence to branch and scale variables produces the principal increase; adding a jump statistic produces only a small additional increase because ρ already records most of its variation.

Table SM8

Restricted exact weak-generator and Monte-Carlo audit. Dominance is defined row-wise across the 384 generator rows by the largest absolute telescoping component; local-magnitude medians are taken across the 60 observable-by-observable grouped summaries.

Audit item	Numerator	Denominator	Result
LOO/global replacement dominant	378	384	98.44%
Attendance heterogeneity dominant	4	384	1.04%
Finite-jump truncation dominant	2	384	0.52%
Median LOO/global	1.494×10^{-2}		local magnitude
Median heterogeneity	1.039×10^{-4}		local magnitude
Median finite jump	3.209×10^{-8}		local magnitude
Independent exact–MC rows	64	64	completed
Maximum absolute standardised residual	≈ 2.38		cross-check
Maximum Monte-Carlo standard error	0.0357		precision audit




Figure SM8. Independent validation of the exact microscopic generator. **(a)** Exact Poisson–binomial values against Monte Carlo estimates from 10^5 – 10^6 one-step simulations; the diagonal is perfect agreement and the intervals are Monte Carlo uncertainty. **(b)** Distribution of the standardised differences Z_{MC} ; the dotted lines mark ± 2 . The largest absolute value is approximately 2.38 over 64 checks, which is compatible with ordinary simulation variability.

SM5. Supplementary return-to-capacity diagnostics. This supplement expands the robustness and scope checks supporting Section 6. It reports the time-resolved deterministic return paths and initial-shape sensitivity, the finite-population particle–PDE bridge, and censoring-aware persistent-entry diagnostics. These analyses distinguish endpoint contraction of the systematic capacity error from ordinary finite- M fluctuations and from agreement of the complete propensity law.




Figure SM9. The left panels show the complete horizon and the right panels enlarge the initial transient. The vertical axis is the remaining fraction $|a(\theta)|/|a(0)|$ of the initial aggregate capacity error. The line at one marks the starting error and the line at 10^{-2} marks the strong-return level used in the endpoint screen. The near-vertical appearance on the full horizon reflects rapid adjustment relative to the plotting window, not an instantaneous jump.

Table SM9

Initial-shape sensitivity of the final-to-initial aggregate capacity-error ratio. The PDE and microscopic rows use the fixed designs described in the text and are not pooled into one fit.

System	Branch	ρ	M	Initial shape	C_a
PDE	A	0.02	30	bimodal	0.000625
PDE	A	0.02	30	tilt	0.000391
PDE	A	0.02	30	translated	0.000308
PDE	A	0.02	300	bimodal	0.000035
PDE	A	0.02	300	tilt	0.000031
PDE	A	0.02	300	translated	0.00003
PDE	B	0.02	30	bimodal	0.001233
PDE	B	0.02	30	tilt	0.001091
PDE	B	0.02	30	translated	0.001026
PDE	B	0.02	300	bimodal	0.000169
PDE	B	0.02	300	tilt	0.000155
PDE	B	0.02	300	translated	0.00015
Microscopic seed mean	A	0.1	30	bimodal	0.1831
Microscopic seed mean	A	0.1	30	tilt	0.1806
Microscopic seed mean	A	0.1	30	translated	0.1456
Microscopic seed mean	B	0.1	30	bimodal	0.7191
Microscopic seed mean	B	0.1	30	tilt	0.2832
Microscopic seed mean	B	0.1	30	translated	0.1384

The Branch B microscopic bimodal value is a single finite-design observation, not a general shape theorem.

Table SM10

Principal finite-population bridge at $\rho = 0.1$, $\delta_\alpha = 0.05$, and tilted initial law. Aggregate, local-window, and law-level metrics are shown separately.

Branch	M	C_a	Centred α RMSE	Window S_R RMSE	Law W_1 RMSE	Late RMS of a
A	30	0.1806	0.0371	0.4519	0.0678	0.0066
A	100	0.0558	0.0275	0.3685	0.0509	0.0035
A	300	0.0152	0.0172	0.3209	0.038	0.0007
B	30	0.2832	0.0435	0.2253	0.079	0.0184
B	100	0.1434	0.0377	0.1637	0.0564	0.0056
B	300	0.0382	0.0167	0.098	0.0361	0.002




Figure SM10. Representative finite-population aggregate trajectories for the principal displaced cases. Solid curves are microscopic seed means, shaded bands are one seed standard deviation, and dashed curves are the case-matched closed PDE. Increasing M reduces the aggregate path error and stochastic spread, although Branch B retains a larger late fluctuation level.




Figure SM11. Censoring-aware persistent entry of displaced microscopic paths. (a) Entry into the fixed deterministic tube $|a| \leq 10^{-3}$. (b) Entry into the independently calibrated finite- M noise tube. Curves show the cumulative fraction of seeds that have entered and remained inside the relevant tube. Curves ending below one indicate right-censoring, not events assigned at the final horizon.

Table SM11

Independent finite- M fluctuation calibration and censoring-aware persistent entry for the principal displaced cases. A missing median means that fewer than half of the seeds entered before the displayed horizon.

Branch	M	Late RMS	Noise threshold	Horizon	Fixed entry	Fixed median	Noise entry	Noise median
A	30	0.002419	0.005185	20	0%	not reached	0%	not reached
A	100	0.001834	0.003555	20	0%	not reached	100%	6.208
A	300	0.000517	0.000855	2.9	100%	0.55	100%	1.7
B	30	0.012256	0.015296	20	0%	not reached	20%	not reached
B	100	0.005119	0.006316	20	0%	not reached	100%	0.05
B	300	0.002002	0.00227	20	35%	not reached	55%	19.378

Table SM12

Targeted grid and domain repetitions. Rate differences are interpreted only for retained finite-window fits; entry differences are evaluated at the stored temporal resolution.

Branch	ρ	M	Variant	$ \Delta a_{\text{final}} $	Rate rel. diff.	Entry diff.	Domain safe	Status
A	0.0025	30	dx_factor_0.5	3.03×10^{-9}	0.01333	0	yes	success
A	0.0025	30	domain_factor_1.25	6.29×10^{-13}	0	0	yes	success
A	0.0025	300	dx_factor_0.5	6.88×10^{-11}	0.00261	0	yes	success
A	0.0025	300	domain_factor_1.25	1.31×10^{-10}	0.00022	0	yes	success
A	0.0025	3000	dx_factor_0.5	2.59×10^{-12}	0.00081	0	no	success
A	0.0025	3000	domain_factor_1.25	7.18×10^{-10}	0.00916	0	no	success
A	0.02	30	dx_factor_0.5	2.43×10^{-8}	0.01334	0	yes	success
A	0.02	30	domain_factor_1.25	9.87×10^{-12}	0.00001	0	yes	success
A	0.02	300	dx_factor_0.5	3.08×10^{-12}	0.00258	0	yes	success
A	0.02	300	domain_factor_1.25	2.65×10^{-10}	0.00009	0	yes	success
A	0.02	3000	dx_factor_0.5	2.60×10^{-11}	0.00002	0	no	success
A	0.02	3000	domain_factor_1.25	5.99×10^{-9}	0.00788	0	no	success

SM6. Supplementary long-time sorting diagnostics. This supplement expands the robustness and scope checks supporting Section 7. It reports the boundary-safe bounded-window crossings and nested-window trajectories, the population-size scaling and spatial-resolution sensitivity of their timing, and the associated late-time mechanism and analytical diagnostics. These analyses distinguish rapid aggregate learning from slower outward redistribution while monitoring finite-domain effects.

Table SM13

Boundary-safe persistent-crossing summary for the eight Branch A long-time trajectories, $\rho \in \{0.0025, 0.02\}$ and $M \in \{30, 300, 3000, 30000\}$. A clean case crosses the threshold persistently while $B_L(\theta) < 5 \times 10^{-3}$.

Window	Threshold	Clean cases	Scaled crossing-time range
S_{10}	0.10	8/8	2.551×10^5 – 2.704×10^5
S_{10}	0.05	6/8	1.068×10^6 – 1.079×10^6
S_{50}	0.10	0/8	–
S_{100}	0.10	0/8	–


Corrected-clock long-time hierarchy: Branch A, $\rho = 0.0025$, $M = 3000$


Figure SM12. Long-time observable hierarchy for the representative Branch A trajectory $\rho = 0.0025$, $M = 3000$, restricted to the boundary-safe part of the calculation. **(a)** Aggregate capacity error $|a|$. **(b)** Within-agent uncertainty β . **(c)** Nested bounded-window masses S_{10} , S_{50} , and S_{100} . The smaller interval responds first because less outward propagation is required for the density to leave it. **(d)** Analytical proof diagnostic $\phi = \beta g_2^3$, appearing in Lemma 5.4 and Corollary 5.5 of [1]. The panels use the scaled time $\theta = (M - 1)t$ and illustrate the ordering of the observables; no post-boundary crossing time is assigned to the wider windows.

Table SM14

Physical-time scaling of the direct Branch A S_{10} crossings. A slope of -1 corresponds to an approximately population-independent crossing time on the scaled variable $\theta = (M - 1)t$.

Threshold	ρ	n	Slope of $\log t$ on $\log(M - 1)$	R^2
0.10	0.0025	4	-0.992274	0.999996
0.10	0.02	4	-0.992339	0.999996
0.05	0.0025	3	-0.997719	1.000000
0.05	0.02	3	-0.997781	1.000000

Table SM15

Spatial-resolution audit for the direct Branch A S_{10} events at $\rho = 0.0025$. The mean crossing time averages the $M = 3000$ and $M = 30000$ values. The population spread is their absolute difference divided by their mean.

Threshold	Δx	Mean $\Theta_{10,q}$	M -spread	Maximum boundary mass
0.10	1	2.3221×10^5	0.0301%	1.017×10^{-3}
0.10	2	2.6388×10^5	0.0298%	1.018×10^{-3}
0.10	4	2.0394×10^5	0.0301%	1.020×10^{-3}
0.05	1	9.4199×10^5	0.0293%	1.017×10^{-3}
0.05	2	1.0709×10^6	0.0296%	1.018×10^{-3}
0.05	4	8.2905×10^5	0.0295%	1.020×10^{-3}

SM6.1. Late-time mechanism and analytical diagnostic. The scaled PDE (7.12) contains the diffusion contributions $(M - 1)a^2$ and b . Their ratio

$$(SM6.1) \quad \mathcal{R}_{\text{diff}} = \frac{(M - 1)a^2}{b}$$

shows which contribution is active at a direct S_{10} crossing. At $S_{10} = 0.10$, its median and maximum are 2.65×10^{-11} and 8.99×10^{-11} ; at $S_{10} = 0.05$, they are 9.02×10^{-12} and 2.67×10^{-11} . The corresponding maxima of the transport coefficient $|a|/\sqrt{\tau}$ are 4.40×10^{-4} and 1.19×10^{-4} . Thus the observed late redistribution is governed by the surviving b -driven diffusion, not by the residual capacity error.

For comparison with the analytical argument, we also record

$$(SM6.2) \quad g_2(\theta) = \int_{\mathbb{R}} (p(x)f(x, \theta))^2 dx, \quad \phi(\theta) = \beta(\theta)g_2(\theta)^3.$$

Under the hypotheses of Lemma 5.4 and Corollary 5.5 of [1],

$$(SM6.3) \quad S_R(\theta) \leq C_R \phi(\theta)^{1/7}.$$

In Branch A, ϕ decreases while S_{10} passes through both direct milestones (Figure SM12). The reported crossing times are nevertheless measured from S_{10} itself, not inferred from (SM6.3).

For Branch B, $p_B(x) > 0$ at every finite x but $\inf_{\mathbb{R}} p_B = 0$. No comparable boundary-safe S_{10} crossing is resolved on the available horizon. Its direct sorting time is therefore unresolved; this is not evidence that Branch B fails to sort.

REFERENCES

- [1] E. BOU DAGHER, M. PEREPELTSIA, AND E. ZATORSKA, *On a PDE model for learning in stochastic market entry games*, 2026, <https://doi.org/10.48550/arXiv.2603.06514>, <https://arxiv.org/abs/2603.06514>, <https://arxiv.org/abs/2603.06514>.
- [2] P. K. SWEBY, *High resolution schemes using flux limiters for hyperbolic conservation laws*, SIAM Journal on Numerical Analysis, 21 (1984), pp. 995–1011, <https://doi.org/10.1137/0721062>.
- [3] B. VAN LEER, *Towards the ultimate conservative difference scheme. V. a second-order sequel to Godunov's method*, Journal of Computational Physics, 32 (1979), pp. 101–136, [https://doi.org/10.1016/0021-9991\(79\)90145-1](https://doi.org/10.1016/0021-9991(79)90145-1).
- [4] P. VIRTANEN ET AL., *SciPy 1.0: Fundamental algorithms for scientific computing in python*, Nature Methods, 17 (2020), pp. 261–272, <https://doi.org/10.1038/s41592-019-0686-2>.