Corrigendum: Equilibrium Technology Diffusion, Trade, and Growth

Jesse Perla

University of British Columbia

Christopher Tonetti

Stanford Graduate School of Business and NBER

Michael E. Waugh

Federal Reserve Bank of Minneapolis and NBER

August 2026

ABSTRACT

This note corrects the solution of the model with geometric Brownian motion (GBM) productivity shocks in Perla, Tonetti, and Waugh (2021). Two boundary conditions were missing: smooth pasting of the value function at the export threshold and the absorbing-barrier condition that the stationary productivity density vanish at the adoption threshold. The corrected value function adds the option value of becoming an exporter to the value of non-exporting firms. The corrected stationary distribution is a difference of two power laws with a boundary layer at the adoption threshold. Both corrections vanish without GBM, so all theoretical results in the published paper are unaffected. We also update the numerical methods used to solve the corrected model. The welfare gains from trade in the corrected model are smaller than in the published paper. The paper's central mechanism and conclusion remain: trade affects incentives to adopt technology, generating growth effects and welfare gains about eight times the static benchmark.

We thank Jake Bradley for identifying the missing smooth-pasting condition at the export threshold and Danial Lashkari for alerting us to the absorbing-barrier condition and the corrected shape of the stationary distribution. Corrected code and replication materials are available at `https://github.com/HighDimensionalEconLab/perla_tonetti_waugh`.

1. Introduction

This note corrects errors in Perla, Tonetti, and Waugh (2021), "Equilibrium Technology Diffusion, Trade, and Growth," *American Economic Review*, 111(1), 73–128—henceforth PTW. The quantitative analysis in the published paper uses a version of the model in which firm productivity evolves according to geometric Brownian motion (GBM). On a balanced growth path (BGP), the normalized continuation value $v(z)$ and the normalized productivity distribution $F(z)$ satisfy

$$(r-g)\,v(z) = \pi(z) + \left(\mu + \tfrac{\upsilon^2}{2} - g\right) z\,v'(z) + \tfrac{\upsilon^2}{2}\,z^2\,v''(z), \tag{1}$$

$$0 = S\,F(z) - S + \left(g - \mu + \tfrac{\upsilon^2}{2}\right) zF'(z) + \tfrac{\upsilon^2}{2}\,z^2F''(z), \tag{2}$$

where $z := Z/M(t) \geq 1$ is productivity relative to the adoption threshold, flow profits are $\pi(z) = \bar{\pi}_{\min} z^{\sigma-1}$ for non-exporters ($z < \hat{z}$) and $\pi(z) = \bar{\pi}_{\min}(1 + (N-1)d^{1-\sigma})z^{\sigma-1} - (N-1)\kappa$ for exporters ($z \geq \hat{z}$), the export threshold is $\hat{z} = d(\kappa/\bar{\pi}_{\min})^{1/(\sigma-1)}$, and S is the flow of adopters.

When $\upsilon > 0$, equations (1) and (2) are second-order differential equations, and the solutions in the published paper omit one boundary condition each:

1. The value function must be continuously differentiable at the export threshold $\hat{z}$ (*smooth pasting*). The published value function has a kink at $\hat{z}$; the corrected value function adds the option value of becoming an exporter to the value of non-exporting firms (Section 2).

2. The density must vanish at the adoption threshold, $f(1) = 0$, because with a diffusion the adoption threshold is an *absorbing barrier*. The published truncated Pareto distribution violates this condition; the corrected stationary distribution is a difference of two power laws (Section 3).

Both corrections disappear as $\upsilon \to 0$. With deterministic dynamics, (1) and (2) are first-order equations for which the published boundary conditions are complete and the published solutions are correct. The corrected value function, distribution function, and all admissible moments converge to the published deterministic solutions. The density converges pointwise everywhere except at the threshold itself, where $f(1) = 0$ for every $\upsilon > 0$ while the limiting Pareto has $f(1) = \theta$; this boundary layer shrinks with υ. All theoretical results in PTW were derived for the $\mu = \upsilon = \delta = 0$ case and are therefore unaffected.

Quantitatively, the corrected model preserves the paper's central message with smaller magnitudes. The welfare gain from the 10 percent trade-cost reduction is 6.9 percent both across steady states and including transition dynamics, compared with 11.2 and 10.8 percent in the published paper. The gain from trade relative to autarky falls from 28.2 to 16.8 percent. The welfare gains nevertheless remain about eight times the static benchmark, and the growth channel continues to account for almost all of them.

2. Correction 1: Smooth Pasting at the Export Threshold

Substituting z^x into the homogeneous part of (1) yields the characteristic equation $\frac{\upsilon^2}{2}x^2 + (\mu - g)x - (r - g) = 0$, with one negative and one positive root, $-\nu < 0 < \beta$:

$$\nu = \frac{\mu - g}{\upsilon^2} + \sqrt{\left(\frac{g-\mu}{\upsilon^2}\right)^2 + \frac{r-g}{\upsilon^2/2}}, \qquad \beta = \nu + \frac{2(g-\mu)}{\upsilon^2}, \tag{3}$$

and, as in the published appendix, $a = \left[r - g - (\sigma - 1)\left(\mu - g + (\sigma-1)\upsilon^2/2\right)\right]^{-1} > 0$, which implies $\beta > \sigma - 1$. The published solution keeps only the decaying root $z^{-\nu}$ on both sides of $\hat{z}$:

$$v_d^0(z) = a\bar{\pi}_{\min}\left(z^{\sigma-1} + \tfrac{\sigma-1}{\nu}z^{-\nu}\right), \qquad 1 \le z \le \hat{z}, \tag{4}$$

$$v_x^0(z) = a\bar{\pi}_{\min}\left(\left(1 + (N-1)d^{1-\sigma}\right)z^{\sigma-1} + \tfrac{\sigma-1}{\nu}z^{-\nu}\right) + \tfrac{(N-1)\bar{\pi}_{\min}}{r-g}\left(b\,z^{-\nu} - \tfrac{\kappa}{\bar{\pi}_{\min}}\right), \qquad z \ge \hat{z}, \tag{5}$$

with $b = (1 - a(r-g))d^{1-\sigma}\hat{z}^{\,\nu+\sigma-1}$ chosen so that v^0 is continuous at $\hat{z}$. This solution solves (1) on each region and satisfies smooth pasting at the adoption threshold, $v'(1) = 0$, but its derivative jumps at the export threshold (Appendix A):

$$v_x^{0\prime}(\hat{z}) - v_d^{0\prime}(\hat{z}) = \frac{(N-1)\kappa}{\hat{z}} \cdot \frac{a\,\nu\,(\sigma-1)(\nu+\sigma-1)\,\upsilon^2/2}{r-g} \; > \; 0 \quad \text{iff } \upsilon > 0. \tag{6}$$

Because exporting is a static decision with a flow fixed cost, profits are continuous at $\hat{z}$, so for $\upsilon > 0$ the value function is continuously differentiable there; a kinked candidate cannot be the firm's value. Excluding z^β is justified by transversality only where z is unbounded: on the bounded region $[1, \hat{z}]$ it belongs to the solution, with its coefficient pinned by the missing smooth pasting condition at $\hat{z}$.

Corrected value function. Imposing $v'(1) = 0$ and both continuity and smooth pasting at $\hat{z}$ yields (Appendix A):

$$v_d(z) = v_d^0(z) + c_2\left(z^\beta + \tfrac{\beta}{\nu}z^{-\nu}\right), \qquad 1 \le z \le \hat{z}, \tag{7}$$

$$v_x(z) = v_x^0(z) + c_2\left(\tfrac{\beta}{\nu} + \hat{z}^{\,\nu+\beta}\right)z^{-\nu}, \qquad z \ge \hat{z}, \tag{8}$$

$$c_2 = \frac{(N-1)\kappa\,a\,\nu\,(\sigma-1)(\nu+\sigma-1)\,\upsilon^2/2}{(r-g)\,(\nu+\beta)}\,\hat{z}^{-\beta} \; > \; 0. \tag{9}$$

The new term $c_2 z^\beta$ is the *option value of becoming an exporter*: with $\upsilon > 0$, positive shocks can carry a non-exporting firm's relative productivity above $\hat{z}$, and this option—increasing in z (and convex at the calibrated parameters, where $\beta > 1$), largest just below the threshold—was omitted from the published value of non-exporting firms. (Deterministically, z only drifts down, a non-exporter can never become an exporter, and the option is worthless: as $\upsilon \to 0$, $\beta \to \infty$ and $c_2 \to 0$, recovering (4)–(5) exactly.) The accompanying $z^{-\nu}$ terms preserve smooth pasting at the adoption threshold and continuity at $\hat{z}$. The value of the marginal adopter, which enters the value matching and free entry conditions, becomes

$$v(1) = a\bar{\pi}_{\min}\left(1 + \frac{\sigma-1}{\nu}\right) + c_2\,\frac{\nu+\beta}{\nu}. \tag{10}$$

Since $c_2 \propto \upsilon^2 \hat{z}^{-\beta}$, this correction is quantitatively negligible at the calibrated parameters ($c_2 \approx 2.2 \times 10^{-10}$). At the calibrated point it becomes larger at higher volatility or a lower export threshold, though we make no global comparative-statics claim because ν, β, a, and $\hat{z}$ all move in equilibrium. We carry the exact formulas throughout.

We thank Jake Bradley for identifying the missing smooth-pasting condition at the export threshold. You can see his comment at `https://www.aeaweb.org/content/file?id=22306`.

3. Correction 2: The Absorbing Barrier and the Stationary Distribution

The published solution to (2) imposes $F(1) = 0$ and $F(\infty) = 1$ and selects the truncated Pareto distribution, $F(z) = 1 - z^{-\theta}$, with adopter flow $S = \theta(g - \mu - \theta\upsilon^2/2)$. For $\upsilon > 0$, however, (2) is second order and a third condition applies. Firms adopt exactly when z hits the adoption threshold, so the threshold is an absorbing barrier of a non-degenerate diffusion, and the stationary density of an absorbed-and-reinjected diffusion must vanish there:

$$f(1) = 0. \tag{11}$$

The truncated Pareto solves (2) pointwise but is not the distribution of the underlying stochastic process (Appendix B). The intuition is the contrast between how firms reach the barrier. With $\upsilon = 0$, firms slide toward the threshold at rate $g - \mu$; a firm at $z = 1 + \varepsilon$ survives for time $\log(1 + \varepsilon)/(g - \mu) \approx \varepsilon/(g - \mu)$, the barrier collects whatever sits on it, and $f(1) = \theta > 0$ is admissible, with $S = (g - \mu)f(1)$. With $\upsilon > 0$, diffusion dominates drift on small scales: within distance ε of the barrier the process moves on the Brownian time scale ε^2/υ^2, which vanishes faster than the drift scale $\varepsilon/(g - \mu)$, so mass cannot accumulate against an absorbing barrier—the stationary density must slope down to zero, and adoption is the diffusive flux $S = \frac{\upsilon^2}{2} f'(1) > 0$.

Corrected stationary distribution. Power solutions $z^{-\xi}$ of (2) satisfy $\frac{\upsilon^2}{2}\xi^2 - (g - \mu)\xi + S = 0$. Imposing $F(1) = 0$, $f(1) = 0$, and $F(\infty) = 1$, and labeling the tail root θ, the stationary distribution is a difference of two power laws:

$$F(z) = 1 - \frac{\xi_2\, z^{-\theta} - \theta\, z^{-\xi_2}}{\xi_2 - \theta}, \qquad f(z) = \frac{\theta\, \xi_2}{\xi_2 - \theta}\left(z^{-\theta-1} - z^{-\xi_2-1}\right), \qquad \xi_2 = \frac{2(g - \mu)}{\upsilon^2} - \theta. \tag{12}$$

The flow of adopters is unchanged, $S = \frac{\upsilon^2}{2}\theta\xi_2 = \theta\left(g - \mu - \theta\frac{\upsilon^2}{2}\right)$, and is now positive by construction on the admissible parameter set. Figure 1 plots the correction at the calibrated parameters: the two-power-law density inserts a boundary layer in which the density rises from zero and is otherwise close to the Pareto. Every moment is rescaled: for $p < \theta$,

$$\mathbb{E}\left[z^p\right] = \underbrace{\frac{\theta}{\theta - p}}_{\text{Pareto}} \times \underbrace{\frac{\xi_2}{\xi_2 - p}}_{\to 1 \text{ as } \upsilon \to 0}, \qquad 1 - F(\hat{z}) = \frac{\xi_2\, \hat{z}^{-\theta} - \theta\, \hat{z}^{-\xi_2}}{\xi_2 - \theta}, \tag{13}$$

and correspondingly for the exporter aggregates, the home trade share, adoption labor, and the value matching integral (Appendix C). As $\upsilon \to 0$, $\xi_2 \to \infty$ and (12) converges to the truncated Pareto.

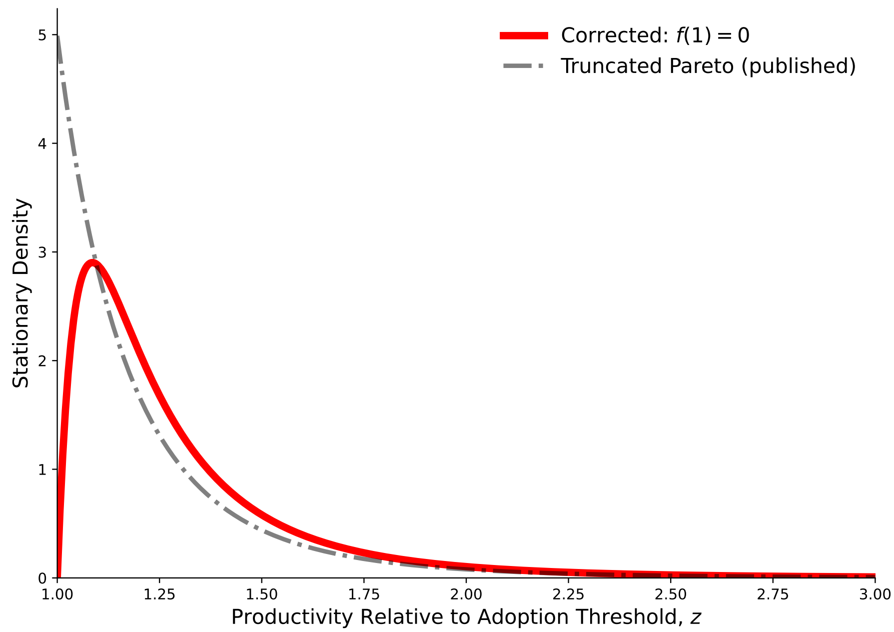

Figure 1: The Corrected Stationary Productivity Distribution

Existence and the initial condition. The inherited-tail solution (12) requires $\theta < \xi_2$—the condition that θ is the *smaller* of the two characteristic roots, so that the inherited exponent θ, rather than ξ_2, is the tail index of the stationary distribution—that is,

$$\upsilon^2 \; < \; \frac{g - \mu}{\theta}. \tag{14}$$

The interpretation is that $\bar{\xi} := (g - \mu)/\upsilon^2$ is the thinnest tail the barrier–GBM mechanism can sustain. Multiplicative shocks and reinjection proportional to the current distribution cannot thin a tail, so an initial distribution with tail $\theta < \bar{\xi}$ passes its tail to the stationary distribution. The correction changes only the shape near the barrier. An initial distribution thinner than the floor, including bounded support, instead converges to the critical solution with tail exactly $\bar{\xi}$ and a logarithmic correction, $f(z) = \bar{\xi}^2 \log(z)\, z^{-\bar{\xi}-1}$, as in Luttmer (2012) and Benhabib, Perla, and Tonetti (2021); see Appendix B. As in the published paper, we maintain the fat-tailed initial condition, with θ fixed at its published estimate. Condition (14) is comfortably satisfied at the calibrated parameters below: $\upsilon^2 = 0.0032$ and $(g - \mu)/\theta = 0.0081$.[1]

4. Recalibration

We re-estimate the corrected model with the same empirical targets, weights, and data as the published paper. The corrected model introduces an identification issue for θ and σ. With the published moments, the two parameters are at best weakly identified because both help determine the firm-size distribution. Estimating both freely would require additional moments that separately discipline the productivity tail and the elasticity of substitution. Rather than add new moments in a corrigendum, we hold θ and

[1] The published calibration table reports 0.048 as the *variance* υ^2; the value used in computation is the standard deviation, $\upsilon = 0.048$ ($\upsilon^2 \approx 0.0023$)—an erratum in the table label. We report standard deviations throughout.

σ at their published estimates and re-estimate $(d, \kappa, \chi, \mu, \upsilon)$ using the published SMM procedure. This isolates the correction from a change in the two published elasticities.

We compute the static firm moments analytically and the transition moments with a mass-conservative discrete operator. Appendix D describes the numerical methods. The corrected model matches the interest rate, productivity growth, and import share exactly. It implies an exporter share of 3.4 percent, compared with 3.3 percent in the data, and a relative exporter size of 5.0, compared with 4.8. The correlation between the data and corrected five-year transition matrix is 0.94, compared with 0.98 for the published model.

Table 1: Calibration: Parameters and Values, Published and Corrected

Parameter	Published	Corrected
Technology Adoption Cost, ζ	1.0	1.0
Number of Countries, N	10	10
Discount Rate, ρ	0.0203	0.0203
Pareto Shape Parameter, θ	4.99	4.99
Variety Elasticity of Substitution, σ	3.17	3.17
Drift of GBM Process μ	-0.031	-0.032
Std. Deviation of GBM Process υ	0.048	0.056
Death Rate of Firms δ	0.020	0.020
Iceberg Trade Cost, d	3.02	3.05
Export Fixed Cost, κ	0.104	0.090
Entry Cost Relative to Adoption Cost $1/\chi$	7.88	6.52

Note: The Published column reproduces the published calibration table. In the Corrected column, θ and σ are held at their published AER estimates (displayed rounded), and $(d, \kappa, 1/\chi, \mu, \upsilon)$ are estimated by SMM on the same targets and weights as the published paper. The GBM parameter is reported as a standard deviation in both columns, correcting the variance label in the published table.

Table 2: Aggregate Moments: Data, Published Model, and Corrected Model

Moment Description	Data	Published	Corrected
U.S. Real Interest Rate	2.83	2.83	2.83
U.S. Productivity Growth	0.79	0.79	0.79
U.S. Import/GDP	10.63	10.63	10.63
Share of Exporting Establishments	3.3	3.3	3.4
Relative Size of Exporting Establishments	4.8	4.8	5.0

Moment construction is as in the published paper. The real interest rate, productivity growth rate, and import/GDP ratio are in percent and averages over the 1977–2000 time period. The Published column reproduces the published paper's model column.

Table 3: Establishment Size Dynamics, Data, Corrected and Published Models

Transition Matrix of Relative Size: largest, quartile 1; smallest, quartile 4

Data

	1	2	3	4
1	50.1	27.0	13.6	9.4
2	16.1	28.5	29.0	26.0

Corrected Model

	1	2	3	4
1	44.0	29.9	16.5	9.6
2	12.6	23.9	31.1	32.3

Corr(Data, Corrected Model) = 0.94

Published Model

	1	2	3	4
1	45.7	31.5	14.2	8.6
2	17.1	28.8	28.2	25.9

Corr(Data, Published Model) = 0.98

Employment Share of New Establishments

Data: 0.02 Corrected Model: 0.02

Note: In each panel, rows represent the establishment-size quartile in period t; columns represent the establishment-size quartile in period $t+5$. Data source and construction as in the published paper; the Published Model panel reproduces the published paper's model panel.

5. Results

Table 4 revisits the central counterfactual of the published paper: a 10 percent reduction in the net-of-one iceberg cost, $d_T - 1 = 0.9(d_0 - 1)$. We refer to this as a 10 percent reduction in trade costs and present the results in the format of the published paper's Table 5.

Table 4: 10 Percent Reduction in Trade Costs: Growth, Trade, and Welfare

	Published		Corrected	
	Baseline BGP	New BGP	Baseline BGP	New BGP
Growth	0.79	1.03	0.79	0.94
Imports/GDP	10.6	14.4	10.6	14.3
Welfare				
Transition Path:	10.8		6.9	
SS to SS:	11.2		6.9	

Note: All values are in percent. Consumption-equivalent is the permanent percent increase in consumption a household requires in the old regime to be indifferent between the new and old regimes. The Published columns reproduce the published paper's Table 5 values at the published calibration, computed under the published package's welfare aggregation, whose terminal continuation substitutes $T\,g(T)$ for the accumulated $\log M(T)$; the same published path under the corrected aggregation gives a transition value of 10.2 rather than 10.8 percent (details in the numerical appendix). The Corrected columns use the corrected model, calibration, and welfare aggregation, so part of the difference between the transition entries reflects the aggregation correction rather than the corrected economics.

The corrected model preserves the paper's central message with smaller magnitudes. A 10 percent reduction in trade costs raises the growth rate by 0.15 percentage points, compared with 0.24 in the published model, and yields a steady-state consumption-equivalent welfare gain of 6.9 percent, compared with 11.2. The ACR benchmark at the corrected calibration is 0.85 percent, so the corrected gain is about 8.1 times ACR. Relative to autarky, the corrected gains from trade are 16.8 percent.[2]

The revision operates through the corrected stationary distribution and the re-estimated parameters. The iceberg cost rises from 3.02 to 3.05, the export fixed cost falls from 0.104 to 0.090, the relative entry cost $1/\chi$ falls from 7.88 to 6.52, the GBM drift moves from -0.031 to -0.032, and its standard deviation rises from 0.048 to 0.056. Together, these changes dampen the response of growth to trade costs. The value-function correction in Section 2, while necessary for internal consistency, remains quantitatively negligible at the calibrated parameters.

[2] As in the published paper, autarky is computed by raising the net-of-one trade cost by a factor of 2.9. This lowers the import share to 0.2 percent. Moving from the baseline to autarky lowers welfare by 14.4 percent, about 2.1 times the gain from the 10 percent trade-cost reduction, while ACR predicts a 2.2 percent loss. The 16.8 percent gain from autarky and the 14.4 percent loss moving to autarky differ because consumption-equivalent changes are multiplicative.

5.1. The Sources of the Welfare Gains from Trade

Section 7.3 of the published paper decomposes the welfare change induced by a small change in trade costs into economically meaningful components. Describing a feasible steady-state allocation by $(\Omega, \hat{z}, g)$ given d, with f_c the consumption level it delivers and $f_\Omega(d)$, $f_{\hat{z}}(d)$, $f_g(d)$ the equilibrium mappings, the total derivative of steady-state utility $\bar{U}(c, g)$ is (equation (60) of the published paper)

$$\frac{\mathrm{d}\bar{U}}{\mathrm{d}d} = \underbrace{\bar{U}_1 \frac{\partial f_c}{\partial d}}_{\text{direct consumption}} + \underbrace{\bar{U}_1 \frac{\partial f_c}{\partial \Omega} \frac{\mathrm{d}f_\Omega}{\mathrm{d}d}}_{\text{variety}} + \underbrace{\bar{U}_1 \frac{\partial f_c}{\partial \hat{z}} \frac{\mathrm{d}f_{\hat{z}}}{\mathrm{d}d}}_{\text{export threshold}} + \underbrace{\bar{U}_1 \frac{\partial f_c}{\partial g} \frac{\mathrm{d}f_g}{\mathrm{d}d}}_{\text{growth, via consumption}} + \underbrace{\bar{U}_2 \frac{\mathrm{d}f_g}{\mathrm{d}d}}_{\text{direct growth}}. \tag{15}$$

In the corrected model the feasibility map f_c depends on g through the stationary distribution as well as through adoption labor, since given g the distribution (12) is pinned by the Kolmogorov forward equation. The decomposition is otherwise unchanged. Evaluating each component at the corrected calibration and expressing it as a percentage of the total change gives

$$\underbrace{\frac{\mathrm{d}\bar{U}(c, g)}{\mathrm{d}d}}_{\textbf{100\%}} = \underbrace{\bar{U}_1 \frac{\partial f_c}{\partial d}}_{\textbf{13.27\%}} + \underbrace{\bar{U}_1 \frac{\partial f_c}{\partial \Omega} \frac{\mathrm{d}f_\Omega}{\mathrm{d}d}}_{\textbf{-9.20\%}} + \underbrace{\bar{U}_1 \frac{\partial f_c}{\partial \hat{z}} \frac{\mathrm{d}f_{\hat{z}}}{\mathrm{d}d}}_{\textbf{0\%}} + \underbrace{\bar{U}_1 \frac{\partial f_c}{\partial g} \frac{\mathrm{d}f_g}{\mathrm{d}d}}_{\textbf{-15.42\%}} + \underbrace{\bar{U}_2 \frac{\mathrm{d}f_g}{\mathrm{d}d}}_{\textbf{111.35\%}}. \tag{16}$$

The composition of the gains is essentially unchanged by the corrections. The direct growth effect accounts for 111.4 percent of the total, compared with 108.9 percent in the published model. This term is nonzero because the decentralized growth rate is inefficiently low. The measure of that inefficiency is 42.4 in the corrected model, compared with 44.7 in the published model and a free-growth benchmark of $1/\rho = 49.2$. The export-threshold term remains exactly zero, consistent with the efficiency of the export margin. The direct consumption effect is 101.3 percent of the local ACR benchmark. What changes is the strength of the growth channel, not its role: the semi-elasticity of growth with respect to trade costs falls from -0.028 to -0.018.

5.2. Transition Dynamics

The published paper's preferred welfare numbers include the transition path following the unanticipated trade-cost reduction. In the published solver, the transition is tractable precisely because of the missing boundary condition: the truncated Pareto solves the Kolmogorov forward equation for *any* path $g(t)$ once $f(1) = 0$ is dropped, so the distribution could be held fixed and only the value function solved backward. In the corrected model the distribution is a state variable. The density $f(z, t)$ solves the forward KFE with the absorbing barrier, the adoption flow is its diffusive flux $S(t) = \frac{v^2}{2}\partial_z f(1, t)$, and the equilibrium is a coupled forward–backward system.[3]

Figures 2–6 present the transition in the format of the published Section 7.4, with the corrected model in solid lines and the published path in dashed lines. The corrected path retains the qualitative shape

[3] The solver combines an implicit finite-difference HJB for the rescaled value with a Scharfetter–Gummel discretization of the forward KFE and iterates jointly on growth and entry using Anderson acceleration. We report the regular completion at the impact instant. Appendix D documents the method.

of the published one; the values that follow update those reported in the body of that section. Imports jump on impact and then converge slowly to their new level. The measure of domestic varieties declines gradually to 95.7 percent of its initial level. At the first reported quarter, labor in entry is 1.66 percentage points below its initial level, compared with 2.2 at impact in the published path, while labor in export fixed costs and adoption is 0.69 and 0.21 percentage points above, compared with 0.7 and 0.4. Normalized consumption rises by 2.8 percent on impact and remains 2.7 percent above its initial level at the first reported quarter.

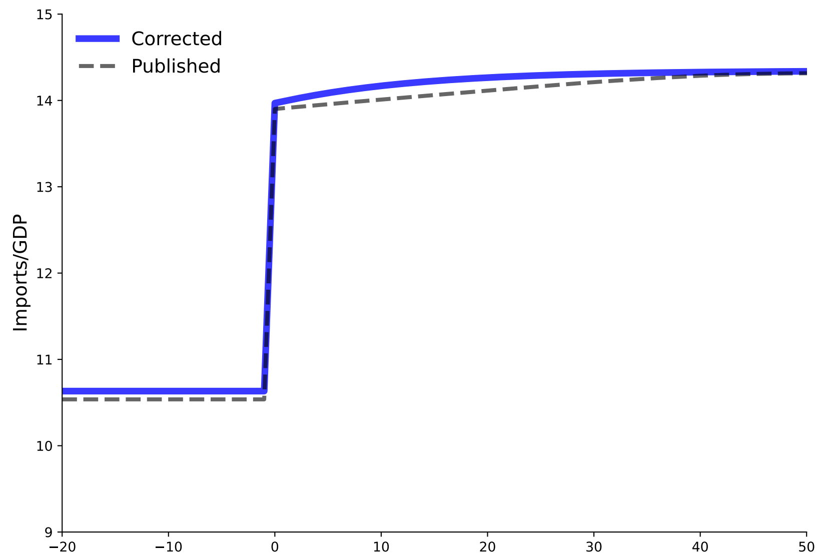

Figure 2: Imports/GDP: $1 - \lambda_{ii}(t)$

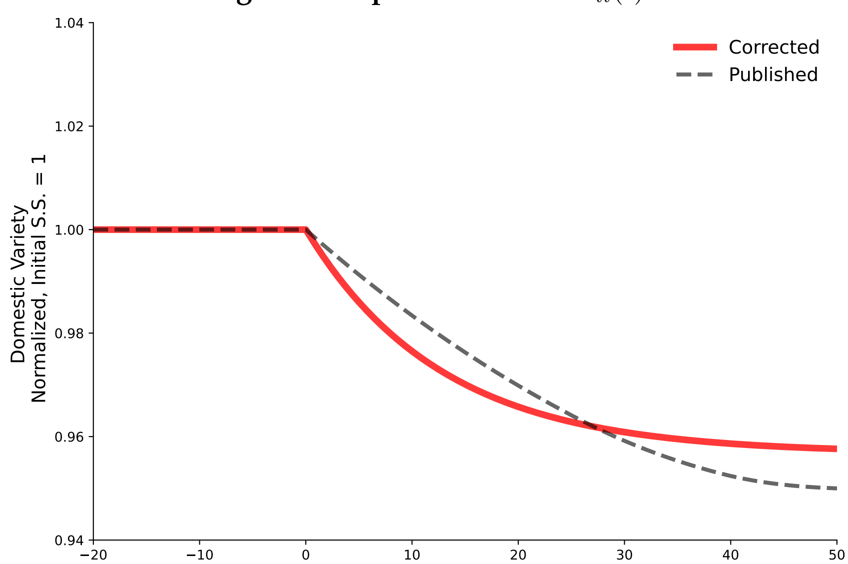

Figure 3: Domestic Variety: $\Omega(t)$

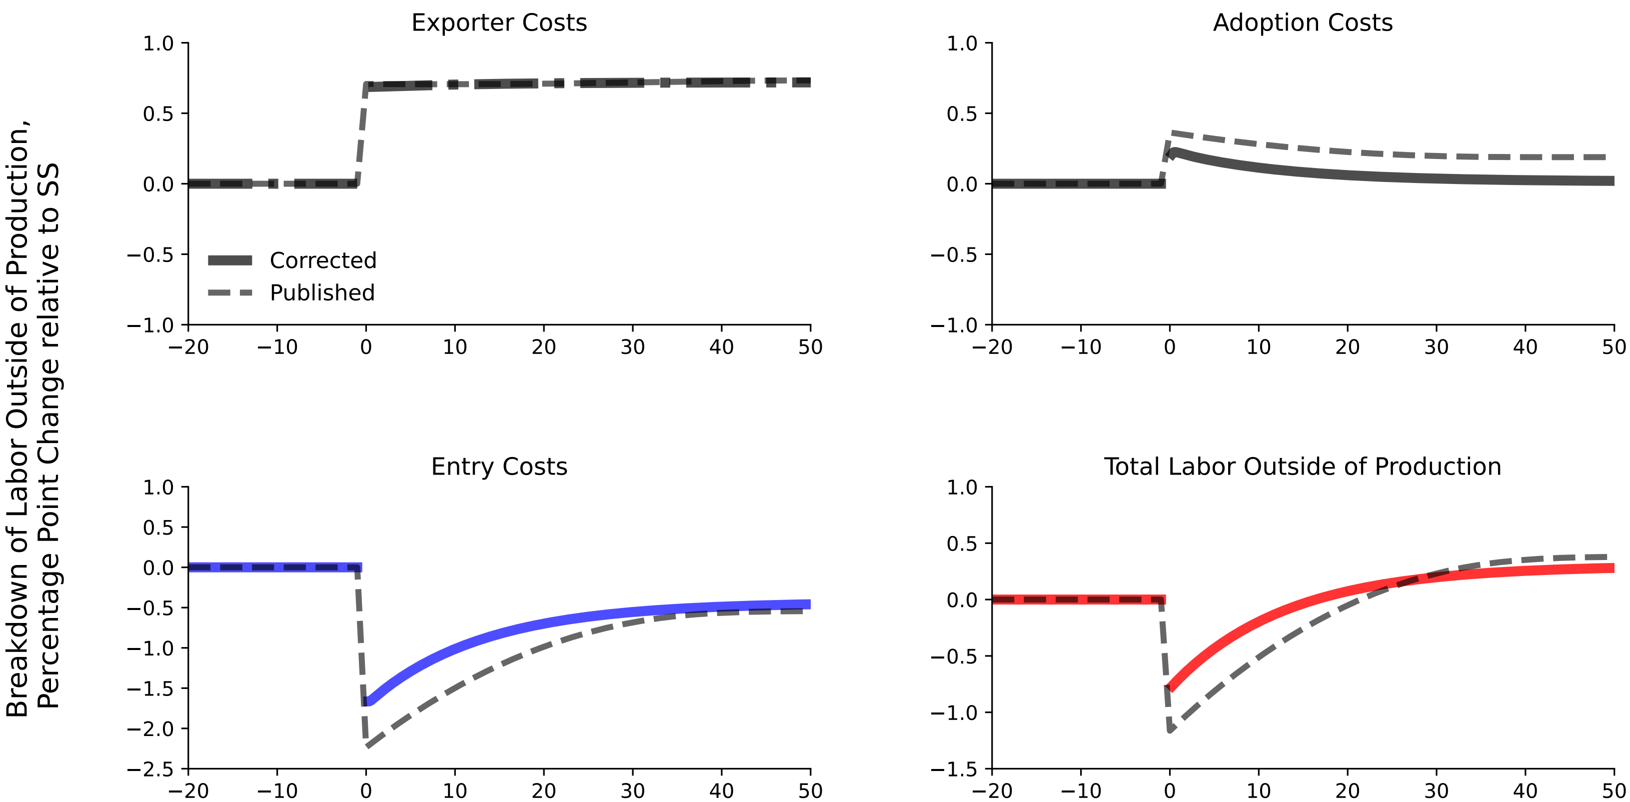

Figure 4: Reallocation Effects in Fixed Costs

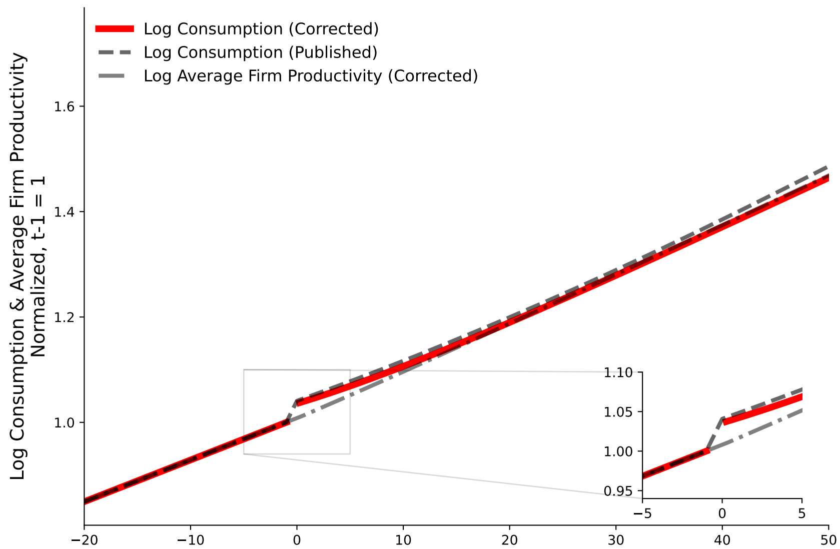

Figure 5: Log Consumption: $C(t)$

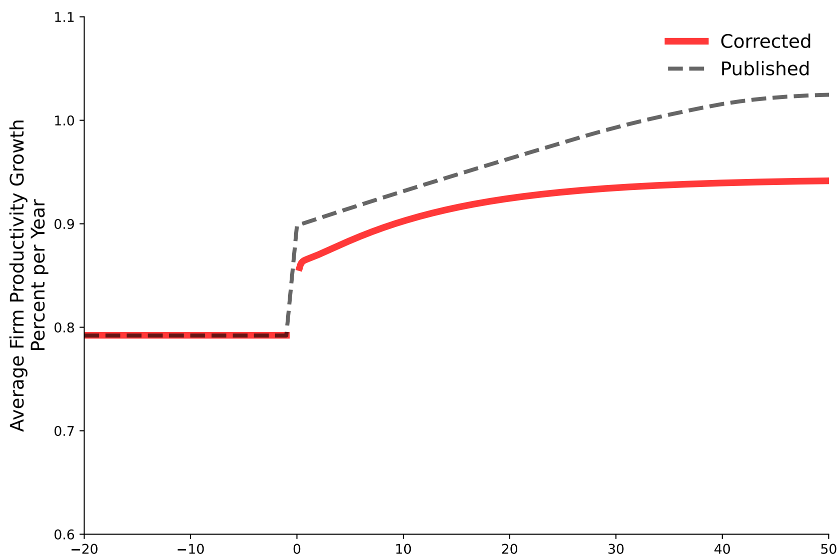

Figure 6: Average Firm Productivity Growth

Average firm productivity is $M(t)\mathbb{E}_t[z]$, where $\mathbb{E}_t[z]$ is mean productivity relative to the adoption threshold. In the published model, its growth rate equals $g(t)$ because the normalized Pareto distribution is fixed. In the corrected model, average firm-productivity growth rises over the reported transition because the adoption flow rises, and this increase more than offsets the small decline in the average productivity gain from each adoption.

The consumption-equivalent gain including the transition is 6.949 percent, while the gain across the production-grid stationary endpoints is 6.936 percent. Both report as 6.9 percent in Table 4. The published transition value is 10.8 percent. Reaggregating that same published path using its actual accumulated $\log M(t)$ gives 10.2 percent, a difference of 0.6 percentage points. This comparison holds the published path fixed; it is not a decomposition of the total correction.

6. The Role of Firm Dynamics and Adoption Costs

The published paper's Section 7.5 examines how firm dynamics and adoption costs shape the growth response to trade costs and the welfare gains from trade. Figures 7 and 8 repeat those experiments with the corrected model. Holding all other parameters at the corrected calibration, we scale μ and υ jointly or scale the exit shock δ. The published text describes the experiment as varying the variance alone, but the code that generated the published Figure 6 scales the drift and the volatility by a common factor, and we replicate that experiment. Each figure reports the calibrated $1/\chi = 6.52$ and values of $1/\chi$ that are 10 percent smaller and larger. Each point compares BGPs before and after a 10 percent reduction in trade costs, without recalibrating other parameters. All GBM points shown are admissible. At sufficiently low exit rates, the initial BGP has nonpositive growth, so the corresponding δ lines begin at the first admissible point.

Three results organize Figure 7. First, holding χ fixed, joint drift–volatility scaling has a modest effect on the response to trade costs. The change in growth ranges from 0.11–0.12, 0.15–0.16, and 0.19–0.21 percentage points for large, calibrated, and small χ. Second, adoption costs have a larger effect. Welfare ranges from 4.9–5.5 percent for large χ, 6.8–7.5 percent at the calibration, and 9.1–9.9 percent for small χ. Third, the level of initial growth is sensitive to the scale of firm-level shocks. Thus, as in the published paper, firm dynamics affect the gains from trade primarily through the adoption cost inferred from the calibration.

Figure 8 shows that, as in the published paper, raising the exit rate increases initial growth, the growth response to lower trade costs, and the welfare gain. At the calibrated χ, the change in growth ranges from about 0.08 to 0.39 percentage points over the admissible displayed range. The sensitivity is greatest near the low-exit boundary, where the initial BGP approaches zero growth.

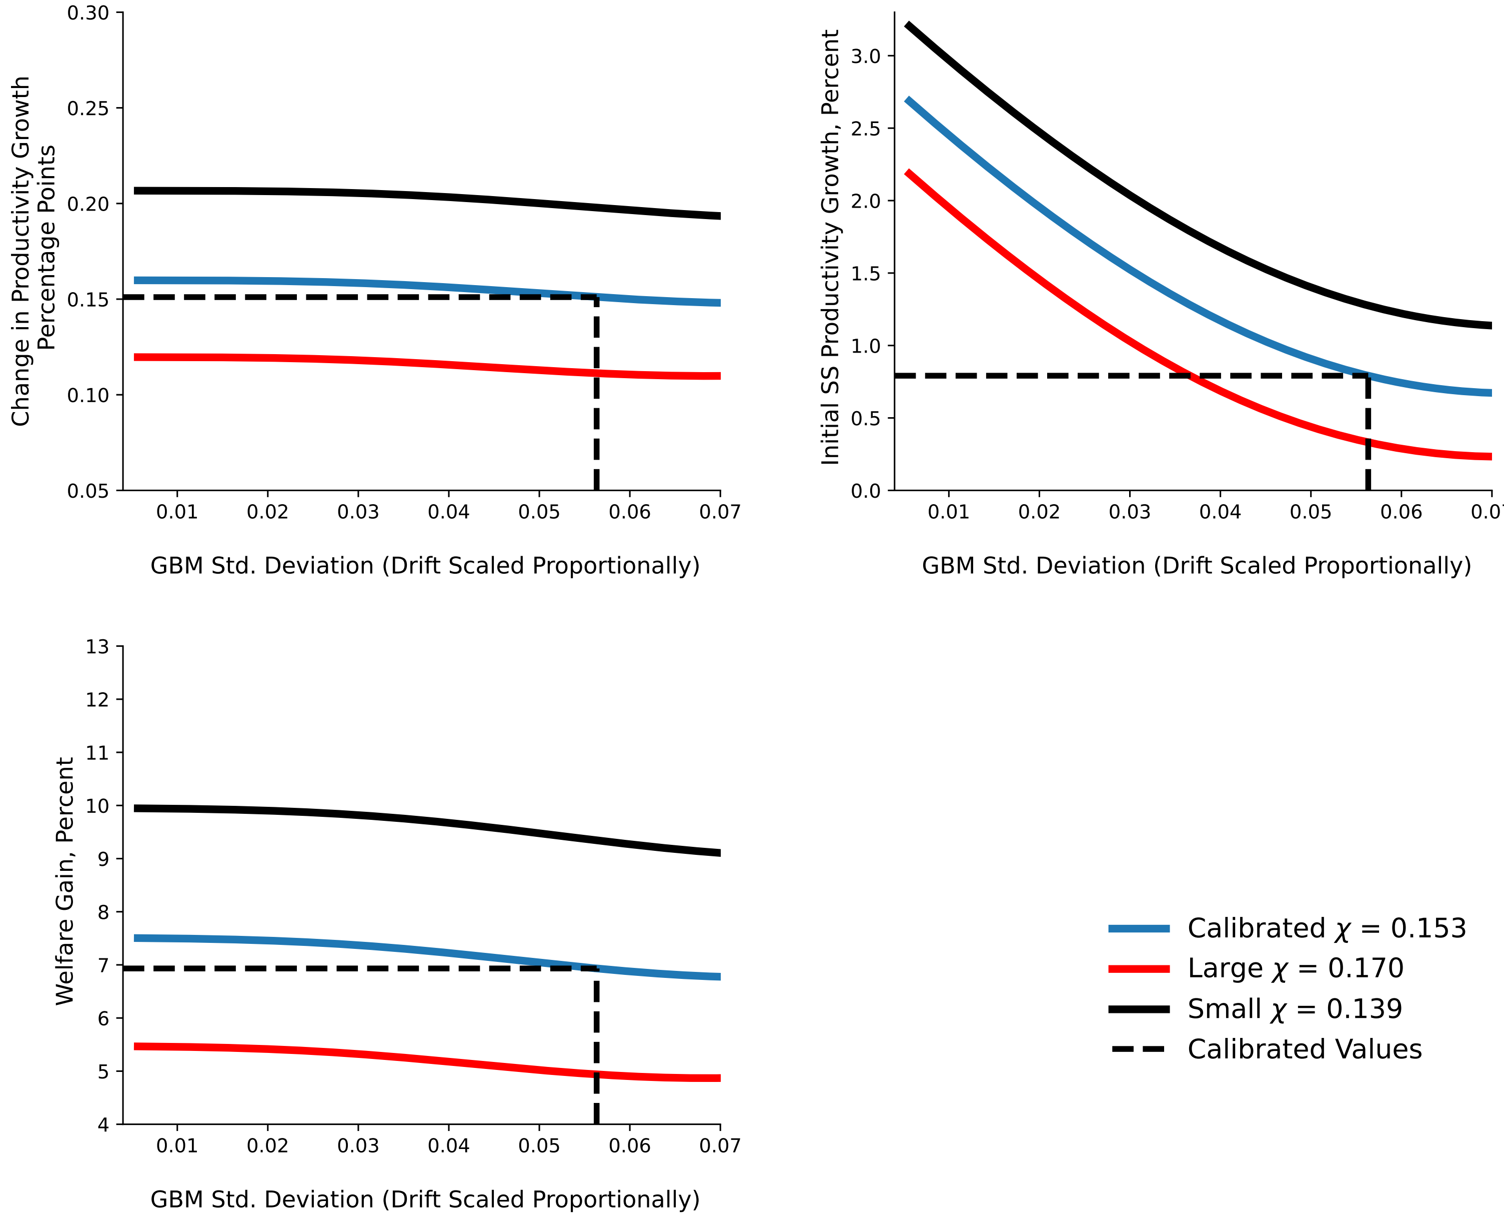

Figure 7: Joint GBM Drift–Volatility Scaling, Adoption Costs, Growth, and Welfare

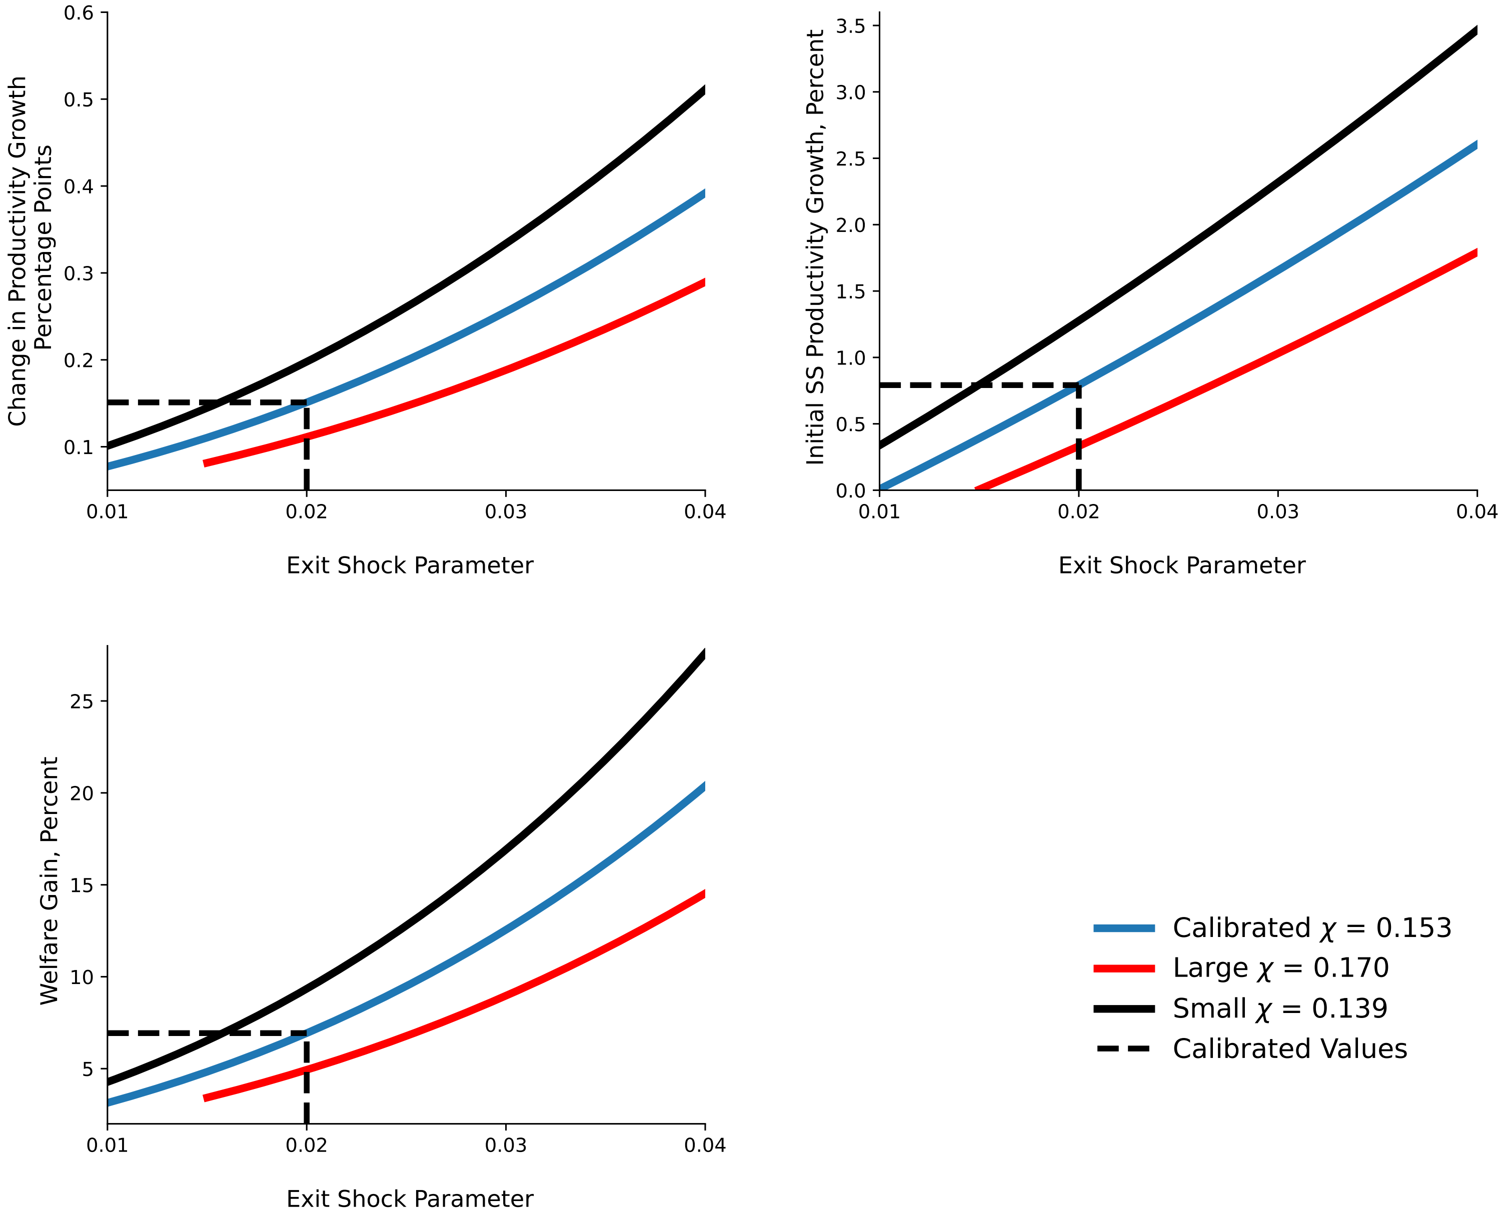

Figure 8: Exit Shocks, Adoption Costs, Growth, and Welfare

References

Achdou, Han, Lasry, Lions, and Moll (2022): "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *Review of Economic Studies*, 89(1), 45–86.

Arkolakis, Costinot, and Rodríguez-Clare (2012): "New Trade Models, Same Old Gains?" *American Economic Review*, 102(1), 94–130.

Benhabib, Perla, and Tonetti (2021): "Reconciling Models of Diffusion and Innovation: A Theory of the Productivity Distribution and Technology Frontier," *Econometrica*, 89(5), 2261–2301.

Luttmer (2007): "Selection, Growth, and the Size Distribution of Firms," *Quarterly Journal of Economics*, 122(3), 1103–1144.

Luttmer (2012): "Eventually, Noise and Imitation Implies Balanced Growth," Federal Reserve Bank of Minneapolis Working Paper 699.

Perla, Tonetti, and Waugh (2021): "Equilibrium Technology Diffusion, Trade, and Growth," *American Economic Review*, 111(1), 73–128.

Peskir and Shiryaev (2006): *Optimal Stopping and Free-Boundary Problems*, Birkhäuser.

Scharfetter and Gummel (1969): "Large-Signal Analysis of a Silicon Read Diode Oscillator," *IEEE Transactions on Electron Devices*, 16(1), 64–77.

Walker and Ni (2011): "Anderson Acceleration for Fixed-Point Iterations," *SIAM Journal on Numerical Analysis*, 49(4), 1715–1735.

A. Derivation of the Corrected Value Function

The general solution of (1) on each region is the sum of a particular solution and the two homogeneous solutions $z^{-\nu}$ and z^{β} from (3):

$$v_d(z) = a\bar{\pi}_{\min}\, z^{\sigma-1} + c_1 z^{-\nu} + c_2 z^{\beta}, \qquad 1 \le z \le \hat{z}, \tag{A.1}$$

$$v_x(z) = a\bar{\pi}_{\min}\left(1 + (N-1)d^{1-\sigma}\right) z^{\sigma-1} - \frac{(N-1)\kappa}{r-g} + c_3 z^{-\nu} + c_4 z^{\beta}, \qquad z \ge \hat{z}. \tag{A.2}$$

Since $\beta > \sigma - 1$, the no-bubble (transversality) condition at $z \to \infty$ requires $c_4 = 0$; it places no restriction on c_2, which multiplies a bounded function on $[1, \hat{z}]$. Three free constants remain, matched by three conditions: smooth pasting at the adoption threshold and continuity and smooth pasting at the export threshold,

$$v_d'(1) = 0, \qquad v_d(\hat{z}) = v_x(\hat{z}), \qquad v_d'(\hat{z}) = v_x'(\hat{z}). \tag{A.3}$$

Smooth pasting at $\hat{z}$ is not a separate high-contact optimality condition: $\hat{z}$ is a static export/no-export interface, not an optimal-stopping boundary. Instead, it is the standard C^1 pasting requirement for a second-order ODE with continuous, piecewise-smooth forcing. Indeed the solution is smoother still: since v, v', and π are continuous at $\hat{z}$, the HJB itself implies that v'' is also continuous there. Differentiating the equation once shows that the first discontinuity appears in v''', with

$$v'''(\hat{z}^{+}) - v'''(\hat{z}^{-}) = -\frac{\pi'(\hat{z}^{+}) - \pi'(\hat{z}^{-})}{(\upsilon^2/2)\hat{z}^2}.$$

Probabilistically, a kinked candidate value function generates a local-time term in the Itô–Meyer formula that the HJB does not account for, so it cannot equal the expected discounted profit stream; by a standard verification argument (e.g., Peskir and Shiryaev, 2006), the C^1 solution of (A.1)–(A.3) does, stopped at the adoption threshold.

The key algebraic identity is the following.

Lemma 1. *Let $-\nu$ solve the characteristic equation and let a be as in Section 2. Then*

$$a\,(\sigma - 1 + \nu) - \frac{\nu}{r-g} \;=\; \frac{a\,\nu\,(\sigma-1)\,(\nu+\sigma-1)\,\upsilon^2/2}{r-g}. \tag{A.4}$$

Proof. Write $s := \sigma - 1$ and $\psi(x) := \frac{\upsilon^2}{2}x^2 + (\mu - g)x - (r-g)$, so $1/a = -\psi(s)$ and $\psi(-\nu) = 0$. Multiplying the left side of (A.4) by $(r-g)/a$:

$$(s+\nu)(r-g) + \nu\,\psi(s) = s(r-g) + \nu\left(\tfrac{\upsilon^2}{2}s^2 + (\mu-g)s\right) = s\left[(r-g) + \nu(\mu-g) + \tfrac{\upsilon^2}{2}\nu s\right].$$

From $\psi(-\nu) = 0$, $(r-g) + \nu(\mu - g) = \frac{\upsilon^2}{2}\nu^2$, so the bracket equals $\frac{\upsilon^2}{2}\nu(\nu+s)$. □

Derivative jump of the published solution. Differentiating (4)–(5) at $\hat{z}$, using the static export-threshold condition $\kappa = \bar{\pi}_{\min} d^{1-\sigma} \hat{z}^{\sigma-1}$ and the published b, gives

$$v_x^{0\prime}(\hat{z}) - v_d^{0\prime}(\hat{z}) = \frac{(N-1)\kappa}{\hat{z}} \left[a(\sigma - 1 + \nu) - \frac{\nu}{r-g} \right],$$

which is equation (6) by Lemma 1: strictly positive for $\upsilon > 0$ and zero at $\upsilon = 0$. (At $\upsilon = 0$ equation (1) is first order, continuity of v at $\hat{z}$ is the only admissible pasting condition, and C^1 holds automatically because π is continuous; the published solution is exactly right in that case.)

Solving for the constants. Substituting (A.1)–(A.2) with $c_4 = 0$ into the two pasting conditions at $\hat{z}$, and defining $K := (N-1)\kappa$ and $Q := K(1 - a(r-g))/(r-g)$:

$$[C^0] \quad (c_3 - c_1)\,\hat{z}^{-\nu} - c_2\,\hat{z}^{\beta} = Q, \qquad [C^1] \quad \nu(c_3 - c_1)\,\hat{z}^{-\nu} + \beta c_2\,\hat{z}^{\beta} = (\sigma - 1)\,aK, \tag{A.5}$$

using $(N-1)\,a\bar{\pi}_{\min}\,d^{1-\sigma}\hat{z}^{\sigma-1} = aK$. Solving (A.5) and applying Lemma 1 to the numerator $(\sigma-1)aK - \nu Q = K\left[a(\sigma - 1 + \nu) - \frac{\nu}{r-g}\right]$ yields c_2 in (9) and

$$c_1 = \frac{(\sigma - 1)\,a\bar{\pi}_{\min} + \beta\,c_2}{\nu}, \qquad c_3 = c_1 + \left(Q + c_2\,\hat{z}^{\beta}\right)\hat{z}^{\nu}, \tag{A.6}$$

where c_1 imposes $v_d'(1) = (\sigma-1)a\bar{\pi}_{\min} - \nu c_1 + \beta c_2 = 0$. Writing (A.1) with (A.6) as the published solution plus c_2-proportional terms gives (7)–(8), and evaluating at $z = 1$ gives (10). All expressions were verified against an independent numerical solution of (1) as a boundary value problem with the kinked profit function ($v'(1) = 0$ and a far-field growth condition), with relative errors below 10^{-8}; the published solution's kink matches (6) to ten digits.

B. Derivation of the Corrected Stationary Distribution

Why $f(1) = 0$. In the model, a firm adopts exactly when its normalized productivity hits the threshold $z = 1$ (the optimal stopping rule), and it is instantaneously reinjected as a draw from the cross-sectional distribution f. For $\upsilon > 0$ the threshold is therefore an absorbing barrier of a non-degenerate diffusion. The transition density of a diffusion killed at an absorbing barrier vanishes linearly at the barrier in the forward variable: by the reflection principle, the Dirichlet heat kernel is the difference of two Gaussian kernels centered at mirror-image points, and that difference vanishes linearly in the distance to the barrier (correspondingly, the integrated probability of ending within ε of the barrier without having touched it is of order ε^2). The stationary density of the absorbed-and-reinjected process is a superposition of such kernels. Hence $f(1) = 0$, with the outflow through the barrier given by the diffusive flux $S = \frac{\upsilon^2}{2} f'(1)$. The published boundary conditions ($F(1) = 0$, $F(\infty) = 1$) are complete only for the first-order ($\upsilon = 0$) equation; for $\upsilon > 0$ they leave a one-parameter family of pointwise solutions, which is precisely the spurious freedom that allowed a truncated Pareto with any θ to appear to solve the system for any υ.

The published formula for the flow of adopters makes the problem visible when decomposed as

$$S \;=\; \underbrace{\theta\,(g - \mu)}_{\text{sweeping}} \;-\; \underbrace{\theta^2 \upsilon^2/2}_{\text{diffusive escape}},$$

where the second term is diffusion carrying mass from the dense region near the barrier up into the thin tail. If υ^2 were large enough, holding a thin-tailed Pareto stationary would require the barrier to *emit* firms ($S < 0$). Since adoption is one-way absorption, $S \geq 0$ is structural; with the corrected boundary condition, $S = \frac{\upsilon^2}{2} f'(1) > 0$ holds by construction.

The corrected solution. The stationary KFE (2) is a linear second-order equation of Euler (equidimensional) type: every term is of the form $z^k F^{(k)}(z)$ with constant coefficients, so in $x = \log z$ it becomes a constant-coefficient equation whose solutions are exponentials in x—that is, powers of z. Its general solution is therefore a particular solution plus a two-dimensional family of powers. The constant $F = 1$ solves the equation (the SF and $-S$ terms cancel and the derivative terms vanish), and substituting a candidate power $z^{-\xi}$ into the homogeneous part yields the characteristic quadratic $\frac{\upsilon^2}{2}\xi^2 - (g-\mu)\xi + S = 0$. Hence the general solution with $F(\infty) = 1$ is

$$F(z) = 1 + C_1 z^{-\xi_1} + C_2 z^{-\xi_2},$$

where ξ_1, ξ_2 are the two roots of the quadratic, which satisfy $\xi_1 + \xi_2 = 2(g-\mu)/\upsilon^2$ and $\xi_1\xi_2 = 2S/\upsilon^2$. The published solution used one boundary condition ($F(1) = 0$) to pin one of the two constants and set the other to zero; the absorbing barrier supplies the second condition. Imposing $F(1) = 0$ and $F'(1) = 0$ and labeling the tail (smaller) root θ gives $C_1 = -\xi_2/(\xi_2 - \theta)$, $C_2 = \theta/(\xi_2 - \theta)$, which is (12). The flux condition then holds automatically: $f'(1) = \frac{\theta\xi_2}{\xi_2 - \theta}[(\xi_2 + 1) - (\theta + 1)] = \theta\xi_2$, so $S = \frac{\upsilon^2}{2}\theta\xi_2 = \theta(g - \mu - \theta\upsilon^2/2)$, identical to the published expression.

Moments. For $p < \theta$ and $1 \leq \ell < u \leq \infty$, with $C := \frac{\theta\xi_2}{\xi_2 - \theta}$,

$$\int_\ell^u z^p f(z)\,\mathrm{d}z = C\left[\frac{\ell^{p-\theta} - u^{p-\theta}}{\theta - p} - \frac{\ell^{p-\xi_2} - u^{p-\xi_2}}{\xi_2 - p}\right], \tag{B.1}$$

where terms with $u = \infty$ vanish (and $p > \theta$ is admissible on bounded intervals, as needed for the z^β term of the value function). Equation (B.1) delivers all objects in the equilibrium system: $\mathbb{E}\left[z^{\sigma-1}\right]$ and $1 - F(\hat{z})$ in (13), the exporter aggregate $\int_{\hat{z}}^\infty z^{\sigma-1} f\,\mathrm{d}z$, and the value matching integral $\int_1^\infty v f\,\mathrm{d}z$ computed piecewise with (7)–(8).

Existence, the tail floor, and the critical solution. The two-root solution requires $\theta < \xi_2$—the root ordering that makes θ the tail—i.e., condition (14). Since the two roots sum to $2\bar{\xi}$ with $\bar{\xi} = (g-\mu)/\upsilon^2$ and the tail is the smaller root, the sustainable stationary tail exponents are exactly the interval $(0, \bar{\xi}]$: the interior $(0, \bar{\xi})$ for the strict two-root inherited solutions, and the endpoint attained by the critical solution below. At the boundary $\theta = \xi_2 = \bar{\xi}$ the two power laws merge and the stationary distribution acquires a logarithmic correction:

$$F(z) = 1 - \left(1 + \bar{\xi}\log z\right) z^{-\bar{\xi}}, \qquad f(z) = \bar{\xi}^{\,2}\log(z)\, z^{-\bar{\xi}-1}, \qquad S = \frac{(g-\mu)^2}{2\upsilon^2}, \tag{B.2}$$

which satisfies $F(1) = 0$, $f(1) = 0$, and $\int_1^\infty f\,\mathrm{d}z = 1$ automatically, with moments $\mathbb{E}\left[z^p\right] = \left(\bar{\xi}/(\bar{\xi} - p)\right)^2$ for $p < \bar{\xi}$ (as with the two-root moments, which require $p < \theta$; full-support moments diverge at p equal to the tail exponent, while on bounded intervals exponent collisions are evaluated by their logarithmic antiderivative limits). An initial distribution with a tail thinner than the floor, including one with bounded support, converges to this critical solution, while an initial tail fatter than the floor persists and leads to the two-root solution (12); this is the front-selection result of Luttmer (2012) and Benhabib, Perla, and Tonetti (2021). All expressions were verified against direct numerical evaluation of (2) and quadrature of the moment formulas.

C. The Corrected Equilibrium System and the Origin of the Tail

Equilibrium system. For reference against the code, the corrected BGP system for $\kappa > 0$ solves for $(g, \hat{z}, \Omega)$ from: (i) the export threshold condition $\hat{z} = d(\kappa/\bar{\pi}_{\min})^{1/(\sigma-1)}$; (ii) free entry, $\int_1^\infty v f \,\mathrm{d}z = x/\chi$; and (iii) value matching, $v(1) = \int_1^\infty v f \,\mathrm{d}z - x$. Here $\bar{z}$ is the aggregate productivity index (consumption per capita is $c = (1 - \tilde{L})\,\bar{z}$) and γ is the coefficient of relative risk aversion, with $\gamma = 1$ (log utility) throughout the quantitative analysis; the remaining objects are

$$\xi_2 = \tfrac{2(g-\mu)}{\upsilon^2} - \theta, \qquad S = \theta\left(g - \mu - \theta\tfrac{\upsilon^2}{2}\right), \qquad \tilde{L} = \Omega\left[(N-1)(1-F(\hat{z}))\kappa + \zeta\left(S + \tfrac{\delta}{\chi}\right)\right],$$

$$\bar{z}^{\,\sigma-1} = \Omega\left(\mathbb{E}\left[z^{\sigma-1}\right] + (N-1)d^{1-\sigma}\textstyle\int_{\hat{z}}^\infty z^{\sigma-1} f \,\mathrm{d}z\right), \qquad \bar{\pi}_{\min} = \frac{1-\tilde{L}}{(\sigma-1)\bar{z}^{\,\sigma-1}},$$

$$r = \rho + \gamma g + \delta, \qquad \lambda_{ii} = \left[1 + (N-1)d^{1-\sigma}\tfrac{\int_{\hat{z}}^\infty z^{\sigma-1} f \,\mathrm{d}z}{\mathbb{E}[z^{\sigma-1}]}\right]^{-1},$$

with f, F, and all moments from (12) and (B.1) (defined for exponents below the tail index θ), the value function from (7)–(10), and $x = \zeta$ under the baseline adoption-cost specification. The published quantitative system is recovered by substituting the truncated-Pareto distribution and its moments for the corrected two-root ones and setting the omitted option-value coefficient c_2 to zero; away from the boundary layer at the adoption threshold, these substitutions coincide with the corrected system in the $\upsilon \to 0$ limit. The equilibrium must satisfy (14), $S > 0$, $\theta > \sigma - 1$, $a > 0$, and $r > g$; the corrected code enforces these restrictions, which were not all active in the published code.

D. Numerical Methods for the Corrected Model

Estimation and stationary counterfactuals use the corrected closed-form balanced growth path (BGP), including the two-power productivity distribution and the smooth value function. The transition endpoints use a fully discretized BGP built from the same operators used along the transition. The transition couples an implicit finite-difference Hamilton–Jacobi–Bellman (HJB) equation to a Scharfetter–Gummel finite-volume Kolmogorov forward equation (KFE). Anderson acceleration jointly solves the complete growth and entry paths. The implementation uses 64-bit JAX, automatic differentiation for nonlinear derivatives, and tridiagonal linear algebra.

D.1. The system to be solved

Time is continuous, $t \in [0, T]$. At $t = 0$ an unanticipated permanent trade-cost reduction moves the iceberg cost from d_0 to

$$d_T = 1 + 0.9(d_0 - 1).$$

Productivity is normalized by the adoption threshold $M(t)$, and $x = \log z \in [0, \infty)$ is log productivity relative to that threshold. Let $v^{\text{lev}}(z, t)$, $F^{\text{lev}}(z, t)$, and $f^{\text{lev}}(z, t)$ denote the value function, CDF, and density in the level coordinate z. These are the objects denoted by v, F, and f in the main text. The corresponding log-coordinate objects are

$$v(x, t) = v^{\text{lev}}(e^x, t), \qquad \mathcal{F}(x, t) = F^{\text{lev}}(e^x, t), \qquad f(x, t) = \partial_x \mathcal{F}(x, t) = e^x f^{\text{lev}}(e^x, t).$$

The quantitative model has log utility, labor-denominated adoption costs, and fixed adoption cost $\zeta = 1$. Let $g = \dot{M}/M$, $D = \upsilon^2/2$, and $s = \sigma - 1$. The objects S and E are adoption and entry rates per active domestic variety, and Ω is the measure of varieties; ΩS and ΩE are the corresponding aggregate flows. On the post-shock transition, $d = d_T$.

Backward block. We solve for the rescaled value $\tilde{v}(x, t) = e^{-sx} v(x, t)$:

$$\begin{aligned}
\frac{\partial \tilde{v}}{\partial t} &= A(t)\tilde{v} - \tilde{\pi}(x, t), \qquad A(t) = \tilde{\rho}(t)\mathcal{I} - b_v(t)\partial_x - D\partial_{xx}, \\
b_v(t) &= \mu - g(t) + s\upsilon^2, \\
\tilde{\rho}(t) &= \rho + \delta + \frac{\mathrm{d}}{\mathrm{d}t}\log(1 - \tilde{L}(t)) - s\left(\mu - g(t) + s\frac{\upsilon^2}{2}\right), \\
\tilde{\pi}(x, t) &= \bar{\pi}_{\min}(t)\left[1 + (N-1)d^{1-\sigma}\mathbb{1}\left\{x \ge \log \hat{z}(t)\right\}\right] - (N-1)\kappa e^{-sx}\mathbb{1}\left\{x \ge \log \hat{z}(t)\right\}.
\end{aligned} \tag{D.1}$$

Here $\mathcal{I}$ is the identity operator. Smooth pasting at adoption and the far-field condition on the rescaled value are

$$\partial_x \tilde{v}(0, t) + s\tilde{v}(0, t) = 0, \qquad \lim_{x \to \infty} \partial_x \tilde{v}(x, t) = 0. \tag{D.2}$$

The terminal value is the value function of the terminal discrete BGP.

Forward block. Write $b(t) = \mu - g(t) < 0$ and $J = bf - D\partial_x f$. The normalized unit-mass distribution solves

$$\frac{\partial f}{\partial t} = -\partial_x J + S(t)f, \qquad f(0,t) = 0, \qquad S(t) = -J(0,t) = D\partial_x f(0,t), \tag{D.3}$$

$$f(x,t) \sim C(t)e^{-\theta x} \quad \text{as } x \to \infty. \tag{D.4}$$

Entry and exogenous exit cancel when the physical firm measure is converted into this conditional distribution. The initial condition is the initial discrete-BGP distribution. The upper condition in (D.4) carries the inherited Pareto exponent along the finite transition.

Statics and equilibrium conditions. Given the distribution and $(g, \hat{z}, E, \Omega)$ at an instant, define

$$\begin{aligned}
\tilde{L} &= \Omega\left[(N-1)(1-\mathcal{F}(\log\hat{z}))\kappa + \zeta\left(S + \frac{E}{\chi}\right)\right], \\
\bar{z}^{\,s} &= \Omega\left(\int_0^\infty e^{sx} f(x)\,\mathrm{d}x + (N-1)d^{-s}\int_{\log\hat{z}}^\infty e^{sx} f(x)\,\mathrm{d}x\right), \\
\bar{\pi}_{\min} &= \frac{1-\tilde{L}}{s\bar{z}^{\,s}}.
\end{aligned}$$

The three conditions imposed at every nonterminal time node are value matching, export indifference, and active free entry:

$$\begin{aligned}
\tilde{v}(0,t) &= \int_0^\infty \tilde{v}(x,t)e^{sx} f(x,t)\,\mathrm{d}x - \zeta, \\
\hat{z}(t)^s &= \frac{\kappa d^s}{\bar{\pi}_{\min}(t)}, \qquad \tilde{v}(0,t) = \zeta\frac{1-\chi}{\chi}.
\end{aligned} \tag{D.5}$$

The variety stock is predetermined at impact and evolves according to

$$\Omega(t) = \Omega_0 \exp\left\{\int_0^t (E(u) - \delta)\,\mathrm{d}u\right\}. \tag{D.6}$$

The corrected transition is therefore a forward–backward system: the KFE distribution is a state variable, while the HJB and the three equilibrium conditions determine the controls.

D.2. Spatial discretization

The continuum domain is truncated at $\bar{x} = 5$. The production spatial grid is piecewise uniform on $[0, 0.1]$, $[0.1, 1]$, and $[1, 5]$. The three segments are created with 1,440, 1,920, and 960 endpoint-inclusive nodes before duplicate junctions are removed. Write the resulting extended grid as $0 = y_0 < y_1 < \cdots < y_{P+1} = \bar{x}$, where $P = 4{,}316$. The HJB unknowns are at $y_1, \ldots, y_P$. The KFE excludes the absorbing barrier and includes the upper node, so it has $P + 1 = 4{,}317$ mass states.

D.2.1 The rescaled HJB

The first derivative in (D.1) uses the backward upwind difference because $b_v(t) < 0$ at the fitted parameters and along the accepted transition. The second derivative uses the centered formula for a nonuniform grid. Both operators are tridiagonal, yielding the standard monotone HJB discretization (Achdou et al., 2022). Let $h_0 = y_1 - y_0$. At the lower boundary, smooth pasting gives the ghost relation

$$\tilde{v}_0 = \Xi_1 \tilde{v}_1, \qquad \Xi_1 = \frac{1}{1 - sh_0}. \tag{D.7}$$

At high productivity, the transversal solution has the expansion

$$\tilde{v}(x,t) = a_\infty(t) + b_\infty(t)e^{-sx} + o(e^{-sx}).$$

Its derivative vanishes as $x \to \infty$. The finite-domain approximation in (D.2) is therefore implemented as

$$\tilde{v}_{P+1} = \tilde{v}_P. \tag{D.8}$$

This is a far-field condition on the rescaled value. It is not a reflecting boundary for the productivity process.

Exporter mass and productivity moments use fractional weights for the cell that contains $\log \hat{z}$. This makes the static aggregates continuous as the threshold moves across the grid. The profit vector in (D.1) uses the sharp export indicator. At export indifference the two flow-profit formulas coincide, so the indicator does not introduce a payoff jump at the root. The upper KFE mass is folded into the last HJB value in the value-matching quadrature using (D.8). All discrete statics and value-matching terms are finite sums over the represented KFE masses; the transition solver does not append an analytic tail beyond $\bar{x}$.

D.2.2 The Scharfetter–Gummel KFE

For the KFE, relabel its mass nodes as $x_i = y_{i+1}$ for $i = 0, \dots, P$ and treat the barrier as $x_{-1} = y_0 = 0$. The state f_i is probability mass assigned to x_i, not a pointwise density. Let $h_i^- = x_i - x_{i-1}$, let $h_i^+ = x_{i+1} - x_i$, and define $w_i = (h_i^- + h_i^+)/2$. At the upper node we set $h_P^+ = h_P^-$. With

$$B(y) = \frac{y}{e^y - 1}, \qquad B(0) = 1,$$

evaluated with `expm1`, the Scharfetter–Gummel (1969) jump rates are

$$\mathrm{down}_i = \frac{D\,B(bh_i^-/D)}{h_i^- w_i}, \qquad\qquad \mathrm{up}_i = \frac{D\,B(-bh_i^+/D)}{h_i^+ w_i}. \tag{D.9}$$

These rates are positive and reproduce a local exponential profile exactly. They avoid the artificial diffusion generated by first-order upwinding in an advection–diffusion equation.

For the interior balance rows, $1 \le i \le P-1$, the transport operator is

$$(Mf)_i = \mathrm{up}_{i-1}\, f_{i-1} + \mathrm{down}_{i+1}\, f_{i+1} - (\mathrm{down}_i + \mathrm{up}_i) f_i.$$

At the lower mass node,

$$(Mf)_0 = \mathrm{down}_1\, f_1 - (\mathrm{down}_0 + \mathrm{up}_0) f_0.$$

The lower absorption is $S_{\mathrm{abs}} = \mathrm{down}_0\, f_0$. Off the unit-mass manifold, the reinjection rate is evaluated as

$$S_h(f) = \frac{S_{\mathrm{abs}}}{m}, \qquad m = \sum_i f_i. \tag{D.10}$$

This normalization exactly cancels lower absorption in the represented mass. It does not remove the separate truncation defect at the upper row.

D.2.3 The inherited-tail row

On an inherited-tail BGP, the stationary log density is proportional to

$$e^{-\theta x} - e^{-\xi_2 x}, \qquad \xi_2 = \frac{2(g-\mu)}{v^2} - \theta.$$

The upper KFE row imposes the inherited exponent:

$$f_P = r_{\mathrm{tail}} f_{P-1}, \qquad r_{\mathrm{tail}} = e^{-\theta h_P^-} \frac{w_P}{w_{P-1}}. \tag{D.11}$$

This row fixes the tail decay rate, not its amplitude. It selects the inherited-tail solution on the truncated domain. A compact-support or zero-flux closure would instead select the critical stationary tail (Luttmer, 2012; Benhabib, Perla, and Tonetti, 2021). Holding the fitted analytic BGP growth rate fixed, the inherited and critical adoption rates are 0.162001 and 0.257016, respectively, so the boundary choice is economically material.

Replacing the final balance equation with (D.11) leaves a represented-mass error of order $e^{-\theta \bar{x}}$. The transition KFE neither clips nor renormalizes the density. It measures the mass error directly and rejects a transition that exceeds its tolerance. At $\bar{x} = 5$, the second power in the stationary density is negligible: $\xi_2 = 20.4662$ and $e^{-\xi_2 \bar{x}} = 3.6 \times 10^{-45}$ at the fitted analytic BGP.

D.3. Estimation and stationary balanced growth paths

We use two stationary backends. The analytic backend is efficient and differentiable and is used inside estimation, the stationary counterfactuals, the comparative statics, and the welfare decomposition. The discrete backend uses the transition operators and supplies the initial and terminal conditions for the transition. A separate post-estimation check compares the two.

D.3.1 Analytic backend and SMM

We hold $\theta = 4.988976587938262$ and $\sigma = 3.166924135838110$ at their published AER values. Under log utility, ρ is the data interest rate less data growth, and δ is set directly to the entry-rate target. We estimate

$$(d, \kappa, 1/\chi, \mu, \upsilon).$$

Every objective evaluation solves the corrected analytic BGP in

$$\phi = (g, \log(\hat{z} - 1), \log \Omega).$$

The BGP uses the corrected two-power distribution, the closed-form value function that is smooth at the export threshold, and analytic moments. The transformations impose $\hat{z} > 1$ and $\Omega > 0$ throughout the Newton solve.

The 12 targets are productivity growth, the home trade share, the exporter fraction, exporter size relative to nonexporter size, and the third and fourth rows of the five-year employment-quartile transition matrix. Exporter relative size follows the data's domestic-shipment convention:

$$\frac{\mathbb{E}[z^s \mid z \geq \hat{z}]}{\mathbb{E}[z^s \mid z < \hat{z}]}.$$

The criterion is

$$Q = \frac{1}{12}\left(\log m^{\text{data}} - \log m^{\text{model}}\right)' W \left(\log m^{\text{data}} - \log m^{\text{model}}\right), \qquad W = \operatorname{diag}(100, 100, 100, 100, 1, \ldots, 1). \tag{D.12}$$

The optimizer works in

$$(\log(d-1), \log \kappa, \log(1/\chi), \mu, \log \upsilon)$$

and uses BFGS with derivatives supplied by JAX. Five deterministic starts are used: a corrected-model neighborhood, the published parameter vector, and three seeded perturbations. Each nested BGP must also pass the economic admissibility screen: small equilibrium residuals; positive growth, adoption, and consumption; a valid trade share and value coefficient; a finite s-moment; and inherited-tail root ordering. Inadmissible trials return a nonfinite objective and are rejected by the line search. The best successful and admissible fit is retained.

The first four model moments are analytic. The eight firm-transition moments use a separate 500-node uniform grid on $[0, 7]$. The invariant weights are computed with the Scharfetter–Gummel rates and the inherited-tail stationary row. The five-year firm generator uses the same interior rates, a down-only upper row, and reinjection of adoption-barrier absorption according to those invariant weights. If L denotes this generator, the five-year transition operator is $\exp(5L)$. Cells that straddle employment-quartile boundaries are split fractionally, so each origin quartile contains exactly one quarter of stationary mass. The model chain omits the exogenous exit shock. The empirical rows are adjusted additively to sum to one, so both model and data moments condition on survival. The third- and fourth-quartile rows enter (D.12).

At the accepted fit,

$$d = 3.0482708, \quad \kappa = 0.0901186, \quad 1/\chi = 6.5197113, \quad \mu = -0.0324729, \quad \upsilon = 0.0563311.$$

The fitted point is well inside the inherited-tail region: $\xi_2 - \theta = 15.4773$. A post-estimation verifier evaluates the same point with numerically guarded analytic formulas.

D.3.2 Discrete stationary backend

For a trial g, the discrete quasi-stationary distribution (QSD) f and adoption rate S solve

$$\begin{aligned} &(M(g) + S\mathcal{I})f = 0 && \text{on the balance rows,} \\ &f_P - r_{\text{tail}} f_{P-1} = 0, && \sum_i f_i = 1. \end{aligned} \tag{D.13}$$

This is a square nonlinear system. Its Jacobian has a tridiagonal core, a border column equal to f on the balance rows and zero on the tail row, and the normalization as a border row. Each Newton step reduces to two tridiagonal solves and one scalar Schur complement. We take eight Newton steps from the analytic two-power density. At finite truncation, the solved quasi-stationary adoption rate S and the normalized lower absorption rate $S_h(f)$ differ by the small upper-row bookkeeping defect. The stationary BGP statics and the transition's impact node use the solved S. Positive-time transition statics use $S_h(f)$ from (D.10).

The discrete BGP then solves the HJB and the three conditions in (D.5) in $(g, \log(\hat{z} - 1), \log \Omega)$. Each trial nests the quasi-stationary solve and a tridiagonal stationary HJB solve. Automatic differentiation supplies the Newton Jacobian. The endpoint solver accepts a BGP only when both the QSD residual and the three-equation BGP residual are below 10^{-8}.

After estimation, we compare the accepted fit with numerically guarded analytic formulas and the discrete KFE–HJB backend. The comparison covers all 12 moments, the main stationary aggregates, the discrete QSD and BGP residuals, and the value-function shape. The discrete backend does not use the optimizer's closed-form density, value function, or analytic tail integrals.

D.3.3 Other reported stationary calculations

The welfare decomposition differentiates the corrected analytic BGP. Let $R(\phi, d) = 0$ denote its three equilibrium residuals. Automatic differentiation constructs the equilibrium response

$$\frac{\mathrm{d}\phi}{\mathrm{d}d} = -R_\phi^{-1} R_d,$$

and the partial derivatives of the feasible-consumption map $f_c(\Omega, \hat{z}, g; d)$. These derivatives deliver the direct-consumption, variety, export-threshold, consumption-through-growth, and direct-growth components reported in Section 5.1. We check that the linearized residual is second order in the perturbation and that the components sum to a separately differentiated total to relative error below 10^{-10}.

To reaggregate the published transition, we hold its path fixed and change only the terminal continuation. Replacing $Tg(T)$ with the path's accumulated $\log M(T)$ lowers utility by

$$e^{-\rho T}\frac{Tg(T)-\log M(T)}{\rho}.$$

For the GBM- and exit-scaling exercises, we anchor both BGPs at unit scale and continue outward using the nearest admissible solutions as warm starts. We stop when either BGP fails the admissibility checks.

D.4. Transition algorithm

The production time grid has 461 shared nodes: 161 nodes on $[0,10]$, 81 on $[10,20]$, and 221 on $[20,75]$, with duplicate junctions removed. The corresponding time steps are 0.0625, 0.125, and 0.25 years. Growth, entry, varieties, the KFE, and the HJB all use this grid.

D.4.1 Forward and backward steps

Given growth at the arrival node, the forward step is implicit Euler:

$$[\mathcal{I}-\Delta t_k\left(M(g_{k+1})+S_{h,k+1}\mathcal{I}\right)]\,f_{k+1}=f_k, \qquad S_{h,k+1}=\frac{\mathrm{down}_0\, f_{0,k+1}}{\sum_i f_{i,k+1}}, \tag{D.14}$$

with the last row replaced by (D.11). Conditional on $S_{h,k+1}$, this is one tridiagonal solve. We alternate the solve and the scalar update 12 times. The reported KFE diagnostic is the final Picard increment. It is supplemented by separate positivity and represented-mass checks.

The variety path in (D.6) is evaluated exactly for the piecewise-linear entry path. Given (f_k, S_k, Ω_k) and the next value function, the backward step solves for

$$\vartheta_k=(g_k, \log(\hat{z}_k-1), E_k).$$

For each trial ϑ_k, the static equilibrium is evaluated, and

$$\frac{\mathrm{d}}{\mathrm{d}t}\log(1-\tilde{L})_k \approx \frac{\log(1-\tilde{L}_{k+1})-\log(1-\tilde{L}_k)}{\Delta t_k}.$$

The implicit HJB step is

$$(\mathcal{I}+\Delta t_k A_k)\tilde{v}_k=\tilde{v}_{k+1}+\Delta t_k\tilde{\pi}_k. \tag{D.15}$$

The tridiagonal solve in (D.15) is nested inside a three-dimensional Newton solve for the conditions in (D.5). JAX differentiates through the linear solve and the static equations. The node root is warm-started from the next time node.

D.4.2 Outer fixed point and final certification

The outer unknown contains the complete node paths of (g, E). A map evaluation marches the KFE forward, constructs Ω from E, and marches the HJB backward. The initial entry guess is shifted once so that its integrated variety growth reaches the terminal-BGP variety level. No level restriction is imposed on later iterates. The terminal-variety discrepancy is an output of the free fixed point.

We begin with one plain undamped map evaluation. Thereafter, Anderson acceleration uses all available history up to a depth of 10 (Walker and Ni, 2011). The regularization ridge is 10^{-12}, the Euclidean fixed-point tolerance is 10^{-7}, and the maximum number of iterations is 80.

After convergence, the solver performs three certification steps. First, one undamped nonlinear map evaluation produces the reported controls. Second, the KFE and Ω are re-marched under those controls. Third, the backward equations are reevaluated at fixed controls on those re-marched states. The assembled residual is evaluated on this returned set of controls and re-marched states and measures their remaining inconsistency. The solver accepts the transition only if the final movement of both outer paths, the nonlinear node roots, the assembled residual, the KFE increment, density positivity, represented mass, entry positivity, and the terminal-variety gap pass their tolerances.

The node-root tolerance is 10^{-9}, and every node solve must report success. The assembled-equilibrium tolerance is 2×10^{-5}, and the final KFE Picard increment must be below 10^{-12}. Represented-mass error must be below 10^{-6}, density masses must exceed -10^{-10}, and entry must be nonnegative. Final growth- and entry-path movements must be below 10^{-7}, and the absolute terminal log-variety gap must be below 5×10^{-4}.

The terminal BGP supplies the terminal value function and terminal controls. The forward states $f(T)$, $S(T)$, and $\Omega(T)$ are marched to the horizon. At the terminal node, the solver combines those marched states with terminal-BGP g, $E = \delta$, $\hat{z}$, $\tilde{L}$, $\bar{z}$, and the home share for continuation and reporting. These terminal-BGP statics are not recomputed from the marched terminal distribution and variety stock. The terminal-variety discrepancy

$$\log \frac{\Omega(T)}{\Omega_T}$$

is measured and gated rather than imposed during the outer iteration.

D.4.3 Welfare

Let $c_k = (1 - \tilde{L}_k)\bar{z}_k$. For reporting and welfare, define the adoption-threshold growth path

$$g_0^W = g_{\text{initial BGP}}, \qquad g_k^W = g_k \quad (k \geq 1). \tag{D.16}$$

Appendix D.5 explains this impact convention. The adoption threshold is accumulated by trapezoidal integration:

$$\log M_k = \sum_{j<k} \frac{\Delta t_j}{2}(g_j^W + g_{j+1}^W).$$

Lifetime utility at impact is

$$
\begin{aligned}
U(0) = &\int_0^T e^{-\rho t}\left[\log M(t) + \log c(t)\right] \mathrm{d}t \\
&+ e^{-\rho T}\left[\frac{\log M(T) + \log c(T)}{\rho} + \frac{g_T}{\rho^2}\right],
\end{aligned}
\tag{D.17}
$$

where the finite integral is evaluated by the trapezoidal rule. The continuation uses the actual accumulated $\log M(T)$ and the terminal-BGP growth rate. For log utility, the reported consumption-equivalent gain is

$$\mathrm{CE} = \exp\{\rho[U_{\text{new}}(0) - U_{\text{old}}]\} - 1.$$

D.5. The impact reporting convention

We report the regular completion at the impact instant. This selection keeps the adoption threshold continuous, rules out an atom of adoption, maintains active entry, and assumes a finite right-hand growth limit and ordinary nondegenerate C^3 contact of the value function at the adoption boundary. Because physical productivities are predetermined, the normalized density is inherited at impact. Within this regularity class, the leading $\sqrt{t}$ term in the adoption flow is proportional to $g(0^+) - g(0^-)$. The coupled HJB, active-entry, and static equilibrium conditions require this term to vanish. Hence

$$g(0^+) = g(0^-), \qquad S(0^+) = S(0^-).$$

This conclusion is conditional on the regular completion.

For the accumulated adoption threshold and welfare, equation (D.16) replaces only the raw impact-node growth reading. All positive-time solver output is unchanged. Using the raw reading instead changes welfare by only 0.008 percentage points.