Online Appendix to "Platform Fee Floors, Market Power, and Missing Markets in Digital Payments: An Industrial-Organisation Calibration"

Craig S. Wright

University of Exeter Business School, University of Exeter, Exeter, United Kingdom

Corresponding author: cw881@exeter.ac.uk

ORCID: 0000-0001-9374-0507

Abstract

This online appendix contains the formal proofs, derivations, robustness details, and extended literature positioning for the manuscript.

Keywords payment systems; two-sided markets; platform market power; fixed fees; digital transition; industrial policy

---

This appendix provides full proofs of the equilibrium and welfare results (Section A.1), the full derivation of the feasible-flow-envelope integrals (Section A.2), the complete numerical derivation and endpoint-population discussion of the feasible-flow envelope (Section A.3), the robustness analysis for the welfare floor (Section A.4), and the benchmark-versus-production cost discussion (Section A.5), together with the market-power counterfactual (Section A.6), Budish's framework in full (Section A.7), and the detailed literature positioning (Section A.8), all referenced in the main text. Equation and section numbers prefixed "A" refer to this appendix; unprefixed references are to the main manuscript. Notation follows the main text: a unit mass of buyers with adoption-cost distribution $\tilde{B}$ and density $\tilde{\beta}$ on $[0, \bar{b}]$; a unit mass of sellers with acceptance-cost distribution H and density h on $[0, \bar{a}]$; transaction values $v \sim G$ with density g; per-transaction platform cost $q_\theta(v) = s_\theta + \tau_0^\theta + \tau_1^\theta v$ under architecture $\theta \in \{C, H\}$; fee schedule $f(v) = \phi_0 + \phi_1 v$; consumer reward rate $\beta\alpha$; per-pair transaction intensity μ_α; minimum viable transaction $v^* = \phi_0/(p_S - \phi_1)$.

A.1 Equilibrium and Welfare: Full Proofs

A.1.1 Adoption conditions and existence of a platform equilibrium

A buyer of type b adopts iff her expected net surplus over viable transactions, scaled by seller participation, covers her adoption cost:

$$\mu_\alpha \, n_S \int_{v^*}^{\infty} \big[(c_0 - \phi_0) + (\beta\alpha - \phi_1')v\big] g(v)\, dv \;\geq\; b, \tag{1}$$

and a seller of type a accepts iff

$$\mu_\alpha n_B \int_{v^*}^{\infty} \big[v(p_S - \phi_1) - \phi_0\big] g(v)\, dv \;\geq\; a. \tag{2}$$

Integrating (1) against $\tilde{B}$ and (2) against H gives the participation masses used in the main text. Define the buyer best-response participation level as a function of seller participation n_S, and the seller best-response as a function of buyer participation n_B:

$$\Phi_B(n_S) = \tilde{B}\Big(\mu_\alpha n_S \int_{v^*}^{\infty} \big[(c_0 - \phi_0) + (\beta\alpha - \phi_1')v\big]\, g(v)\, dv\Big), \tag{3}$$

$$\Phi_S(n_B) = H\Big(\mu_\alpha n_B \int_{v^*}^{\infty} \big[v(p_S - \phi_1) - \phi_0\big]\, g(v)\, dv\Big). \tag{4}$$

Proposition A.1 (Existence)**.** *Suppose $\tilde{B}$ and H are continuous and non-decreasing on $[0, \bar{b}]$ and $[0, \bar{a}]$ respectively, and that the bracketed integrals in* (3)–(4) *are non-negative at $v = v^*$. Then the composite map $T(n_B, n_S) = (\Phi_B(n_S), \Phi_S(n_B))$ has a fixed point in $[0,1]^2$, and hence a platform equilibrium exists.*

Proof. The proof has four steps: the two cross-side integrals are finite, so Φ_B and Φ_S are well-defined; T maps $[0,1]^2$ into itself; T is order-preserving; and Tarski's theorem then delivers a fixed point. Normalisation: $n_B, n_S \in [0,1]$ denote the *fractions* of the unit buyer and seller masses that participate, so the relevant domain is the unit square; this is the normalisation used throughout the main text.

Step 1: the integrals are finite, so the maps are well-defined. Write

$$I_B = \int_{v^*}^{\infty} \big[(c_0 - \phi_0) + (\beta\alpha - \phi_1')v\big] g(v)\, dv, \qquad I_S = \int_{v^*}^{\infty} \big[v(p_S - \phi_1) - \phi_0\big] g(v)\, dv.$$

Each integrand is affine in v, so $|\cdot| \leq K_1 + K_2 v$ for constants K_1, K_2 depending only on the fee and reward parameters. Hence

$$|I_B| \leq K_1\big[1 - G(v^*)\big] + K_2 \int_{v^*}^{\infty} v\, g(v)\, dv \leq K_1 + K_2\, \mathbb{E}[v],$$

and identically for I_S. The maintained assumption that the value distribution G has a finite mean, $\mathbb{E}[v] = \int_0^\infty v\, g(v)\, dv < \infty$ (stated in the main-text model and required for the welfare integrals to exist), makes both I_B and I_S finite real numbers for every admissible (ϕ_0, ϕ_1). Therefore the arguments $\mu_\alpha n_S I_B$ and $\mu_\alpha n_B I_S$ of $\tilde{B}$ and H in (3)–(4) are finite, and Φ_B, Φ_S are well-defined single-valued functions on $[0,1]$.

Step 2: self-map. A cumulative distribution function takes values in $[0,1]$. Since $\tilde{B}$ and H are CDFs, $\Phi_B(n_S) = \tilde{B}(\mu_\alpha n_S I_B) \in [0,1]$ and $\Phi_S(n_B) = H(\mu_\alpha n_B I_S) \in [0,1]$ for every $(n_B, n_S) \in [0,1]^2$, regardless of the sign of I_B, I_S (a CDF evaluated at a negative argument is simply 0). Hence $T = (\Phi_B \circ \pi_S,\ \Phi_S \circ \pi_B)$ maps $[0,1]^2$ into $[0,1]^2$, where π_S, π_B are the coordinate projections.

Step 3: monotonicity. $[0,1]^2$ under the coordinatewise order is a complete lattice. Take

$n'_S \geq n_S$. Because I_B does not depend on n_S, the hypothesis $I_B \geq 0$ (which follows from the bracketed integrand being non-negative at the lower limit $v = v^*$ together with $\beta\alpha - \phi'_1 \geq 0$ for a buyer rewarded net of buyer-side ad valorem cost, so that the integrand is non-negative on all of $[v^*, \infty)$) gives $\mu_\alpha n'_S I_B \geq \mu_\alpha n_S I_B$; monotonicity of the CDF $\tilde{B}$ then yields $\Phi_B(n'_S) \geq \Phi_B(n_S)$. For the seller side, $I_S \geq 0$ because its integrand $v(p_S - \phi_1) - \phi_0$ is non-negative exactly for $v \geq \phi_0/(p_S - \phi_1) = v^*$, the lower limit; hence $n'_B \geq n_B$ gives $\Phi_S(n'_B) \geq \Phi_S(n_B)$. Both components of T are non-decreasing in the relevant coordinate, so T is order-preserving on the lattice.

Step 4: conclusion. Tarski's fixed-point theorem states that an order-preserving self-map of a complete lattice has a non-empty set of fixed points, itself a complete lattice. (Tarski requires only monotonicity, not continuity; finiteness from Step 1 and the self-map property from Step 2 are what make T a well-defined monotone self-map.) Hence a fixed point $(n_B^*, n_S^*) \in [0,1]^2$ exists, and by the definition of platform equilibrium in the main text it is an equilibrium. □

Proposition A.2 (Uniqueness under a contraction condition)**.** *If in addition $\tilde{B}$ and H are continuously differentiable with bounded densities, and the cross-side feedback is weak in the sense that*

$$\sup_{(n_B,n_S)\in[0,1]^2} \left|\Phi'_B(n_S)\right| \cdot \left|\Phi'_S(n_B)\right| < 1, \tag{5}$$

then the equilibrium is unique.

Proof. Step 1: differentiability of the best-response maps. By hypothesis $\tilde{B}$ and H are continuously differentiable with densities $\tilde{\beta} = \tilde{B}'$ and $h = H'$ bounded by some $\bar{\beta}, \bar{h} < \infty$. The arguments $\mu_\alpha n_S I_B$ and $\mu_\alpha n_B I_S$ are affine (hence C^∞) in n_S and n_B respectively, with I_B, I_S finite constants by Step 1 of Proposition A.1. Composition of a C^1 CDF with an affine map is C^1, so Φ_B, Φ_S are continuously differentiable with

$$\Phi'_B(n_S) = \tilde{\beta}(\mu_\alpha n_S I_B)\,\mu_\alpha I_B, \qquad \Phi'_S(n_B) = h(\mu_\alpha n_B I_S)\,\mu_\alpha I_S,$$

and these derivatives are bounded: $|\Phi'_B| \leq \bar{\beta}\mu_\alpha |I_B|$ and $|\Phi'_S| \leq \bar{h}\mu_\alpha |I_S|$ on $[0,1]$.

Step 2: reduction to a scalar fixed point. (n_B, n_S) is a fixed point of T iff $n_B = \Phi_B(n_S)$ and $n_S = \Phi_S(n_B)$ hold simultaneously. Since Φ_S is a single-valued function (Step 1 of Proposition A.1), the second equation determines n_S uniquely as $n_S = \Phi_S(n_B)$ for any given n_B. Substituting into the first equation, (n_B, n_S) is a fixed point of T iff n_B solves

$$n_B = \Phi_B\big(\Phi_S(n_B)\big) =: \psi(n_B),$$

with n_S then recovered as $\Phi_S(n_B)$. The correspondence between fixed points of T and fixed points of ψ on $[0,1]$ is therefore exact and bijective: distinct fixed points of T have distinct n_B-coordinates (else equal n_B would force equal $n_S = \Phi_S(n_B)$), and each fixed point of ψ yields exactly one fixed point of T.

Step 3: ψ is a contraction. $\psi = \Phi_B \circ \Phi_S$ is C^1 on $[0,1]$ as a composition of C^1 maps, and by the chain rule

$$\psi'(n_B) = \Phi'_B\big(\Phi_S(n_B)\big) \cdot \Phi'_S(n_B).$$

Taking absolute values and bounding each factor by its supremum,

$$|\psi'(n_B)| \le \Big(\sup_{[0,1]} |\Phi_B'|\Big)\Big(\sup_{[0,1]} |\Phi_S'|\Big) < 1$$

by the cross-feedback condition (5). Since $|\psi'|$ is bounded by a constant $L < 1$ on the interval, the mean value theorem gives $|\psi(x) - \psi(y)| \le L|x - y|$ for all $x, y \in [0, 1]$: ψ is an L-Lipschitz contraction.

Step 4: conclusion. $[0, 1]$ is a complete metric space and ψ maps it into itself (Step 2 of Proposition A.1 gives $\Phi_B, \Phi_S \in [0, 1]$, so $\psi \in [0, 1]$). By the Banach fixed-point theorem ψ has exactly one fixed point n_B^*. By the bijection of Step 2 the platform equilibrium $(n_B^*, \Phi_S(n_B^*))$ is therefore unique. □

A.1.2 Total-margin characterisation

The platform chooses (ϕ_0, ϕ_1) to maximise

$$\pi(\phi_0, \phi_1) = n_B^* n_S^* \, \mu_\alpha \int_{v^*}^{\infty} \big[\phi_0 + \phi_1 v - q_\theta(v)\big] g(v)\, dv - F_\theta, \tag{6}$$

where (n_B^*, n_S^*) depends on (ϕ_0, ϕ_1) through (3)–(4) and $v^* = \phi_0/(p_S - \phi_1)$.

Before the first-order condition, the participation response $\partial(n_B^* n_S^*)/\partial\phi_0$ that enters it must be computed from the equilibrium system rather than treated as a primitive. The following lemma does so by the implicit function theorem; the contraction condition (5) is exactly what guarantees the relevant Jacobian is invertible.

Lemma A.3 (Participation comparative statics)**.** *Under the hypotheses of Proposition A.2, the equilibrium (n_B^*, n_S^*) is continuously differentiable in ϕ_0, and*

$$\frac{\partial n_B^*}{\partial \phi_0} = \frac{\Phi_{B,\phi_0} + \Phi_B'\, \Phi_{S,\phi_0}}{1 - \Phi_B'\Phi_S'}, \qquad \frac{\partial n_S^*}{\partial \phi_0} = \frac{\Phi_{S,\phi_0} + \Phi_S'\, \Phi_{B,\phi_0}}{1 - \Phi_B'\Phi_S'}, \tag{7}$$

where $\Phi_B' = \partial\Phi_B/\partial n_S$ and $\Phi_S' = \partial\Phi_S/\partial n_B$ are evaluated at the equilibrium, and $\Phi_{B,\phi_0} = \partial\Phi_B/\partial\phi_0$, $\Phi_{S,\phi_0} = \partial\Phi_S/\partial\phi_0$ are the direct (own-fee) partials holding the other side's participation fixed. The denominator $1 - \Phi_B'\Phi_S'$ is strictly positive by (5)*. Consequently*

$$\frac{\partial(n_B^* n_S^*)}{\partial\phi_0} = n_S^* \frac{\partial n_B^*}{\partial \phi_0} + n_B^* \frac{\partial n_S^*}{\partial \phi_0}. \tag{8}$$

Proof. Equilibrium is the solution of the system

$$\mathcal{F}(n_B, n_S; \phi_0) = \begin{pmatrix} n_B - \Phi_B(n_S; \phi_0) \\ n_S - \Phi_S(n_B; \phi_0) \end{pmatrix} = 0.$$

Each Φ is C^1 in (n_B, n_S, ϕ_0): differentiability in the participation arguments is Step 1 of Proposition A.2; differentiability in ϕ_0 holds because ϕ_0 enters Φ_B, Φ_S only through the finite integrals $I_B(\phi_0), I_S(\phi_0)$ and the threshold $v^*(\phi_0)$, all C^1 (the integrals by differentiation under the integral sign, valid since the integrands and their ϕ_0-derivatives are dominated by an integrable function under the finite-mean assumption; the threshold because $v^* = \phi_0/(p_S - \phi_1)$ is affine in ϕ_0). Hence $\mathcal{F}$ is C^1. Its Jacobian in (n_B, n_S) is

$$D_{(n_B,n_S)}\mathcal{F} = \begin{pmatrix} 1 & -\Phi_B' \\ -\Phi_S' & 1 \end{pmatrix}, \qquad \det = 1 - \Phi_B'\Phi_S'.$$

By (5), $|\Phi_B'\Phi_S'| < 1$, so $\det \neq 0$ and the matrix is invertible with

$$\left(D_{(n_B,n_S)}\mathcal{F}\right)^{-1} = \frac{1}{1 - \Phi_B'\Phi_S'} \begin{pmatrix} 1 & \Phi_B' \\ \Phi_S' & 1 \end{pmatrix}.$$

The implicit function theorem applies: (n_B^*, n_S^*) is C^1 in ϕ_0 and

$$\begin{pmatrix} \partial n_B^*/\partial\phi_0 \\ \partial n_S^*/\partial\phi_0 \end{pmatrix} = -\left(D_{(n_B,n_S)}\mathcal{F}\right)^{-1} \partial_{\phi_0}\mathcal{F} = \frac{1}{1 - \Phi_B'\Phi_S'} \begin{pmatrix} 1 & \Phi_B' \\ \Phi_S' & 1 \end{pmatrix} \begin{pmatrix} \Phi_{B,\phi_0} \\ \Phi_{S,\phi_0} \end{pmatrix},$$

since $\partial_{\phi_0}\mathcal{F} = (-\Phi_{B,\phi_0},\, -\Phi_{S,\phi_0})^\top$ and the two minus signs cancel. Carrying out the matrix product gives (7). Equation (8) is the product rule applied to $n_B^*(\phi_0)n_S^*(\phi_0)$. □

Proposition A.4 (Total margin equals inverse semi-elasticity)**.** *At an interior optimum* (ϕ_0^*, ϕ_1^*) *with* $(n_B^*, n_S^*) \in (0,1)^2$*, the total per-transaction margin satisfies*

$$\bar{f} - \bar{q} \;=\; \frac{\Lambda^*}{-\,\partial\Lambda^*/\partial\phi_0} \;=\; \frac{\Lambda^*}{-\,\partial\Lambda^*/\partial\phi_1^\dagger}, \tag{9}$$

where $\Lambda^* = n_B^* n_S^* \mu_\alpha \left[1 - G(v^*)\right]$ *is equilibrium volume,* $\bar{f}$ *is the volume-weighted average fee,* $\bar{q}$ *the volume-weighted average cost, and* $\phi_1^\dagger$ *denotes a unit value-weighted perturbation of the ad valorem rate. The two semi-elasticities are equated at the optimum.*

Proof. Write $\Pi(\phi_0, \phi_1) = \pi + F_\theta$ for the gross operating profit. Total differentiation of (6) with respect to ϕ_0, accounting for the dependence of (n_B^*, n_S^*) and v^* on ϕ_0, gives

$$\begin{aligned} \frac{\partial\Pi}{\partial\phi_0} = \; & \underbrace{n_B^* n_S^* \mu_\alpha \big[1 - G(v^*)\big]}_{=\Lambda^* \text{ (mechanical revenue effect)}} \; + \; \underbrace{\frac{\partial(n_B^* n_S^*)}{\partial\phi_0}\, \mu_\alpha \int_{v^*}^{\infty} \big[\phi_0 + \phi_1 v - q_\theta(v)\big]\, g\, dv}_{\text{participation effect}} \\ & - \; \underbrace{n_B^* n_S^* \mu_\alpha \left[\phi_0 + \phi_1 v^* - q_\theta(v^*)\right] g(v^*) \frac{\partial v^*}{\partial\phi_0}}_{\text{viability-margin effect}}. \end{aligned} \tag{10}$$

The first term is the mechanical gain from charging $d\phi_0$ more on every one of the Λ^* inframarginal transactions. The second is the change in margin on the transactions added or shed as participation responds, with $\partial(n_B^* n_S^*)/\partial\phi_0$ now given explicitly by (8) of Lemma A.3—it is not a free parameter but the equilibrium response of the two-sided system, finite because the contraction condition makes the Jacobian determinant $1 - \Phi_B'\Phi_S'$ strictly positive. The third is the loss from transactions that drop below the viability threshold v^* as it rises, with $\partial v^*/\partial\phi_0 = 1/(p_S - \phi_1) > 0$. Collecting the participation and viability terms, each is proportional to the response of volume to the fee; define the total volume response

$$\frac{\partial \Lambda^*}{\partial \phi_0} = \frac{\partial (n_B^* n_S^*)}{\partial \phi_0}\, \mu_\alpha [1 - G(v^*)] - n_B^* n_S^* \mu_\alpha\, g(v^*) \frac{\partial v^*}{\partial \phi_0},$$

in which the first summand is supplied by Lemma A.3. Setting $\partial \Pi / \partial \phi_0 = 0$ in (10) and writing the bracketed margin evaluated at the volume-weighted average as $\bar{f} - \bar{q}$, the first-order condition rearranges to

$$\Lambda^* + (\bar{f} - \bar{q}) \frac{\partial \Lambda^*}{\partial \phi_0} = 0 \quad \Longrightarrow \quad \bar{f} - \bar{q} = \frac{\Lambda^*}{- \partial \Lambda^* / \partial \phi_0}.$$

Repeating the identical computation for the value-weighted ad valorem perturbation $\phi_1^\dagger$—which enters (6) through the same three channels, with $\partial v^* / \partial \phi_1 = \phi_0 / (p_S - \phi_1)^2$ and the participation response again given by the analogue of (8) with ϕ_1-partials—gives the second equality in (9). Equating the two expressions for $\bar{f} - \bar{q}$ shows the two semi-elasticities coincide at the optimum. □

The welfare results in the main text use only the total margin and the volume response, not the split of (ϕ_0, ϕ_1) across the two sides, which depends on the relative densities $\tilde{\beta}, h$ and is not needed for what follows.

A.1.3 Welfare decomposition

Proposition A.5 (Intensive/extensive decomposition)**.** *Let W_θ be total surplus net of platform cost under architecture θ, holding participation (n_B, n_S) fixed. Then the gain from switching from C to H is*

$$W_H - W_C = \underbrace{n_B n_S \mu_\alpha \int_{v_C^*}^{\infty} \big[q_C(v) - q_H(v)\big] g(v)\, dv}_{\textit{intensive margin}} + \underbrace{n_B n_S \mu_\alpha \int_{v_H^*}^{v_C^*} \big[v p_S - q_H(v)\big] g(v)\, dv}_{\textit{extensive margin}}. \tag{11}$$

If trust costs are equal across architectures ($\tau_0^H = \tau_0^C,\ \tau_1^H = \tau_1^C$), the intensive margin equals $\Lambda_0 (s_C - s_H)$ with $\Lambda_0 = n_B n_S \mu_\alpha [1 - G(v_C^*)]$.

Proof. Total surplus net of platform cost under architecture θ, with viability threshold v_θ^*, is

$$W_\theta = n_B n_S \mu_\alpha \int_{v_\theta^*}^{\infty} \big[v p_S - q_\theta(v)\big] g(v)\, dv,$$

up to fixed and adoption costs that are differenced away below. Because lower cost weakly lowers the viability threshold, $v_H^* \le v_C^*$. Write the difference and split the H-integral at v_C^*:

$$\begin{aligned}
W_H - W_C &= n_B n_S \mu_\alpha \left[\int_{v_H^*}^{\infty} (vp_S - q_H) g\, dv - \int_{v_C^*}^{\infty} (vp_S - q_C) g\, dv \right] \\
&= n_B n_S \mu_\alpha \left[\int_{v_C^*}^{\infty} (vp_S - q_H) g\, dv + \int_{v_H^*}^{v_C^*} (vp_S - q_H) g\, dv \right. \\
&\qquad \left. - \int_{v_C^*}^{\infty} (vp_S - q_C) g\, dv \right].
\end{aligned}$$

The two integrals over $[v_C^*, \infty)$ combine: the vp_S terms cancel, leaving $\int_{v_C^*}^{\infty} (q_C - q_H) g\, dv$, the intensive margin. The remaining integral over $[v_H^*, v_C^*]$ is the extensive margin, the net surplus on transactions viable only under H. This is (11).

For the special case, $q_C(v) - q_H(v) = (s_C - s_H) + (\tau_0^C - \tau_0^H) + (\tau_1^C - \tau_1^H)v$. Under equal trust costs the last two parentheses vanish, so $q_C(v) - q_H(v) = s_C - s_H$, a constant. The intensive margin becomes

$$n_B n_S \mu_\alpha (s_C - s_H) \int_{v_C^*}^{\infty} g(v)\, dv = n_B n_S \mu_\alpha (s_C - s_H) \big[1 - G(v_C^*)\big] = \Lambda_0 (s_C - s_H). \qquad \square$$

Pass-through. The intensive margin in (11) is a resource saving and accrues as welfare regardless of how it is split between platform and users; it does not depend on pass-through. The extensive margin requires lower user-facing fees to induce the new transactions and therefore does depend on pass-through, vanishing at zero pass-through while the intensive margin is unchanged. The calibration multiplies only the extensive term, not the rectangle, by the pass-through-dependent factor.

Which trust components change. The switching cost s_θ falls by orders of magnitude under the low-fee settlement rail; the trust components τ_θ are less clear-cut. Fraud prevention (τ^{fraud}) may fall if on-chain auditability reduces detection cost, or be unchanged if detection technology is architecture-invariant. Compliance (τ^{comp}) may rise if regulators impose additional requirements on novel infrastructure, or fall if automated on-chain reporting reduces manual burden. Dispute resolution (τ^{disp}) is likely architecture-invariant for routine transactions; cryptographic finality may reduce disputes over settlement but not over goods or service quality. Because the direction of the net trust-cost change is ambiguous and not identified from public filings, the calibration in the main text varies the ratio τ_H/τ_C across its full range $[0, 1]$ rather than committing to a point estimate.

A.2 The Suppressed-Market Integrals: Full Derivation

This section derives in full the feasible relationship-flow envelope and count reported in Section 5.9 of the main text, under the Pareto relationship-flow density $f(x) = A\, x^{-(\theta+1)}$ on $x > 0$ and the calibration link (Assumption 5 in the main text). By the sequential-settlement assumption of the main text (Assumption 4), participation is decided on the viability condition $mx \ge \phi$ alone, so a relationship of flow x participates iff $x \ge \phi/m$, independently of the timing parameter b (which governs only the stage-two settlement frequency of relationships that already clear). The set excluded by the incumbent fee ϕ_C but admitted by the low-fee rail fee $\phi_H < \phi_C$ is $\{x : \phi_H/m \leq x < \phi_C/m\}$.

Proposition A.6 (Suppressed flow value)**.** *For $\theta \neq 1$, the total flow value of the excluded set is*

$$V_{\text{supp}} = \int_{\phi_H/m}^{\phi_C/m} x\, f(x)\, dx = \frac{A}{\theta - 1}\left[(\phi_H/m)^{-(\theta-1)} - (\phi_C/m)^{-(\theta-1)}\right], \tag{12}$$

which is finite for all strictly positive ϕ_H, ϕ_C. For $\theta = 1$ the value is $A\log(\phi_C/\phi_H)$.

Proof. Substitute the density:

$$V_{\text{supp}} = \int_{\phi_H/m}^{\phi_C/m} x \cdot Ax^{-(\theta+1)}\, dx = A\int_{\phi_H/m}^{\phi_C/m} x^{-\theta}\, dx.$$

For $\theta \neq 1$, the antiderivative of $x^{-\theta}$ is $x^{1-\theta}/(1-\theta)$, so

$$V_{\text{supp}} = A\left[\frac{x^{1-\theta}}{1-\theta}\right]_{\phi_H/m}^{\phi_C/m} = \frac{A}{1-\theta}\left[(\phi_C/m)^{1-\theta} - (\phi_H/m)^{1-\theta}\right].$$

Multiplying numerator and denominator by -1 and writing $1 - \theta = -(\theta - 1)$ gives

$$V_{\text{supp}} = \frac{A}{\theta - 1}\left[(\phi_H/m)^{-(\theta-1)} - (\phi_C/m)^{-(\theta-1)}\right],$$

which is (12). Both endpoints ϕ_H/m and ϕ_C/m are strictly positive and finite whenever $\phi_H, \phi_C > 0$, so each power term is finite and the integral is finite, irrespective of the value of θ relative to 1 or 2. For $\theta = 1$ the integrand is Ax^{-1}, whose antiderivative is $A\log x$, giving $V_{\text{supp}} = A[\log(\phi_C/m) - \log(\phi_H/m)] = A\log(\phi_C/\phi_H)$. □

Proposition A.7 (Suppressed count and its floor-sensitivity)**.** *For $\theta \neq 0$, the number of excluded relationships is*

$$N_{\text{supp}} = \int_{\phi_H/m}^{\phi_C/m} f(x)\, dx = \frac{A}{\theta}\left[(\phi_H/m)^{-\theta} - (\phi_C/m)^{-\theta}\right]. \tag{13}$$

As $\phi_H \to 0$ with $\theta > 0$, $N_{\text{supp}} \to \infty$. With ϕ_H held at any strictly positive value, both the count and the flow value are finite. As the low-fee floor tends to zero, the count diverges faster; when $\theta \geq 1$, the value integral also diverges, but more slowly than the count.

Proof. Substituting the density and integrating $x^{-(\theta+1)}$, whose antiderivative is $x^{-\theta}/(-\theta)$,

$$N_{\text{supp}} = A\int_{\phi_H/m}^{\phi_C/m} x^{-(\theta+1)}\, dx = A\left[\frac{x^{-\theta}}{-\theta}\right]_{\phi_H/m}^{\phi_C/m} = \frac{A}{\theta}\left[(\phi_H/m)^{-\theta} - (\phi_C/m)^{-\theta}\right],$$

which is (13). The leading term is $(A/\theta)(\phi_H/m)^{-\theta}$. For $\theta > 0$, $(\phi_H/m)^{-\theta} \to \infty$ as $\phi_H \to 0$, so $N_{\text{supp}} \to \infty$: the count is dominated by the lower limit and diverges as the low-fee rail fee approaches zero. The value integral (12) has leading term $(A/(\theta-1))(\phi_H/m)^{-(\theta-1)}$; the value remains finite at any strictly positive ϕ_H and grows, as $\phi_H \to 0$, at the slower rate $\phi_H^{-(\theta-1)}$ when $\theta > 1$ (logarithmically when $\theta = 1$), versus the count's $\phi_H^{-\theta}$. This is the count-divergence/value-slower-growth contrast reported in the text; the reported figures hold ϕ_H at its strictly positive calibrated value, where both are finite. □

A.2.1 Transaction count and the two settlement channels

The total number of transactions cleared on a rail at fee ϕ, used in Section 5 of the main text, integrates the per-relationship settlement frequency over participating relationships. A relationship of flow x participates iff $x \geq \phi/m$ and, when interior, settles at frequency $n^*(x) = bx/\phi$ transactions per period (one transaction whenever $bx/\phi < 1$, i.e. $x < \phi/b$). The closed form therefore depends on the ordering of the participation floor ϕ/m and the interiority cutoff ϕ/b, equivalently on whether $b \geq m$:

$$N(\phi) = \begin{cases} \displaystyle\int_{\phi/m}^{\infty} \frac{bx}{\phi} f(x)\, dx = \frac{b}{\phi} V(\phi), & b \geq m, \\ \displaystyle\int_{\phi/m}^{\phi/b} f(x)\, dx + \int_{\phi/b}^{\infty} \frac{bx}{\phi} f(x)\, dx, & b < m, \end{cases} \tag{14}$$

where $V(\phi) = \int_{\phi/m}^{\infty} x f(x)\, dx$ is value settled. When $b \geq m$, the participation floor lies at or above the interiority cutoff ($\phi/m \geq \phi/b$), so every participating relationship satisfies $x \geq \phi/m \geq \phi/b$, giving $bx/\phi \geq 1$: all participants are interior and the count collapses to the product form $(b/\phi)V(\phi)$. When $b < m$, relationships in $[\phi/m, \phi/b)$ participate but settle exactly once (the first integral, contributing unit count each), while relationships above ϕ/b settle at the interior frequency (the second integral). Differentiating either branch in ϕ exhibits the two channels reported in the text: the lower limit ϕ/m falls as $\phi \to 0$ (the extensive participation channel, new relationships entering the feasible set), and the integrand frequency bx/ϕ rises as $\phi \to 0$ (the intensive divisibility channel, interior relationships settling existing flow more finely). Both vanish at any $\phi > 0$ bounded away from zero and are released only as $\phi \to 0$.

A.3 The $18 Trillion Figure: Derivation, Decomposition, and the Endpoint Population

This section gives the full numerical derivation of the headline figure summarised in the estimation section of the main text, decomposes it across the surplus-rate and tail-index grid, states precisely what it does and does not represent, and develops the machine-to-machine (M2M) and artificial-intelligence (AI) agent settlement classes as the largest observable candidate endpoint populations consistent with it. Throughout, the figure is a model-implied Pareto extrapolation under the calibration-link assumption (Assumption 5 of the main text), not a measurement of an existing market, and the device and agent populations below are candidate *endpoint* universes, not estimates of realised payment demand, settlement frequency, or relationship value.

A.3.1 Step-by-step numerical derivation

The feasible-flow value is (12), with the logarithmic $\theta \to 1$ limit $V_{\text{supp}}|_{\theta=1} = A\log(\phi_C/\phi_H)$ relevant at the central calibration. The inputs are: the fitted tail index $\hat{\theta} = 1.02$ ($R^2 = 0.99$) and scale $\hat{A} = 1.59 \times 10^3$ (billions of relationships per dollar at $x = \$1$), both estimated on the Federal Reserve transaction-value tail over $x \geq \$15$; the incumbent per-transaction fee floor $\phi_C = \$0.10$; the low-fee rail fee $\phi_H = \$3 \times 10^{-6}$; and the surplus rate m, varied over a grid. The participation floors are ϕ_C/m and ϕ_H/m.

At $m = 0.10$, the integration interval for the feasible flow is

$$\left[\frac{\phi_H}{m}, \frac{\phi_C}{m}\right] = \left[\frac{3 \times 10^{-6}}{0.10}, \frac{0.10}{0.10}\right] = [3 \times 10^{-5},\ 1.0] \text{ dollars of flow.}$$

Evaluating the value integral at $\hat{\theta} = 1.02$ from (12),

$$V_{\text{supp}} = \frac{\hat{A}}{\hat{\theta} - 1}\left[(\phi_H/m)^{-(\hat{\theta}-1)} - (\phi_C/m)^{-(\hat{\theta}-1)}\right] = \frac{1.59 \times 10^3}{0.02}\left[(3 \times 10^{-5})^{-0.02} - (1.0)^{-0.02}\right].$$

The bracketed term is $(3 \times 10^{-5})^{-0.02} - 1$. Since $(3 \times 10^{-5})^{-0.02} = \exp(-0.02\ln(3 \times 10^{-5})) = \exp(0.02 \times 10.41) = \exp(0.2082) = 1.2315$, the bracket equals 0.2315, and

$$V_{\text{supp}} = \frac{1.59 \times 10^3}{0.02} \times 0.2315 = 7.95 \times 10^4 \times 0.2315 \approx 1.84 \times 10^4,$$

in billions of dollars of flow, i.e. approximately \$18.4 trillion. The logarithmic cross-check at $\theta = 1$, $V_{\text{supp}}|_{\theta=1} = A\log(\phi_C/\phi_H)$, gives $V_{\text{supp}} = \hat{A}\log(\phi_C/\phi_H) = 1.59 \times 10^3 \times \log(0.10/3 \times 10^{-6}) = 1.59 \times 10^3 \times \log(3.33 \times 10^4) = 1.59 \times 10^3 \times 10.41 \approx 1.66 \times 10^4$, i.e. approximately \$16.6 trillion: the $\theta = 1$ and $\theta = 1.02$ evaluations bracket the reported figure, confirming numerical stability across the singular point. The reported central figure of approximately \$18 trillion is the $\hat{\theta} = 1.02$ value; the logarithmic limit is the conservative end of the immediate neighbourhood.

A.3.2 Decomposition across the grid

The figure is sensitive to the tail index and, more weakly, to the surplus rate. The dependence on m enters only through the interval endpoints ϕ/m; because both endpoints scale by $1/m$ and the integrand is a power law, a change in m rescales V_{supp} by $m^{\theta-1}$, which at $\theta \approx 1$ is close to unity. The dependence on θ is the binding sensitivity. In the reported grid, the scale parameter is re-estimated when the calibration window changes, so the joint movement of $(\hat{A}, \hat{\theta})$ rather than θ alone determines the level. Evaluating (12) across the window-implied range gives approximately \$11 trillion at $\theta = 0.92$, the central \$18 trillion at $\theta = 1.02$, and larger values toward the upper end of the grid. The economically relevant statement is therefore directional: under every calibration in the grid, the feasible-flow envelope exceeds the entire \$9.76 trillion current US card market, while the point magnitude is tail-sensitive and should be read as a range rather than a single number.

A.3.3 Fee-floor and lower-tail attenuation bounds

The main manuscript adds a bounds table that varies both the fee-floor pair and the lower-tail continuation assumption. This subsection gives the arithmetic. Let $\rho \in [0,1]$ multiply the fitted relationship-flow intensity only over the unobserved excluded interval:

$$\mu_\rho(dx) = \rho A x^{-(\theta+1)} dx, \qquad x \in [\phi_H/m, \phi_C/m]. \tag{15}$$

The corresponding feasible-flow envelope is simply

$$V_\rho(\phi_C, \phi_H; m, \theta, A) = \rho\, V_{\text{supp}}(\phi_C, \phi_H; m, \theta, A). \tag{16}$$

This attenuation factor is not a probability and not an adoption rate. It is a transparent way to ask how much of the upper-tail-based lower-tail continuation can be removed before the economic implication disappears. The main table reports $\rho \in \{1, 0.50, 0.25, 0.10\}$ with $\theta = 1.02$, $\hat{A} = 1.59 \times 10^3$, and $m = 10\%$. The row $(\phi_C, \phi_H) = (\$0.10, \$0.000003)$ gives $V_{\text{supp}} = \$18.4$ trillion at $\rho = 1$ and therefore $\$9.2$, $\$4.6$, and $\$1.8$ trillion at $\rho = 0.50$, 0.25, and 0.10, respectively. Rows using a higher low-fee floor or a lower incumbent fixed fee reduce the envelope mechanically through the narrower integration interval.

This exercise is not meant to rescue a precise headline number. It does the opposite: it converts the central extrapolation into a transparent bounded object. The policy-relevant comparative static is stable even when the magnitude is stressed. The fixed fee determines the relationship-flow interval that can settle; reducing that fixed fee admits a lower tail of relationships that incumbent payment data cannot directly observe.

A.3.4 What the figure is and is not

The figure is the total *flow value* of economic relationships whose whole-period value flow falls in the interval $[\phi_H/m, \phi_C/m]$ under the fitted intensity measure: relationships that a positive incumbent per-transaction fee excludes from the feasible set, and that a value-independent near-zero fee admits into it. It is not a revenue forecast, not a measure of realised transactions, and not a prediction of adoption. Three boundaries are worth stating explicitly. First, it is bounded above by the maintained Pareto form: if the true relationship-flow distribution is not Pareto below the fitted tail, the extrapolation misstates the figure, and the deviation of the below-range observed densities from the fitted line (Figure 1 of the main text) is consistent with suppression but also with such misspecification. Second, it is a *feasible-flow* envelope: by the settlement proposition of the main text a near-zero fee is necessary for these relationships to settle, but it is not sufficient. Realised adoption in two-sided payment markets is governed by rewards, routing, steering, acceptance incentives and trust, not processing cost alone. Third, the corresponding relationship *count* is not reported as a magnitude, because under $\theta \approx 1$ it is dominated by the lower integration limit and diverges as $\phi_H \to 0$ (Section A.2); only the flow value, which is less floor-sensitive than the count but still depends on the fee ratio, is reported as a calibrated envelope.

A.3.5 The machine-to-machine endpoint population

The relationships in the excluded interval are, by construction, those whose entire periodic flow is below roughly one dollar at a ten-percent surplus rate—too small for any general-purpose card rail to settle at a ten-cent per-transaction floor. The economically natural candidates are not human retail payers but automated endpoints transacting for data, compute, bandwidth, sensor readings, or actuation at sub-cent unit values. The largest readily observable such endpoint population is the installed base of connected Internet-of-Things devices. IoT Analytics reports 18.5 billion connected IoT devices at the end of 2024, 21.1 billion at the end of 2025 (a 14% year-over-year increase), and a forecast of 39 billion by 2030 at a 13.2% compound annual growth rate.[1] Three features of this population make it the appropriate observable instance of the excluded interval. It is composed of fixed-function endpoints rather than general-purpose retail payers, so its natural transaction values are the per-reading, per-query, and per-actuation micro-amounts that the incumbent fee floor excludes. It is already generating the underlying economic activity—an estimated eighty zettabytes of data in 2025—whose monetisation at the level of the individual datum or query would require precisely the sub-cent settlement that a value-independent near-zero fee makes feasible. And its growth is driven in part by the declining cost of edge AI compute, which converts passive sensors into agents capable of initiating and receiving payments autonomously.

Two cautions apply and are stated rather than elided. The device count is a count of *connections*, not of economic relationships: a single relationship may span many devices, and many devices may never transact, so the count neither confirms nor sizes the feasible-flow figure, which comes entirely from the calibrated flow distribution. And the existence of a large endpoint universe does not establish realised demand: whether these endpoints transact depends on application design, standardisation, adoption incentives, routing, and trust mechanisms, none of which this paper forecasts. The device population identifies *where* in the economy the excluded interval would be populated if it is populated at all; it does not establish that it will be.

A.3.6 Autonomous AI agents as a settlement class

A second and overlapping candidate population is autonomous software agents. As edge and cloud AI compute costs decline, a growing share of economic transactions is initiated by software acting on behalf of a principal rather than by a human at a point of sale: an agent purchasing compute or model inference by the call, retrieving a paid data feed by the record, or compensating another agent for a sub-task at a fraction of a cent. These transactions share the defining property of the excluded interval—per-transaction values far below the incumbent fee floor, with no human-scale retail relationship to bundle them against—and so fall outside the feasible set of any general-purpose card rail for the same structural reason as M2M sensor settlement. The agent population is harder to count than connected devices and no comparable installed-base figure is available; it is mentioned here as a qualitatively distinct member of the same excluded class, not as an independent quantification. Its relevance is that it widens the can-

[1] IoT Analytics, "State of IoT 2025." These counts exclude smartphones, tablets, and personal computers; the installed base is therefore composed almost entirely of fixed-function endpoints, the class for which sub-cent autonomous settlement is the relevant payment mode rather than card-based retail payment.

didate endpoint universe beyond physical devices to include purely computational principals, reinforcing the point that the excluded interval is populated by automated sub-cent settlement rather than by human retail payments, while leaving the magnitude of realised demand—as for M2M—entirely outside what this paper estimates.

A.3.7 Why the excluded interval is invisible in payment data

The closing observation of the main text is made precise here. A relationship settles on a rail, and therefore appears in that rail's transaction statistics, only if its flow clears the rail's participation floor. Relationships with flow in $[\phi_H/m, \phi_C/m]$ clear the low-fee rail floor but not the incumbent floor; under the incumbent rail they do not transact, and so they generate no observation in card-payment data. The excluded interval is thus invisible not because the underlying economic activity is absent but because the missing transactions lie outside the feasible set the incumbent fee defines: the data record what cleared, and the cleared set is exactly the complement of the excluded set. This is why the figure cannot be obtained by direct measurement of payment flows and must instead be recovered, under an explicit and labelled assumption, by extrapolating the fitted relationship-flow distribution below the incumbent floor.

A.4 Robustness of the Welfare Floor

The \$700 million switching-cost floor in the main text rests on three inputs: the engineering cost gap $s_H \ll s_C$, US non-cash transaction volume, and the use of Visa's reported "network and processing" expense as a proxy for avoidable switching cost. This section examines the sensitivity of the floor to each.

A.4.1 Independence from the scaling law

The floor does not depend on the measured efficiency law $\eta(M) = 1 - 0.0473 \ln M$ ($R^2 = 0.9997$, bootstrap 95% CI for the slope $[0.0471, 0.0480]$, $n = 5$ measured points at $M \in \{1, 5, 10, 50, 100\}$). That law governs how per-node cost rises as nodes are added, and therefore affects estimates of cost at extreme throughput. The floor uses only the independently verified operating point: at 70,000 TPS the implied switching cost is $s_H \approx 1.6 \times 10^{-8}$ USD per transaction, computed from a per-node monthly cost of \$2,910 divided by the monthly transaction capacity at that rate. The per-tier switching-cost figures are:

Operating point	s_H (USD/tx)	Source
Halborn-verified (70K TPS)	1.58×10^{-8}	Independent verification
Production (0.33B TPS/node)	3.3×10^{-12}	$\frac{1}{3}$ measured peak
Measured peak (1.01B TPS/node)	1.1×10^{-12}	Internal scaling data

Even if the scaling law were entirely misspecified away from the verified point, s_H at the verified point remains roughly five orders of magnitude below $s_C = \$0.0033$. The floor $\Lambda_0(s_C - s_H)$ is therefore insensitive to the scaling law, since s_H is negligible relative to s_C at any plausible operating point.

A.4.2 Sensitivity to the volume input

US non-cash transaction volume is taken as approximately 200 billion per year (Federal Reserve Payments Study, which reports 204.5 billion non-cash payments for 2021). The floor scales linearly in this figure. Using the narrower general-purpose card count (153.3 billion in 2022) rather than all non-cash payments reduces the floor proportionally to approximately \$500 million; using all non-cash payments including ACH and checks raises it. The order of magnitude is robust to the choice of volume base within the range of published Federal Reserve figures.

A.4.3 Sensitivity to the cost proxy

The floor uses Visa's reported "network and processing" expense (\$778 million on 233.8 billion transactions, FY2024) as the avoidable switching cost. This is an accounting average, not a marginal cost, and it may include items that are not avoidable under a different architecture, or exclude switching-related costs booked elsewhere. If the true avoidable switching cost is half the reported figure, the floor halves; if it is twice, the floor doubles. The figure is reported transparently so that the reader can rescale. The qualitative conclusion—that the switching-cost saving is real but economically modest relative to the incumbent's operating margin—holds across this range.

A.5 Benchmark Versus Production Costs

The switching costs in the main text are benchmark figures from a controlled environment. Production-grade transaction processing incurs additional operational costs not present in a benchmark: redundancy and failover, continuous monitoring, operational staffing, software maintenance, governance, and compliance systems. This section explains how these are treated.

These additional costs are classified as trust/operational costs and absorbed into τ_H rather than the switching cost s_H, because they arise from operating a regulated service rather than from the computational act of validating and recording a transaction. If some of them are better classified as processing—for instance, the compute overhead of running redundant validation—then the effective s_H would be higher than the benchmark figure. The relevant question for the welfare floor is whether any such reclassification could raise s_H to a level comparable to $s_C = \$0.0033$. It cannot: the benchmark s_H is on the order of 10^{-8} USD per transaction, and even a reclassification inflating it by several orders of magnitude leaves it far below the incumbent figure. The floor is therefore robust to the benchmark-versus-production distinction. Production overheads matter for the absolute level of total cost, not for the switching-cost gap that the floor measures.

A.6 Upper-Bound Market-Power Counterfactual

This section gives the full accounting decomposition, surplus-stakes table, and entry-barrier discussion summarised in the upper-bound counterfactual subsection of the main text. The welfare estimates in the main calibration hold market structure fixed: the incumbent keeps its margin and only the cost base changes. This section asks a different question—the upper bound on welfare gains if, in addition to lower switching cost, the incumbent's margin were also competed away. This is a counterfactual decomposition, not a prediction about the effects of the technology.

Visa's FY2024 filings provide the accounting inputs. Net revenue was \$35.9 billion on 233.8 billion processed transactions, yielding an average network fee of \$0.154/tx (Visa Q4 FY2024 Earnings Release, October 29, 2024: net revenue \$35,926M and 233.8B processed transactions, pages 1, 3, 6; this is the network's own revenue and excludes interchange paid to issuers and acquirer processing fees paid to acquirers). Operating expenses were \$0.053/tx, giving an operating margin of \$0.101/tx. The margin per transaction is roughly thirty times the reported network-and-processing cost (\$0.0033/tx).

Visa's operating margin is not automatically deadweight loss. It includes return on invested capital, compensation for risk, payment for network investment and brand development, regulatory compliance cost, and monopoly rent. Eliminating the margin is not the same as creating social welfare unless the deadweight-loss component is separately identified. The table below therefore reports *surplus stakes*—the maximum transfer from producer to consumer/merchant surplus if fees fell to the entrant's cost—not welfare gains. The welfare gain from margin compression depends on the deadweight-loss share (output distortion, allocative inefficiency from above-cost pricing) and the resource-waste share (rent-seeking expenditure). This paper does not estimate either.

If fees fell from \$0.154/tx toward the entrant's total cost ($q_H = s_H + \tau_H$), the implied surplus stakes on 200 billion US non-cash transactions would be:

Scenario	\$/tx	\$/year
Cost reduction only (Δs)	0.003	0.7B
Margin transfer (fee $\to q_C$)	0.101	20.2B
Margin + cost (fee $\to q_H, \tau_H = \tau_C$)	0.105	20.9B
Margin + partial trust (fee $\to q_H$, $\tau_H = 0.5\tau_C$)	0.130	25.9B

The first row is a welfare gain (real resource saving). The remaining rows are surplus-transfer upper bounds: only the deadweight-loss component of the margin would constitute additional welfare from competition; the remainder is redistribution. The \$700 million floor and the \$20 billion decomposition are not additive; they answer the same question under different market-structure assumptions. The distance between them measures the market-structure stakes. Neither figure is the feasible-flow estimate: the \$700 million is the resource-saving floor on existing transactions and the \$20 billion is the market-power margin on those transactions, whereas the feasible-flow envelope (\$18 trillion feasible flow value) concerns relationships outside the current feasible set entirely. The three quantities answer three different questions and are not combined.

A.6.1 Why low switching cost does not by itself predict entry

A low switching cost weakens one server-cost rationale for decreasing average cost. It does not establish that entry would follow. At least six barriers remain.

(1) Interchange and rewards. Card platforms use interchange and rewards to solve two-sided adoption problems. A low-cost entrant without rewards, brand trust, routing incentives or merchant acceptance may fail to attract consumers even if its switching cost is negligible. Cost-structure entry and regulatory intervention are not interchangeable.

(2) Fraud-detection scale. The DOJ's 2024 Visa complaint (§108) identifies fraud-detection data as a barrier: smaller networks lack the transaction data needed for robust fraud detection.

(3) Network effects. Two-sided adoption: consumers prefer widely-accepted networks; merchants prefer widely-used ones. Low processing cost helps supply but not demand.

(4) Regulation. PCI-DSS, Durbin routing rules, PSD2 compliance, and POS integration all require regulatory access.

(5) ACH precedent. ACH costs \$0.006/tx—far below Visa's fee—yet has not displaced cards at POS, because it lacks rewards, real-time authorisation, chargebacks, and network effects.

(6) Brand trust. Consumer willingness to adopt an unfamiliar payment mechanism is not determined by processing cost.

A.7 Budish's Framework in Full

This section states in full the three equations and calibration of Budish (2025) summarised in the background section of the main text. Budish studies the cost of providing trust through Nakamoto's proof-of-work mechanism; his framework centres on three relations, stated here in his notation.

Equation 1 (Zero-profit condition on miners). The flow revenue to miners (block rewards p_{block} plus transaction fees) must cover their flow costs. In equilibrium with free entry, miners earn zero economic profit: $N^* p_{\text{block}} = N^* c$, where N^* is the number of miners and c the per-miner flow cost.

Equation 2 (Attacker incentive-compatibility). The cost of a majority attack must exceed the attacker's benefit: $\text{Cost}_{\text{attack}} \geq V_{\text{attack}}$, where V_{attack} is the value extractable through double-spending or reversion. Budish shows V_{attack} is increasing in the economic importance of the system.

Equation 3 (Equilibrium constraint). Combining the two, the flow cost of mining must satisfy $N^* c \geq g(V_{\text{attack}})$ with g increasing. For proof-of-work the relationship is linear: the cost of trust scales proportionally with V_{attack}, with no economies of scale. Doubling the value at risk doubles the required mining expenditure.

Budish's calibration (his Table IV) makes the scale concrete. At the attacker-optimal majority share of 59%, per-transaction security cost is 0.0025% of the value secured against attack. Securing against a \$1 billion attack requires annual security expenditure of \$2.63 trillion; against a \$40 billion attack, \$105 trillion, approximately global GDP. Per transaction, assuming 2,000 transactions per block, securing against a $1 billion attack is $25,000 per transaction. In his appendix discussion of increasing throughput, Budish observes that a $1{,}000\times$ throughput increase would reduce this to $25 per transaction—"a promising response"—but notes the subtlety that V_{attack} might grow with throughput.

Budish also provides the institutional-trust benchmark. In his Section VI he models traditional deterrence (Becker, 1968) as having per-transaction cost $F/V_{\text{honest}} + c_{\text{enf}}$, where F is the fixed investment in courts, police, and regulatory infrastructure and c_{enf} the variable enforcement cost per unit of honest activity. The United States spends approximately $300 billion annually on prisons, courts, and policing, a fixed cost that secures trillions in honest economic activity, so F/V_{honest} is decreasing in V_{honest}: institutional deterrence exhibits economies of scale that Budish's formulation of Nakamoto trust lacks. The low-fee settlement rail recovers the property by operating within the existing legal trust domain, so its per-transaction trust cost τ_H is institutional rather than mechanism-determined.

The intuition is that proof-of-work fuses two functions, processing and security: the same miners who build blocks are compensated through rewards and fees large enough to deter attacks. Adding hash rate increases security but not throughput, which is fixed by protocol block-size and interval parameters; because processing is performed by agents whose compensation must cover security costs, processing cost inherits security's linear scaling with value. The low-fee settlement rail breaks this coupling by separating the processing function from the security function.

A.8 Related Literature in Detail

This section expands the positioning summarised in the introduction of the main text across three literatures.

Blockchain economics. Budish (2025) is the direct predecessor. Huberman et al. (2021) model Bitcoin as a congestion-pricing system; capacity expansion addresses the throughput constraint that drives their congestion result, though their monopoly-pricing analysis remains relevant for any system with capacity limits. Easley et al. (2019) study the evolution of mining fees; this paper separates the fee structure from the security mechanism. Catalini and Gans (2020) identify verification and networking as blockchain's fundamental costs; this paper formalises the verification-cost reduction and embeds it in a welfare model. Hendrickson et al. (2016) and Garratt and Wallace (2018) study digital-currency competition and policy; Auer et al. (2025) study governance of distributed-ledger money; Bolt et al. (2024) analyse the Bank of Amsterdam as historical payment infrastructure.

Two-sided payment markets. Rochet and Tirole (2003) establish the foundational model of two-sided markets, which I adapt. Wright (2004) and Schmalensee (2002) study interchange fee determination; my model extends this to endogenous infrastructure investment. Rysman (2009) surveys the empirical literature on payment card markets. The key difference is that I treat the platform's cost structure as a choice variable, through infrastructure investment, rather than a parameter.

IT and returns to scale. Lashkari et al. (2024) show that IT investment shifts production from variable to fixed costs, generating increasing returns. Elliott et al. (2025) derive economies of scale in mobile networks from engineering first principles (Shannon capacity) and embed them in a structural IO model. This paper performs the analogous operation for payment processing: deriving cost structures from engineering data and embedding them in an equilibrium model. Kalyani et al. (2025) study the diffusion of new technologies across space and skill levels; the adoption dynamics here are a within-industry analogue.

References

Raphael Auer, Cyril Monnet, and Hyun Song Shin. Distributed ledgers and the governance of money. *Journal of Financial Economics*, 167(C), 2025. doi: https://doi.org/10.1016/j.jfineco.2024.103986.

Gary S. Becker. Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2):169–217, 1968. doi: https://doi.org/10.1086/259394.

Wilko Bolt, Jon Frost, Hyun Song Shin, and Peter Wierts. The bank of Amsterdam and the limits of fiat money. *Journal of Political Economy*, 132(12):3919–3941, 2024. doi: https://doi.org/10.1086/732524.

Eric Budish. Trust at scale: The economic limits of cryptocurrencies and blockchains. *The Quarterly Journal of Economics*, 140(1):1–62, 2025. doi: https://doi.org/10.1093/qje/qjae033.

Christian Catalini and Joshua S. Gans. Some simple economics of the blockchain. *Communications of the ACM*, 63(7):80–90, 2020. doi: https://doi.org/10.1145/3359552.

David Easley, Maureen O'Hara, and Soumya Basu. From mining to markets: The evolution of Bitcoin transaction fees. *Journal of Financial Economics*, 134(1):91–109, 2019. doi: https://doi.org/10.1016/j.jfineco.2019.03.004.

Jonathan T. Elliott, Georges V. Houngbonon, Marc Ivaldi, and Paul T. Scott. Market structure, investment, and technical efficiencies in mobile telecommunications. *Journal of Political Economy*, 133(5):1401–1459, 2025. doi: https://doi.org/10.1086/734132.

Rodney Garratt and Neil Wallace. Bitcoin 1, bitcoin 2. . . : An experiment in privately issued outside monies. *Economic Inquiry*, 56(3):1887–1897, 2018. doi: https://doi.org/10.1111/ecin.12569.

Joshua R. Hendrickson, Thomas L. Hogan, and William J. Luther. The political economy of bitcoin. *Economic Inquiry*, 54(2):925–939, 2016. doi: https://doi.org/10.1111/ecin.12291.

Gur Huberman, Jacob D. Leshno, and Ciamac Moallemi. Monopoly without a monopolist: An economic analysis of the Bitcoin payment system. *Review of Economic Studies*, 88(6): 3011–3040, 2021. doi: https://doi.org/10.1093/restud/rdab014.

Aakash Kalyani, Nicholas Bloom, Marcela Carvalho, Tarek Hassan, Josh Lerner, and Ahmed Tahoun. The diffusion of new technologies. *The Quarterly Journal of Economics*, 140(2): 1299–1365, 2025. doi: https://doi.org/10.1093/qje/qjaf002.

Danial Lashkari, Arthur Bauer, and Jocelyn Boussard. Information technology and returns to scale. *American Economic Review*, 114(6):1769–1815, 2024. doi: https://doi.org/10.1257/aer.20220522.

Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003. doi: https://doi.org/10.1162/154247603322493212.

Marc Rysman. The economics of two-sided markets. *Journal of Economic Perspectives*, 23 (3):125–143, 2009. doi: https://doi.org/10.1257/jep.23.3.125.

Richard Schmalensee. Payment systems and interchange fees. *Journal of Industrial Economics*, 50(2):103–122, 2002. doi: https://doi.org/10.1111/1467-6451.00170.

Julian Wright. The determinants of optimal interchange fees in payment systems. *Journal of Industrial Economics*, 52(1):1–26, 2004. doi: https://doi.org/10.1111/j.0022-1821.2004.00214.x.