Online Appendix to Competition Creates Competition

Austin Li

A. Full analytical arguments

This appendix develops the probability space, proves the main results and extensions, and states the regularity conditions used in the sale-design calculations. Sections B and C describe the equilibrium certificates and numerical exercises.

A.1. Probability space, strategies, and conditional information

The benchmark has a risk-neutral seller, an incumbent, a potential challenger, a strategic investor, competitive market makers, and exogenous noise demand. A known target standalone value is normalized to zero. The seller commits before trading to a binding second-price cash auction with reserve p. The incumbent has value $R \sim U[0, r]$, pays no incremental participation cost at the modeled stage, and learns its value before bidding. The challenger has value ℓ or h with equal probability, where $0 < p < \ell < r < h$. Neither bidder trades target equity. An outside investor observes challenger quality and chooses an order in $[-1, 1]$, pays $k|q|$, and earns $q(V_T - P) - k|q|$. Its order unit is not a controlling ownership stake. All payoffs and preparation costs use the same per-share normalization.

Noise has density $f(z) = e^{-|z|/b}/(2b)$, with $b > 1$ and $k > 0$. Market makers observe $X = q + Z$ and price $P(X) = \mathbb{E}[V_T \mid X]$. The challenger observes P but not X, then its independent preparation cost, equal to c_L with probability ρ and c_H otherwise, where $0 < \rho < 1$ and $0 \le c_L < c_H$. Paying the cost reveals its value and permits bidding; declining means absence. Both bidders bid truthfully once informed. A bid meeting the reserve is admissible. I prescribe preparation at equality of gross expected profit and cost. The allocation, cash payment, and stock payoff are then realized. The equilibrium requires Bayesian pricing, optimal preparation at every reached price and cost, and global investor best responses against the fixed equilibrium price and preparation schedules, permitting mixed orders.

I write t_0 for expected target proceeds without challenger entry, t_H, t_L for proceeds conditional on entry and quality, and g_H, g_L for conditional gross challenger profits. Set $\Delta_T = t_H - t_L$, $B_r(\mu) = g_L + \mu(g_H - g_L)$, $m = (1 + e^{2/b})^{-1}$, and $M = 1 - m$. For the comparison $p < \ell < r_0 < r_1 < h$, the hypotheses of Proposition 2, referred to below as (A1) to (A3), are

$$
\begin{aligned}
&0 \le c_L < B_{r_1}(m), \\
&B_{r_0}(1/2) < c_H < B_{r_1}(M), \\
&\Delta_T(r_0) < k < (1 - 1/b)\rho m \Delta_T(r_1).
\end{aligned}
\tag{OA.1}
$$

To define beliefs and deviations, let $\Theta \in \{L, H\}$ have equal probabilities, with acquisition values $v_L = \ell$ and $v_H = h$. Independently draw incumbent value R, preparation cost C, noise demand Z, and an auxiliary random variable U uniform on $[0, 1]$. A mixed trader strategy is a probability kernel σ_Θ on the Borel subsets of $[-1, 1]$. It can be implemented by a Borel quantile map $q_\Theta(U)$. Conditional randomization therefore contains no information about the other primitives beyond the trader's specified signal.

All spaces are standard Borel. Prices, entry policies, and payoffs are Borel functions. Because values and the order interval are bounded, all acquisition and trading payoffs are integrable. A regular conditional distribution of quality given price exists. Expectations conditional on a price atom use its complete preimage, not a pointwise inversion that presumes injectivity.

For the benchmark, define

$$\begin{aligned} a_\theta(x) &= \int_{[-1,1]} f(x-q)\,\sigma_\theta(dq), \\ \nu(dx) &= \frac{a_H(x)+a_L(x)}{2}\,dx, \\ \mu_X(x) &= \frac{a_H(x)}{a_H(x)+a_L(x)}. \end{aligned} \tag{OA.2}$$

The density f is positive everywhere, bounded, continuous, and integrates to unity. Tonelli's theorem gives $\int a_\theta = 1$. Dominated convergence in q implies continuity of $a_\theta(x)$, and positivity gives a continuous posterior. The distribution of X under every pure or mixed unilateral order deviation is absolutely continuous with respect to Lebesgue measure. Moreover, ν is equivalent to Lebesgue measure because its density is strictly positive. Thus an equilibrium identity holding ν-almost surely also holds almost surely under every unilateral trading deviation. No deviation can exploit a different version of prices or beliefs on a null set.

Write $\mathcal{G} = \sigma(P(X))$ and $\mu_P = \mathbb{E}[\mathbf{1}\{\Theta = H\} \mid \mathcal{G}]$. Since $P(X)$ is measurable with respect to X,

$$\mu_P = \mathbb{E}[\mu_X(X) \mid \mathcal{G}]. \tag{OA.3}$$

For any $q, q' \in [-1, 1]$, the triangle inequality implies

$$||x-q| - |x-q'|| \le |q-q'| \le 2. \tag{OA.4}$$

Exponentiating and integrating with respect to $\sigma_H(dq)\sigma_L(dq')$ gives $e^{-2/b}a_L \le a_H \le e^{2/b}a_L$, hence $m \le \mu_X \le M$. Equation (OA.3) preserves these bounds for μ_P. These conclusions do not impose pure trading, monotone orders, or an informative price.

The bidder's preparation rule is $\mathbf{1}\{C \le B_r(\mu_P)\}$, with entry at equality. When $c_L < B_r(m)$, the cost-averaged rule is bounded below by ρ. Conditional on Θ and X, the variables R, C remain independent of the investor's action and retain their specified laws. This is why the auction-stage expectations can be used inside the pricing equation. It also explains why a unilateral order deviation shifts the density of X without changing the conditional auction formulas.

For the complementary-signal extension, I replace the investor's observation of Θ by T and add Y. The randomization U, noise Z, incumbent R, and cost C remain independent of (Θ, T, Y); T and Y are independent conditional on Θ. Sections A.7 and C.3 specify the resulting conditional laws rather than treating the buyer's private signal as public.

A.2. Auction implementation and Proposition 1

Conditional on the highest competing admissible bid, truthful bidding maximizes a private-value bidder's payoff in a second-price auction. A bid above value can turn a loss into an unprofitable win, and a bid below value can discard a profitable win. Conditional on winning, the bidder's own bid does not determine payment. The presence of a public price or information conveyed by entry changes beliefs about other bidders but not this pointwise dominance argument. I use truthful implementation to fix payoff-equivalent bid descriptions.

Without the challenger, the seller receives p if $R \geq p$. A high-value challenger wins almost surely and pays $\max(p, R)$ because h exceeds incumbent support. When the challenger has value ℓ, the winning value is $\max(R, \ell)$ and the runner-up is $\min(R, \ell)$, so seller payment is $\max(p, \min(R, \ell))$. The challenger obtains $(\ell - \max(p, R))_+$.

The uniform formulas follow by splitting at p and ℓ:

$$
\begin{aligned}
t_0 &= \frac{1}{r} \int_p^r p \, du, \\
t_H &= \frac{1}{r} \left[p^2 + \frac{r^2 - p^2}{2} \right], \\
t_L &= \frac{1}{r} \left[p^2 + \frac{\ell^2 - p^2}{2} + \ell(r - \ell) \right], \\
g_L &= \frac{1}{r} \left[p(\ell - p) + \frac{(\ell - p)^2}{2} \right], \qquad g_H = h - t_H.
\end{aligned}
\tag{OA.5}
$$

Evaluating these expressions gives $t_0 = p(1 - p/r)$, $t_H = r/2 + p^2/(2r)$, $t_L = \ell - (\ell^2 - p^2)/(2r)$, $g_H = h - t_H$, and $g_L = (\ell^2 - p^2)/(2r)$. Subtracting gives $\Delta_T = (r - \ell)^2/(2r)$. On $p < \ell < r < h$, the derivatives are $\Delta'_T = 1/2 - \ell^2/(2r^2) > 0$, $g'_H = -1/2 + p^2/(2r^2) < 0$, and $g'_L = -(\ell^2 - p^2)/(2r^2) < 0$. Thus $\partial_r B_r(\mu) < 0$ at every fixed posterior. In particular $g_H - g_L > 0$, since the high challenger has a strictly higher value in a mechanism whose allocation is monotone in its own value.

For a general continuous incumbent CDF F on $[0, \bar{r}]$, the payment difference is zero for $R \leq \ell$ and $R - \ell$ for $R > \ell$. By writing a positive part as an integral of indicators,

$$
(R - \ell)_+ = \int_\ell^{\bar{r}} \mathbf{1}\{R > u\} \, du, \qquad (\theta - \max(p, R))_+ = \int_p^\theta \mathbf{1}\{R < u\} \, du.
\tag{OA.6}
$$

Tonelli applies because the integrands are nonnegative and the domains bounded. CDF continuity removes the distinction between $R < u$ and $R \leq u$ in expectation; even without continuity the Lebesgue integral over u is unaffected by countably many atoms. Thus the two expressions are respectively $\int_\ell^{\bar{r}} (1 - F(u)) du$ and $\int_p^\theta F(u) du$, extending $F = 1$ above support. A first-order stochastic strengthening lowers F pointwise. Nonnegative integral differences prove weak ordering; a strictly positive integral difference proves strictness. Posterior-weighted challenger profits inherit the ordering. This completes Proposition 1.

A.3. Price sufficiency, residuals, and convolution regularity

For any candidate equilibrium with the low-cost entry floor, the independent cost distribution gives a Borel function $e(P) \geq \rho$. Conditional on $X = x$, expected target payoff is

$$P(x) = t_0 + e(P(x))\{t_L - t_0 + \Delta_T \mu_X(x)\}. \tag{OA.7}$$

The function $g(P) = [P - t_0 - e(P)(t_L - t_0)]/[e(P)\Delta_T]$ is Borel on the realized price range because its denominator is strictly positive. It satisfies $\mu_X(X) = g(P(X))$ almost surely. Equation (OA.3) then gives $\mu_P = \mu_X$ almost surely. The entry function is consequently

$$e(\mu) = \rho + (1-\rho)\mathbf{1}\{g_L + \mu(g_H - g_L) \geq c_H\}. \tag{OA.8}$$

For construction, put $w_L = t_L - t_0 > 0$. For $\mu_2 > \mu_1$,

$$\begin{aligned} P(\mu_2) - P(\mu_1) &= e(\mu_2)\Delta_T(\mu_2 - \mu_1) + [e(\mu_2) - e(\mu_1)](w_L + \Delta_T \mu_1) \\ &\geq \rho \Delta_T (\mu_2 - \mu_1) > 0. \end{aligned} \tag{OA.9}$$

Thus $P(\mu)$ is strictly increasing even with an entry jump. Its inverse on its image is measurable: it is the restriction of a generalized inverse defined through infima of upper level sets. Price values inside a skipped interval are not generated by any flow and need not receive artificial posterior mass. By contrast, a flat posterior tail maps into an actual price atom and is conditioned on as an atom.

Conditional on fundamental quality, expected terminal value under a unilateral trading deviation is still $t_0 + e(P(x))(t_\theta - t_0)$. Subtracting (OA.7) gives $A_H = e\Delta_T(1 - \mu_X)$ and $A_L = e\Delta_T \mu_X$. Both residuals are bounded in $[\rho m \Delta_T, \Delta_T]$. The signs exclude incorrectly signed orders. This derivation accounts for the endogenous preparation response before taking the investor's expectation over shifted noise.

Next I justify every differentiation used for continuous deviations. Let $f \in W^{1,1}(\mathbb{R})$, with positive density and $|f'| \leq Lf$ almost everywhere, and let A be bounded, measurable, and nonnegative. For sign $\epsilon \in \{-1, 1\}$, define

$$F_\epsilon(s) = \int_{\mathbb{R}} f(x - \epsilon s) A(x)\, dx. \tag{OA.10}$$

For $s_1 < s_2$, the absolutely continuous representative of f satisfies

$$f(x - \epsilon s_2) - f(x - \epsilon s_1) = -\epsilon \int_{s_1}^{s_2} f'(x - \epsilon u)\, du. \tag{OA.11}$$

The absolute integral of the right-hand side after multiplication by A is at most $\|A\|_\infty (s_2 - s_1)\|f'\|_1$. Fubini therefore permits interchanging the integrals. This proves absolute continuity of F_ϵ and

$$F'_\epsilon(s) = -\epsilon \int f'(x - \epsilon s) A(x)\, dx, \qquad |F'_\epsilon(s)| \leq L F_\epsilon(s) \tag{OA.12}$$

almost everywhere. Translation continuity of f' in L^1 also makes the displayed derivative continuous. One way to see the needed translation continuity without an additional smoothness assumption is to approximate f' in L^1 by compactly supported continuous functions, for which it holds uniformly; the approximation error is unchanged by translation. The difference quotient of (OA.11) then converges in L^1 to $-\epsilon f'(\cdot - \epsilon s)$, identifying the classical derivative everywhere. The almost-everywhere statement is already sufficient for the proof.

The Laplace density is in $W^{1,1}$: it is continuous at zero, absolutely continuous on every bounded interval, and has integrable weak derivative $-\operatorname{sgn}(x)f(x)/b$. Its kink contributes no atom to the first weak derivative because the density itself is continuous. Hence $L = 1/b$. This argument differentiates the translated density, not the discontinuous entry policy.

For $U(s) = sF(s) - ks$, absolute continuity and the product rule imply

$$U'(s) = F(s) + sF'(s) - k \geq (1 - sL)F(s) - k. \tag{OA.13}$$

If the right-hand side is bounded below by a strictly positive constant on the allowed interval, integration shows that $U(s_2) > U(s_1)$ for every $s_2 > s_1$. This is global order optimality, not a local first-order condition.

A.4. Complete proof of Proposition 2

The proof first restricts the trading strategy in any candidate equilibrium, then constructs prices and entry that support it.

Weak incumbent. Proposition A.1 and (A1) give the entry floor under every conditional mixed order. Section A.3 gives the residual upper bound. For a correctly signed order of magnitude $s > 0$,

$$\mathbb{E}[\text{trading profit} \mid \Theta] \leq s\{\Delta_T(r_0) - k\} < 0. \tag{OA.14}$$

An incorrectly signed order has strictly negative gross expected profit. Zero gives zero. Thus each type has the unique best response zero, irrespective of the candidate equilibrium schedules. A mixed distribution placing positive probability on a nonzero order cannot be optimal: its payoff is an average of strictly negative payoffs and zero, and is strictly negative if that probability is positive. There is no issue from arbitrarily small trades; a nonpositive integrable random payoff that is strictly negative on a positive-probability set has a negative expectation.

Under zero orders, $X = Z$ is independent of quality. Set $e = \rho$ and $P = t_0 + \rho[(t_H + t_L)/2 - t_0]$. Condition (A2) excludes high-cost preparation at the prior, and (A1) admits low-cost preparation. Competitive pricing is correct and the previous payoff calculation validates zero. This constructs the unique trading and on-path preparation outcome.

Strong incumbent. Again consider any candidate equilibrium, including mixed orders. Section A.3 gives $F_\theta(s) \geq \rho m \Delta_T(r_1)$, and (OA.13) with $L = 1/b$ gives

$$U'_\theta(s) \geq (1 - 1/b)\rho m \Delta_T(r_1) - k > 0. \tag{OA.15}$$

Every smaller correctly signed magnitude is strictly inferior to the unit magnitude; incorrect signs are inferior to zero. Thus the only possible equilibrium orders are $q_H = 1, q_L = -1$.

Their conditional flow densities are $f(x-1)$ and $f(x+1)$. Bayes' rule gives the piecewise posterior

$$\mu(x) = \begin{cases} m, & x \leq -1, \\ (1 + e^{-2x/b})^{-1}, & -1 < x < 1, \\ M, & x \geq 1. \end{cases} \tag{OA.16}$$

Use the preparation rule (OA.8) and pricing (OA.7). Equation (OA.9) makes the price invertible with respect to this posterior, so the buyer optimizes using exactly the information it observes. The global bounds verify investor best responses, and the construction satisfies competitive pricing and truthful bidding. This proves existence and the uniqueness of the trading and on-path preparation outcome. Price functions can differ at unrealized price values or flow null sets without creating a different economic outcome.

Entry and allocation. Since $g_H > g_L$, expensive preparation requires $\mu \geq \tau$. The strict inequalities $B_{r_1}(1/2) < B_{r_0}(1/2) < c_H < B_{r_1}(M)$ imply $\tau \in (1/2, M)$. Solving (OA.16) gives $x^* = b\,\mathrm{logit}(\tau)/2 \in (0, 1)$. The Laplace survival function is

$$\overline{F}_Z(z) = \begin{cases} 1 - \frac{1}{2}e^{z/b}, & z < 0, \\ \frac{1}{2}e^{-z/b}, & z \geq 0. \end{cases} \tag{OA.17}$$

Hence $\alpha_H = \overline{F}_Z(x^* - 1)$ and $\alpha_L = \overline{F}_Z(x^* + 1)$ give the main-text formulas. Both are positive. State-specific entry is $e_\theta = \rho + (1 - \rho)\alpha_\theta$, total entry is $(e_H + e_L)/2$, and high-quality ownership is $e_H/2$ because $h > r$. Each comparison is strict.

Information ordering. A signal experiment is its pair of conditional distributions given quality. Applying a constant Markov kernel to the strong-economy price produces the weak-economy constant price experiment. Conversely, every state-independent kernel applied to a constant experiment has the same conditional law in both states. The strong experiment does not: on an interior price region its posterior is nonconstant, since the conditional flow likelihood ratio is nonconstant and the price reveals it. Thus no reverse garbling exists. Differences in mean payoffs across economies do not affect this comparison of signal experiments.

Nonemptiness. Section A.6 gives a direct primitive construction for every $h > \ell$. The specified benchmark is also strictly inside the condition region. All inequalities involve continuous primitive expressions away from the support boundaries, so their strict satisfaction persists on an open neighborhood.

A.5. Fixed experiments, threshold distinctions, and finite nonmonotonicity

For Proposition A.3, fix one joint distribution of (S, Θ) for both strengths, and keep the preparation cost C independent of (S, Θ) with its law also fixed as strength varies. Only r changes between the two economies; the signal experiment and the cost law do not. Cost independence is what makes the buyer's posterior $\mu(S) = \Pr(H \mid S)$ the right object: with a cost that carried information about quality, the comparison would need $\Pr(H \mid S, C)$ and an independence condition between the incumbent's value and the buyer's full information, and I do not claim that version. For $r_1 > r_0$, the difference in entry indicators is

$$\mathbf{1}\{C \leq B_{r_0}(\mu(S))\} - \mathbf{1}\{C \leq B_{r_1}(\mu(S))\} = \mathbf{1}\{B_{r_1}(\mu(S)) < C \leq B_{r_0}(\mu(S))\}. \tag{OA.18}$$

Taking expectations proves weak decrease and the exact strictness condition, without requiring an atomless cost distribution. This comparison holds within a fixed full-order regime because both conditional flow laws are unchanged when only r changes. It need not hold across different informative order profiles.

For Proposition A.4, define the maintained domain

$$\mathcal{D} = \{r \in (\ell, h) : c_L < B_r(m),\ B_r(1/2) < c_H\}. \tag{OA.19}$$

Under pooling, the public posterior is $1/2$ and entry is ρ. A unilateral correctly signed order sees a constant residual $\rho\Delta_T/2$ and earns $s(\rho\Delta_T/2 - k)$. Thus pooling is an equilibrium exactly when this coefficient is nonpositive. Equality supports pooling although the trader is indifferent among correctly signed sizes against that constant candidate schedule; it does not make every alternative profile an equilibrium with recomputed prices.

Solving $\Delta_T(r) = d$ gives $(r-\ell)^2 = 2rd$, whose roots are $\ell + d \pm \sqrt{d^2 + 2\ell d}$. The plus root is the one above ℓ. Writing $\mathfrak{r}(d) = \ell + d + \sqrt{d^2 + 2\ell d}$ gives the exact pooling-existence boundary $r_N = \mathfrak{r}(2k/\rho)$, the sufficient pooling-uniqueness boundary $\mathfrak{r}(k)$, and the sufficient full-order uniqueness boundary $r_U = \mathfrak{r}(k/[(1-1/b)\rho m])$. These boundaries are applied only within $\mathcal{D}$. A sufficient global bound does not identify the first existence of an informative equilibrium.

At $r > r_U$ in $\mathcal{D}$, the global derivative bound forces full orders. Expensive entry is feasible only if the attainable posterior reaches the required τ. The uniform mixed-strategy posterior bound makes $B_r(M) < c_H$ sufficient to exclude expensive entry in every equilibrium, regardless of the orders. The relevant equality is

$$Mr^2 - 2(Mh - c_H)r - [(1-M)\ell^2 - p^2] = 0. \tag{OA.20}$$

Because $B_r(M)$ is strictly decreasing on $r > \ell > p$, there is at most one root in the economic domain. The larger quadratic root gives that root when it lies in the domain; if both algebraic roots are positive, the other cannot lie in this strictly decreasing domain as a second crossing. Numerical specifications must verify the discriminant, the root residual, and the support and floor conditions rather than treating the formula as globally applicable.

For Laplace noise, the posterior M is attained on $x \geq 1$. At equality $B_r(M) = c_H$, entry-at-indifference admits the costly buyer on that complete tail. As $\tau \uparrow M$, the flow threshold approaches unity and

$$\alpha_H \to \frac{1}{2}, \qquad \alpha_L \to \frac{1}{2}e^{-2/b}, \qquad \mathsf{E} \to \rho + \frac{1-\rho}{4}(1 + e^{-2/b}). \tag{OA.21}$$

For any strict infeasibility beyond the crossing, costly entry is zero. An alternative tie rule changes the value at the equality itself, not the strict comparisons on either side. With logistic noise the upper posterior is approached only as $x \to \infty$; its survival probability tends to zero, and entry converges continuously to ρ.

For Proposition 2(iii), impose (A1) to (A3) at r_0, r_1, the low-cost floor at r_2, the strong global derivative bound at r_2, and $B_{r_2}(M) < c_H$. Parts (i) and (ii) of Proposition 2 give the first two unique outcomes; the bound and infeasibility condition give unique full orders but entry ρ at r_2. This proves a finite rise and fall without a selection or uniqueness assertion for strengths between them.

A useful full-profile existence test is sharper than the uniform uniqueness bound. At a fixed full-order candidate,

$$J = F_H(1) = F_L(1) = \Delta_T \int e(x)\frac{f(x-1)f(x+1)}{f(x-1)+f(x+1)}\,dx. \tag{OA.22}$$

The equality follows by substituting Bayes' rule into both residuals. The pointwise density-ratio bound implies $F_\theta(s) \geq e^{-(1-s)/b}J$. Since $(1-s/b)e^{-(1-s)/b}$ is decreasing on $[0,1]$, $k < (1-1/b)J$ makes every correctly signed marginal profit positive. It verifies existence of that profile, not uniqueness over other candidate schedules. The distinction is preserved in the numerical classifications.

A.6. Smooth noise, atomless costs, and primitive nonemptiness

For logistic noise,

$$f(z) = \frac{e^{-z/b}}{b(1+e^{-z/b})^2}, \qquad (\log f)'(z) = -\frac{1}{b}\tanh\left(\frac{z}{2b}\right). \tag{OA.23}$$

It is positive and smooth, and $|f'| \leq f/b$ implies $\|f'\|_1 \leq 1/b$. Therefore it belongs to $W^{1,1}$ and all posterior and global trading bounds apply. Under full orders the posterior log odds equal

$$\log\frac{\mu(x)}{1-\mu(x)} = 2\log\cosh\left(\frac{x+1}{2b}\right) - 2\log\cosh\left(\frac{x-1}{2b}\right). \tag{OA.24}$$

Their derivative equals

$$\frac{1}{b}\left[\tanh\left(\frac{x+1}{2b}\right) - \tanh\left(\frac{x-1}{2b}\right)\right] > 0. \tag{OA.25}$$

The limits are $-2/b$ and $2/b$, giving the open posterior range (m, M). Every strictly interior threshold is crossed with positive probability. To derive its inverse, set $y = e^{x/b}$, $A = e^{1/b}$, and $w = \sqrt{\tau/(1-\tau)}$. The square root of the likelihood ratio equals $(Ay+1)/(y+A)$. Solving $w = (Ay+1)/(y+A)$ gives $y = (Aw-1)/(A-w) > 0$. Thus $x^*_{\log} = b\log[(Aw-1)/(A-w)]$. Its conditional survival probabilities are $\alpha_H = [1+e^{(x^*_{\log}-1)/b}]^{-1}$ and $\alpha_L = [1+e^{(x^*_{\log}+1)/b}]^{-1}$. These formulas and $\mathsf{E} = \rho + (1-\rho)(\alpha_H+\alpha_L)/2$ supply the complete entry comparison. The trade and information-ordering arguments complete Proposition A.5.

For Proposition A.6, let $H_C = \rho H_L + (1-\rho)H_H$, with atomless component laws supported within $[c_L - \varepsilon_C, c_L + \varepsilon_C]$ and $[c_H - \varepsilon_C, c_H + \varepsilon_C]$. Retain (A3) and impose $c_L - \varepsilon_C \geq 0$, $c_L + \varepsilon_C < B_{r_1}(m)$, and

$$B_{r_0}(1/2) < c_H - \varepsilon_C < c_H + \varepsilon_C < B_{r_1}(M). \tag{OA.26}$$

Use either the Laplace or logistic noise law. The entry rule is $e(\mu) = H_C(B_r(\mu))$. It is continuous and nondecreasing, and $e \geq \rho$ under the low-cost support restriction. Equation (OA.9) remains strictly positive even if the component densities are discontinuous or vanish on subintervals. In the weak economy, $e(1/2) = \rho$. In the strong economy, a posterior sufficiently close to M satisfies $B_r(\mu) > c_H + \varepsilon_C$; an event with positive probability then gives $e = 1$. All global trading comparisons are unchanged. This proof requires an atomless distribution, not a differentiable density. Differentiability is imposed separately when differentiating a seller objective.

For primitive nonemptiness, fix $h > \ell > p > 0$, $b > 1$, and $0 < \rho < 1$. At the continuous limit $r = \ell$,

$$g_H - g_L = h - \ell > 0, \qquad B_\ell(M) - B_\ell(1/2) = (M - 1/2)(h-\ell) > 0. \tag{OA.27}$$

Choose $r_1 \in (\ell, h)$ close enough to ℓ that $B_{r_1}(M) > B_\ell(1/2)$. Because $B_r(1/2)$ decreases in r, every $r_0 \in (\ell, r_1)$ then satisfies $B_{r_0}(1/2) < B_{r_1}(M)$. Choose r_0 sufficiently near ℓ that $\Delta_T(r_0) < (1-1/b)\rho m \Delta_T(r_1)$. Select k strictly inside this interval and c_H strictly between the two profit bounds. The quantity $B_{r_1}(m)$ is positive, so choose c_L strictly between zero and that bound. The constructed $c_L < c_H$ follows because $B_{r_1}(m) < B_{r_0}(1/2)$, using $m < 1/2$ and the monotonicity in both arguments. All inequalities have strict slack and persist by continuity. This is a mathematical nonemptiness construction, not a statement about empirical effect size at every value ratio.

Two further stability results are recorded for completeness.

Lemma OA.1 (analytical). *Replace the trading cost by $K(s) = ks + \gamma_q s^2/2$, $\gamma_q \geq 0$. Under (A1) and (A2), $\Delta_T(r_0) < k$, and $k + \gamma_q < (1 - 1/b)\rho m \Delta_T(r_1)$, the unique trading outcomes and entry reversal in parts (i) and (ii) of Proposition 2 remain valid.*

Indeed $K(s) \geq ks$ preserves the weak-economy exclusion. At high strength, subtracting $K'(s) = k + \gamma_q s$ from the gross marginal profit leaves a strictly positive bound. The price construction, threshold, and information comparison are unchanged. This perturbation retains the position bound; it does not solve an unbounded-order game.

Lemma OA.2 (analytical). *For prior $\pi \in (0, 1)$, let $\mu_- = \text{logistic}(\text{logit}\,\pi - 2/b)$, $\mu_+ = \text{logistic}(\text{logit}\,\pi + 2/b)$, and $\omega = \min\{\mu_-, 1 - \mu_+\}$. Replace (A1)–(A3) by $c_L < B_{r_1}(\mu_-)$, $B_{r_0}(\pi) < c_H < B_{r_1}(\mu_+)$, and $\Delta_T(r_0) < k < (1 - 1/b)\rho\omega\Delta_T(r_1)$. Then the benchmark entry reversal and unique trading outcomes hold.*

Prior odds multiply the order-flow likelihood ratio. The uniform residual lower bound is therefore $\rho\omega\Delta_T$, while the upper bound is still Δ_T. Every inference and global deviation argument follows with these bounds. Under full Laplace orders, $x^* = b[\text{logit}\,\tau - \text{logit}\,\pi]/2$ lies in $(0, 1)$, and total entry is $\rho + (1 - \rho)[\pi\alpha_H + (1 - \pi)\alpha_L] > \rho$. Strict feasibility at the equal prior persists for nearby priors. Neither supplemental lemma is required to generate a reported numerical value in the present manuscript.

A.7. Full proof and probability formulas for Proposition A.7

The auction, value supports, order interval, noise law, and preparation-cost distribution are those in A.1. Only the information observations change. Let $a, d \in (1/2, 1)$ denote investor and buyer signal accuracy. Define

$$\begin{aligned}
\mu_- &= (1 - a) + (2a - 1)m, & \mu_+ &= (1 - a) + (2a - 1)M, \\
\phi_+(\mu) &= \frac{d\mu}{d\mu + (1 - d)(1 - \mu)}, & \phi_-(\mu) &= \frac{(1 - d)\mu}{(1 - d)\mu + d(1 - \mu)}, \\
w_H(r) &= t_H(r) - t_0(r), & w_L(r) &= t_L(r) - t_0(r).
\end{aligned} \tag{OA.28}$$

The complete sufficient conditions for Proposition A.7 are

$$\begin{aligned}
&c_L < B_{r_1}(\phi_-(\mu_-)), \\
&B_{r_0}(d) < c_H < B_{r_1}(\phi_+(\mu_+)), \\
&(2a - 1)\{\Delta_T(r_0) + (1 - \rho)(2d - 1)w_H(r_0)\} < k \\
&\qquad\qquad < (1 - 1/b)m(2a - 1)\rho\Delta_T(r_1).
\end{aligned} \tag{OA.29}$$

I prove unique zero orders and entry ρ at r_0, unique maximum correctly signed signal-contingent orders at r_1, and strictly higher entry at r_1. The result allows $d > a$ and arbitrary mixed orders conditional on the investor's signal.

I use T for the trader's signal, Y for the buyer's private signal, and Θ for fundamental quality. Conditional on H, the probabilities of $T = +$ and $Y = +$ are a and d; conditional on L they are $1 - a$ and $1 - d$. Conditional independence gives the joint signal law as the product of these probabilities within each fundamental state. Equal fundamental priors imply equal marginal probabilities for each trader signal.

For arbitrary signal-contingent mixed orders, let $a_+(x) = \int f(x-q)d\sigma_+(q)$ and $a_-(x) = \int f(x-q)d\sigma_-(q)$. Then

$$\lambda_X = \frac{a_+}{a_+ + a_-} \in [m, M], \qquad \mu_X = \frac{aa_+ + (1-a)a_-}{a_+ + a_-} = (1-a) + (2a-1)\lambda_X. \tag{OA.30}$$

Since X depends on quality only through T, Θ and X are conditionally independent given T. Also Y is independent of X conditional on Θ. The latter remains true after applying a measurable price function, and implies

$$\Pr(H \mid P, Y = y) = \phi_y(\mu_P). \tag{OA.31}$$

The maps ϕ_+, ϕ_- are continuous and strictly increasing on $(0, 1)$; the favorable private signal gives the larger posterior. The public posterior based on price inherits $[\mu_-, \mu_+]$ by conditional expectation. The lowest possible joint posterior is consequently $\phi_-(\mu_-)$. The low-cost restriction in (OA.29) and declining gross profits in r make low-cost preparation optimal in both economies under any candidate equilibrium.

Define high-cost indicators $I_y(P) = \mathbf{1}\{B_r(\phi_y(\mu_P)) \geq c_H\}$. They obey $I_+ \geq I_-$. Since Y's law conditional on quality does not change with X or the trader's signal, true-state entry conditional on P is

$$\begin{aligned} e_H(P) &= \rho + (1-\rho)\{dI_+(P) + (1-d)I_-(P)\}, \\ e_L(P) &= \rho + (1-\rho)\{(1-d)I_+(P) + dI_-(P)\}. \end{aligned} \tag{OA.32}$$

In particular,

$$e_H - e_L = (1-\rho)(2d-1)(I_+ - I_-), \qquad e_H \geq e_L \geq \rho. \tag{OA.33}$$

Let $D = e_H w_H - e_L w_L$. As $w_H - w_L = \Delta_T$ and $w_L > 0$,

$$D = e_L \Delta_T + (e_H - e_L)w_H \in [\rho\Delta_T,\ \Delta_T + (1-\rho)(2d-1)w_H]. \tag{OA.34}$$

Conditional pricing is $P = t_0 + e_L(P)w_L + \mu_X D(P)$. Both entry probabilities are functions of price, even before proving sufficiency. The positive coefficient $D(P)$ therefore makes μ_X measurable with respect to price by the explicit inversion. The tower property gives $\mu_P = \mu_X$ and hence also reveals λ_X through the affine inverse in (OA.30).

For construction, set the entry indicators using $\phi_y(\mu)$ and write

$$P(\mu) = t_0 + \mu e_H(\mu)w_H + (1-\mu)e_L(\mu)w_L. \tag{OA.35}$$

Both entry probabilities are nondecreasing in μ. For $\mu_2 > \mu_1$, decompose its change into a belief change at the old entry probabilities and an entry change evaluated at the new belief:

$$\begin{aligned} P(\mu_2) - P(\mu_1) = (\mu_2 - \mu_1)D(\mu_1) &+ \mu_2[e_H(\mu_2) - e_H(\mu_1)]w_H \\ &+ (1-\mu_2)[e_L(\mu_2) - e_L(\mu_1)]w_L > 0. \end{aligned} \tag{OA.36}$$

Thus the same measurable-inverse construction works even when the entry indicators jump simultaneously or one never changes. This argument does not rely on differentiating an indicator.

Conditional on $T = +$, the trader assigns quality probability a even after its own randomized order and the resulting flow are specified. Conditional on $T = -$, it assigns $1 - a$. This remains valid under unilateral orders because the deviation introduces no additional quality information conditional on the signal. Its expected terminal target payoff at a flow x is therefore $t_0 + e_L(P(x))w_L + aD(P(x))$ for $T = +$, and the corresponding expression with $1 - a$ for $T = -$. Subtracting (OA.35) and using (OA.30) gives

$$A_+ = (2a - 1)(1 - \lambda_X)D, \qquad A_- = (2a - 1)\lambda_X D. \tag{OA.37}$$

The residuals have strict signs and obey

$$m(2a - 1)\rho\Delta_T \leq A_+, A_- \leq (2a - 1)\{\Delta_T + (1 - \rho)(2d - 1)w_H\} =: \overline{A}. \tag{OA.38}$$

The weak restriction in (OA.29) makes any correctly signed order of magnitude $s > 0$ earn at most $s(\overline{A} - k) < 0$. Hence zero is necessary under all candidate equilibria. With zero orders the public posterior is the prior, the buyer's favorable posterior is d, and the expensive-entry exclusion in (OA.29) gives entry ρ. Constant competitive pricing constructs the weak equilibrium.

At high strength, the convolution regularity applies to either bounded signal residual and gives $U'(s) \geq (1 - 1/b)m(2a - 1)\rho\Delta_T - k > 0$. Thus the necessary and unique trading profile is full correctly signed orders by trader signal. Bayes' rule, (OA.35), and optimal preparation construct the equilibrium. The strict high-cost upper inequality ensures that $Y = +$ and public beliefs sufficiently near μ_+ induce preparation. Full Laplace orders attain those beliefs on a positive-probability event; the conditional probability of $Y = +$ is positive. Total entry strictly exceeds ρ. The proof allows $d > a$ throughout.

For numerical implementation with either private signal, let $\tau = (c_H - g_L)/(g_H - g_L)$ and solve $\phi_y(\mu) = \tau$:

$$\begin{aligned} \mu_{\text{req},+} &= \frac{\tau(1 - d)}{d(1 - \tau) + \tau(1 - d)}, \\ \mu_{\text{req},-} &= \frac{\tau d}{(1 - d)(1 - \tau) + \tau d}, \\ \lambda_{\text{req},y} &= \frac{\mu_{\text{req},y} - (1 - a)}{2a - 1}. \end{aligned} \tag{OA.39}$$

These formulas presume an interior fundamental threshold. Thresholds outside $[0, 1]$ imply always or never entry and are handled directly by the profit comparison. For a requirement strictly between m and M, full Laplace orders give $x_y^* = b\,\text{logit}(\lambda_{\text{req},y})/2$. Requirements below the lower bound or above the upper bound imply always or never entry; equality at a plateau follows the tie rule. The buyer's unfavorable signal need not be assumed to preclude high-cost entry in a parameter sweep.

The fundamental-conditioned flow densities are

$$f_H^X(x) = af(x - 1) + (1 - a)f(x + 1), \qquad f_L^X(x) = (1 - a)f(x - 1) + af(x + 1). \tag{OA.40}$$

Let $Q_{\theta,y} = \Pr(X \geq x_y^* \mid \Theta = \theta)$, with the always/never cases evaluated accordingly. Conditional entry is

$$\begin{aligned} \bar{e}_H &= \rho + (1-\rho)[dQ_{H,+} + (1-d)Q_{H,-}], \\ \bar{e}_L &= \rho + (1-\rho)[(1-d)Q_{L,+} + dQ_{L,-}]. \end{aligned} \tag{OA.41}$$

Thus $\mathsf{E} = (\bar{e}_H + \bar{e}_L)/2$, $\mathsf{O}_H = \bar{e}_H/2$, and $\mathcal{R}_T = t_0 + (\bar{e}_H w_H + \bar{e}_L w_L)/2$. Independently integrating the joint law of (Θ, T, Y, Z, C) must reproduce these expressions. Using the marginal law of Y instead of its conditional law given quality would give the wrong entry and pricing objects.

A.8. Verifiable-value bargaining

The bargaining comparison uses a different acquisition institution, specified here. After diligence, all buyer values are verifiable. The seller can implement a sale to the runner-up at that buyer's value, with acceptance at zero buyer surplus; therefore the disagreement payoff in negotiating with the best buyer is the next-best value z. The best buyer has value $V \geq z$ and disagreement payoff zero. For $0 < \eta < 1$, its Nash transfer solves

$$\max_{P \in [z,V]} (P-z)^{\eta}(V-P)^{1-\eta}. \tag{OA.42}$$

For $V > z$, strict concavity of the log objective gives the unique transfer $P = (1-\eta)z + \eta V$. When $V = z$, feasible transfers collapse to that common value. At $\eta = 0$ I use the continuous limiting transfer z. With no challenger, the same rule has fallback zero and pays the seller ηR. Thus no-entry proceeds are specified by the institution as well as entry proceeds.

The highest-value buyer receives the target. For $R \leq \ell$, the high and low target payments are $(1-\eta)R + \eta h$ and $(1-\eta)R + \eta\ell$. For $R > \ell$, they are $(1-\eta)R + \eta h$ and $(1-\eta)\ell + \eta R$. Their difference is

$$\eta(h-\ell) + (1-2\eta)(R-\ell)_+. \tag{OA.43}$$

A challenger receives $(1-\eta)(\theta - R)$ when $\theta \geq R$ and zero otherwise. Taking expectations gives Proposition A.8 and applying FOSD yields its sign comparisons. For a uniform incumbent,

$$\begin{aligned} \Delta_\eta &= \eta(h-\ell) + (1-2\eta)\frac{(r-\ell)^2}{2r}, \\ G_{H,\eta} &= (1-\eta)(h - r/2), \\ G_{L,\eta} &= (1-\eta)\frac{\ell^2}{2r}. \end{aligned} \tag{OA.44}$$

The transfer is always feasible between fallback and winning value. For $\eta > 1/2$ the competition derivative of the spread is negative; the spread itself remains positive because the high challenger has the higher value. The endpoint $\eta = 1$ extinguishes buyer rents and is not used to infer a positive-cost entry equilibrium. The comparison does not solve a private-value bargaining or first-price auction game; verifiability and fallback enforceability are part of the stated institution.

A.9. Matched welfare and the external-dividend diagnostic

I couple the feedback and price-hidden economies at $r = r_1$ using the same (Θ, R, C, Z) and full informed orders. The strong trading bound applies in the price-hidden economy with entry probability ρ, so this coupling uses equilibrium behavior in both environments. The high-cost type remains out when price information is unavailable because $B_{r_1}(1/2) < c_H$; all low-cost realizations enter in both environments.

Without challenger entry, allocation value is $R\mathbf{1}\{R \geq p\}$. With entry it is $\max(R, \theta)$ because $\theta > p$. Splitting at $R = p$ and $R = \theta$ gives

$$\max(R, \theta) - R\mathbf{1}\{R \geq p\} = (\theta - \max(p, R))_+ + p\mathbf{1}\{R < p\}. \tag{OA.45}$$

If $R < p$, both sides equal θ; if $R \geq p$, both sides equal $(\theta - R)_+$. Conditional on the signal and preparation cost, R retains its independent uniform law, so incremental expected allocation value net of cost is

$$B_r(\mu_X) - C + p\Pr(R < p) = B_r(\mu_X) - C + \frac{p^2}{r}. \tag{OA.46}$$

Additional entry occurs only when its first two terms sum to a nonnegative value. The final term is strictly positive. The statement is about conditional expectations: each additional preparation decision contributes at least p^2/r in conditional expectation, given price information and preparation cost. A realized low-value preparation need not generate a positive realized net gain, since $(\ell - \max(p, R))_+$ can fall short of the cost paid. For atomic costs, integrate against the marginal full-order flow density $g(x) = [f(x-1) + f(x+1)]/2$ to obtain

$$\Delta\mathcal{W} = (1-\rho)\int g(x)\left[B_r(\mu_X(x)) - c_H + \frac{p^2}{r}\right]\mathbf{1}\{B_r(\mu_X(x)) \geq c_H\}\,dx > 0. \tag{OA.47}$$

For an atomless high-cost component H_H, replace the integrand by

$$(1-\rho)\int\left[B_r(\mu_X(x)) - c + \frac{p^2}{r}\right]\mathbf{1}\{c \leq B_r(\mu_X(x))\}\,H_H(dc). \tag{OA.48}$$

The strict-support conditions ensure positive additional entry. Fubini is valid because values and the specified cost supports are bounded. The realized trading cost is the same in both environments; any resource interpretation of that cost therefore cancels. Transfers between bidders, shareholders, market makers, and noise traders are not additional allocation surplus.

Target proceeds increase by averaging $t_\theta - t_0 > 0$ over each state's additional-entry probability. For atomic costs the difference is $(1-\rho)\{\alpha_H(t_H - t_0) + \alpha_L(t_L - t_0)\}/2 > 0$; for atomless costs integrate the corresponding state-specific tail probabilities over the high-cost component. As an independent numerical check, compute total expected allocation minus the full expected preparation cost in each environment:

$$\mathcal{W} = \frac{r^2 - p^2}{2r} + \frac{1}{2}\sum_\theta \bar{e}_\theta\left(g_\theta + \frac{p^2}{r}\right) - \mathbb{E}[C\mathbf{1}\{\text{entry}\}], \tag{OA.49}$$

valid in the benchmark support region. Subtracting the price-hidden value must agree with (OA.47) or (OA.48).

For the matched-level diagnostic, attach a deterministic payoff $D_0 = \mathcal{R}_T^{\text{feedback}} - \mathcal{R}_T^{\text{hidden}}$ to the traded claim in the price-hidden economy. Its competitive price shifts by D_0 and $V_T - P$ is unchanged. The buyer's acquisition profits and information set do not change. This matches mean prices without manufacturing information. The external payoff is not counted as a real surplus gain and is not an available reserve or acquisition payment.

A.10. Continuous acquisition values and sale-design regularity

For realized challenger value v, define pointwise seller revenue

$$T(R, v, p) = \begin{cases} p\mathbf{1}\{R \geq p\}, & v < p, \\ \max\{p, \min(R, v)\}, & v \geq p, \end{cases} \qquad G(R, v, p) = (v - \max(p, R))_+. \tag{OA.50}$$

These formulas include no sale when both values fail the reserve. They are bounded, Borel, and can be integrated first over R and then over any conditional challenger distribution. The latter order is convenient when a reserve cuts through a value band. At a value atom equal to the reserve, the bid meets the reserve even when its acquisition rent is zero; the cost of preparation was paid before the realized value was learned.

For uniform R and a fixed v, a useful complete kernel is

$$\begin{aligned} &v < p: && t_v = t_0, \quad g_v = 0, \\ &p \leq v < r: && t_v = v - \frac{v^2 - p^2}{2r}, \quad g_v = \frac{(v-p)^2}{2r} + \frac{p(v-p)}{r}, \\ &p \leq r \leq v: && t_v = \frac{r}{2} + \frac{p^2}{2r}, \quad g_v = v - t_v, \\ &r < p \leq v: && t_v = p, \quad g_v = v - p, \end{aligned} \tag{OA.51}$$

with $t_0 = p(1 - p/r)$ for $p \leq r$ and $t_0 = 0$ for $p > r$. The third line presumes $p \leq r$. The kernels agree at $v = r$ and $p = r$. At $p = v$, the challenger remains admissible at zero acquisition rent; moving the reserve above a value atom can therefore cause a jump in target proceeds. In the second line $g_v = (v^2 - p^2)/(2r)$. Direct integration of (OA.50), rather than the shortened kernels alone, is the independent numerical oracle.

For the atomless-value diagnostic, a binary class S is equally likely. Conditional on $S = L$, V is uniform on $[\ell - \varepsilon_V, \ell + \varepsilon_V]$; conditional on $S = H$, it is uniform on $[h - \varepsilon_V, h + \varepsilon_V]$. The investor observes S, not the exact V. Preparation reveals V. The independent incumbent remains uniform. Conditional class values have equal width but different means. This is a binary-information economy with atomless acquisition values, not a perfectly informed trader observing a continuous value.

When $p < \ell - \varepsilon_V < \ell + \varepsilon_V < r < h - \varepsilon_V$,

$$\begin{aligned} t_L &= \ell - \frac{\ell^2 + \varepsilon_V^2/3 - p^2}{2r}, \qquad g_L = \frac{\ell^2 + \varepsilon_V^2/3 - p^2}{2r}, \\ t_H &= \frac{r}{2} + \frac{p^2}{2r}, \qquad g_H = h - t_H, \qquad \Delta_T = \frac{(r-\ell)^2}{2r} + \frac{\varepsilon_V^2}{6r}. \end{aligned} \tag{OA.52}$$

When the reserve is above the entire low band but below r, $t_L = t_0$, $g_L = 0$, and the high-class formulas are unchanged. Conditional on class and flow, exact V retains its within-class distribution, so the benchmark price and residual proof applies to these class-averaged payoffs whenever $\Delta_T > 0$ and a positive entry floor holds. In this region $t_L = t_0$ is permitted: the price construction remains strictly increasing on the bounded posterior interval because $\Delta_T \mu > 0$.

For a general reserve, the object the seller faces after choosing p is a complete continuation

$$\sigma = (\sigma_H, \sigma_L;\ P(\cdot);\ \mu_P;\ \text{preparation rule}) \in \mathcal{E}(p, r), \tag{OA.53}$$

the conditional order distributions, the price mapping on flow, the buyer's beliefs at every realized price including price atoms, and the preparation rule at every price and cost. Its beliefs are the public price posterior, not necessarily the raw order-flow posterior. If entry vanishes on a set of flows, the inversion denominator vanishes there, those flows share the no-entry price, and the buyer conditions on the whole preimage of that price. The construction used in C.6, which prices every positive-entry flow by the benchmark inversion and pools exactly the zero-entry preimage at t_0, is one admissible member of $\mathcal{E}(p, r)$ when it validates. It is not the only one. Section A.11 exhibits, in the reserve game itself, a one-parameter family of continuations with the same orders and different pools, each of which validates. Two continuations with the same orders but different price mappings are different elements of $\mathcal{E}(p, r)$, and the numerical record identifies a continuation by the whole object (OA.53). If $\Delta_T = 0$ or all entry is impossible, the informative-price argument cannot be imported. Section C.6 therefore requires direct pricing, posterior, and entry verification at such reserves.

For the local decomposition in the paper, impose a neighborhood I contained strictly in $0 < p < \ell$, a pure strategy branch $q_\theta(p)$ continuously differentiable on I, and an atomless preparation-cost CDF H_C continuously differentiable on the compact profit range. Suppose $\mu_\theta(z; p)$ is differentiable in p for almost every z and there is an integrable function $J(z)$ dominating

$$f(z)\left|h_C(B_{p,r}(\mu_\theta))\left[-p/r + (g_H - g_L)\partial_p\mu_\theta\right]\right| \tag{OA.54}$$

uniformly on a compact neighborhood of the reserve. These assumptions are sufficient, not an assertion that every continuation has them. For positive smooth noise with bounded log-density derivative and a bounded derivative of each pure order, posterior derivatives can be bounded directly by likelihood-ratio differentiation; bounded h_C then supplies domination by a constant times f. For Laplace noise, the same calculation holds away from the finite shifted kinks, a null set in z, and the bounded log slope again supplies domination.

For $\bar{e}_\theta(p) = \int f(z)H_C(B_{p,r}(\mu_\theta(z; p)))\, dz$, dominated differentiation gives

$$\frac{d\bar{e}_\theta}{dp} = \int f(z)h_C(B_{p,r}(\mu_\theta))\left[-\frac{p}{r} + (g_H - g_L)\frac{\partial\mu_\theta}{\partial p}\right] dz. \tag{OA.55}$$

With $\mathsf{E} = (\bar{e}_H + \bar{e}_L)/2$ and $\mathcal{R}_T = t_0 + \sum_\theta \bar{e}_\theta(t_\theta - t_0)/2$, the product rule gives

$$\frac{d\mathcal{R}_T}{dp} = (1 - \mathsf{E})\left(1 - \frac{2p}{r}\right) + \mathsf{E}\frac{p}{r} + \frac{1}{2}\sum_\theta \frac{d\bar{e}_\theta}{dp}(t_\theta - t_0). \tag{OA.56}$$

The calculation uses $\partial_p t_0 = 1 - 2p/r$, $\partial_p t_H = \partial_p t_L = p/r$, and $\partial_p B|_\mu = -p/r$. At an atomic preparation threshold or a nonregular continuation, this differentiation argument is unavailable. The level objective and full conditional laws remain the correct objects.

A seller equilibrium requires a feasible continuation selection after each reserve and a maximizing reserve under that selection. The numerical envelope of found continuations is an exploration of this problem, not a proof of existence, measurable selection, global attainment, or optimality. These are the explicit open parts of the sale-design result.

A.11. Price pooling when preparation can vanish

When preparation can vanish, identical investor orders can support different price pools and different preparation outcomes. This subsection proves the explicit family in Proposition A.10. Dow et al. (2017)

discuss continuations distinguished by their pooling regions following their Lemma 1. The reserve-game example makes the price rule part of the continuation that the seller must face.

Inputs and institution. Take the benchmark declaration of C.0 with $r = r_0$ and change only the reserve, to the value declared as `pool_reserve` in the registry:

```
h = 10; ell = 1; r = 1.2; p = 7;
rho = 0.25; c_L = 1; c_H = 6;
b = 2; k = 0.02; prior_H = 0.5;
q_H = 1; q_L = -1;
entry_at_indifference = yes;
bid_equal_reserve_is_admissible = yes.
```

The reserve exceeds incumbent support and excludes the low-value challenger, so this is the expanded reserve domain of A.10, not the benchmark support $p < \ell$. From (OA.51) with $r < p \le v$ for $v = h$ and $v < p$ for $v = \ell$,

$$t_0 = t_L = g_L = 0, \qquad t_H = 7, \qquad g_H = 3, \qquad \Delta_T = 7.$$

A payoff routine written for $p < \ell$ must not be applied here; C.6b evaluates these objects in exact rational arithmetic and by direct integration of (OA.50).

Densities and the price family. With $b = 2$ and full orders,

$$\begin{aligned} f(z) &= \frac{1}{4}e^{-|z|/2}, \qquad a_H(x) = f(x-1), \qquad a_L(x) = f(x+1), \\ \mu_X(x) &= \left[1 + \exp\left\{-\frac{|x+1| - |x-1|}{2}\right\}\right]^{-1}, \end{aligned} \tag{OA.57}$$

which is (OA.16) at $b = 2$: the plateau $m = (1+e)^{-1}$ on $x \le -1$, logistic(x) on $(-1, 1)$, and $M = 1 - m$ on $x \ge 1$. For each cutoff $c \in [-\log 2, 0]$ define the price rule

$$P_c(x) = \begin{cases} 0, & x < c, \\ \frac{7}{4}\mu_X(x), & x \ge c, \end{cases} \tag{OA.58}$$

with the preparation policy: at price zero neither cost type prepares; at every positive price the low-cost type prepares and the high-cost type does not. The upper plateau $x \ge 1$ maps to the price $\frac{7}{4}M$, a positive-price atom with probability $\frac{1}{2}[S_Z(0) + S_Z(2)]$; the buyer conditions on it as an atom, and its posterior is M because the plateau is a level set of μ_X.

Noise CDF versus survival. I write F_Z for the noise CDF and $S_Z = 1 - F_Z$ for its survival function, and keep the two apart throughout, because the pooled posterior uses the CDF and the outcome quantities use the survival function:

$$F_Z(z) = \begin{cases} \frac{1}{2}e^{z/2}, & z \le 0, \\ 1 - \frac{1}{2}e^{-z/2}, & z > 0, \end{cases} \qquad S_Z(z) = 1 - F_Z(z). \tag{OA.59}$$

Equation (OA.17) is S_Z at general b.

Buyer at price zero. The zero price is one atom whose preimage is the whole pool $\{X < c\}$. Its posterior follows (OA.79) with that preimage:

$$\bar{\mu}_c = \Pr(H \mid X < c) = \frac{\int_{-\infty}^{c} a_H}{\int_{-\infty}^{c} a_H + \int_{-\infty}^{c} a_L} = \frac{F_Z(c-1)}{F_Z(c-1) + F_Z(c+1)}. \tag{OA.60}$$

The likelihood ratio a_H/a_L is nondecreasing, so $\bar{\mu}_c$ is nondecreasing in c and, over the declared family, $\bar{\mu}_c \le \bar{\mu}_0 = \frac{1}{2}e^{-1/2} < \frac{1}{3}$. Gross profit at the pool is $B(\bar{\mu}_c) = g_H\bar{\mu}_c = 3\bar{\mu}_c < 1 = c_L$, so not preparing is strictly optimal at price zero for both cost types. The point that matters for the numerical layer is the case $c = 0$. There the raw flow posterior $\mu_X(x)$ exceeds $\frac{1}{3}$ on $(-\log 2, 0)$, and a buyer who observed those flows one by one would prepare at low cost. Those flows are not individually observed. They all produce price zero, the buyer's information set at price zero is the pool, and the pooled posterior governs. A candidate that lets the buyer act on μ_X inside the pool is not measurable with respect to the price and is rejected for that reason, not for a numerical breach. C.6b carries this as a negative control.

Buyer at positive prices. On $x \ge c$ the price $\frac{7}{4}\mu_X(x)$ is strictly increasing in μ_X, so the buyer recovers μ_X from the price on $(-1, 1)$ and the plateau atom on $x \ge 1$ carries posterior M. Since $c \ge -\log 2$ and μ_X is nondecreasing, $\mu_X(x) \ge \mu_X(-\log 2) = \frac{1}{3}$ on the positive-price region, so $3\mu_X \ge c_L$ and low-cost preparation is optimal under the equality convention, with equality at $c = -\log 2$ exactly at the cutoff. High-cost preparation is impossible at any price: gross profit is at most $g_H = 3 < 6 = c_H$. The zero-price and positive-price regions do not collide, since $\frac{7}{4}\mu_X \ge \frac{7}{12} > 0$ on $x \ge c$.

Conditional pricing. For $x < c$ nobody is prepared, neither the incumbent nor a low-value challenger can meet the reserve, and the terminal value is zero. For $x \ge c$ preparation occurs with probability ρ, and a sale, at exactly p, occurs only when the prepared challenger has value h:

$$\mathbb{E}[V_T \mid X = x] = \rho\, p\, \mu_X(x) = \tfrac{7}{4}\mu_X(x) \quad (x \ge c), \qquad \mathbb{E}[V_T \mid X = x] = 0 \quad (x < c). \tag{OA.61}$$

So P_c is the competitive price pointwise in flow on both regions. This is the test. Average price equal to average payoff is implied by it, and a wrong pointwise schedule can pass the average check, so C.6b verifies (OA.61) on a flow mesh that includes the cutoff from both sides and every breakpoint, and treats the average identity as a separate necessary check.

Global investor validation. Against a fixed member of the family the residuals of A.3, after the preparation response, are

$$A_H(x) = \rho p[1 - \mu_X(x)]\mathbf{1}\{x \ge c\}, \qquad A_L(x) = \rho p\, \mu_X(x)\mathbf{1}\{x \ge c\}. \tag{OA.62}$$

They vanish inside the pool, where the price carries no information and no entry follows, and are nonnegative elsewhere; wrong-signed orders are therefore dominated by zero. Let $J_c = F_H(1) = F_L(1)$ as in (OA.22). The equality holds pointwise inside the integral by Bayes' rule, since $f(x-1)[1-\mu_X(x)] = f(x+1)\mu_X(x)$. The Laplace density-ratio bound gives $F_\theta(s) \ge e^{-(1-s)/b}J_c$ and $|F'_\theta(s)| \le F_\theta(s)/b$ by (OA.12), hence by (OA.13)

$$U'_\theta(s) \ge \left(1 - \frac{1}{b}\right) J_c - k$$

for every $s \in [0,1]$. Every cutoff in the family satisfies $c \le 0 < 1$, so the tail $x \ge 1$ is active with entry ρ under each member. On that tail $1 - \mu_X = m = (1+e)^{-1} > \frac{1}{4}$, and its probability in state H is $S_Z(0) = \frac{1}{2}$. Dropping the rest of the active region,

$$J_c \ge \frac{\rho p\, m}{2} > \frac{7}{32}, \qquad U'_\theta(s) > \frac{7}{64} - \frac{1}{50} = \frac{143}{1600} > 0. \tag{OA.63}$$

The rational bound (OA.63) is the proof, common to every member of the family and to both types. A positive derivative found on a finite grid is not the proof; it is an implementation check, and C.6b keeps the two apart by writing the rational bound and the mesh minimum in separate columns. Full correctly signed orders are globally optimal against every price rule in the family. This is analytical existence for the whole family. It says nothing about whether other continuations exist in $\mathcal{E}(7, r_0)$.

Proposition OA.3 (price pooling at fixed orders; analytical). *For every $c \in [-\log 2, 0]$, the price rule* (OA.58) *with the stated preparation policy and full orders* $(1, -1)$ *is a continuation equilibrium in* $\mathcal{E}(7, r_0)$*: prices are competitive pointwise in flow, beliefs at the zero-price atom are* (OA.60)*, the buyer's preparation decision is optimal at every price information set, and full correctly signed orders are globally optimal for both investor types against the fixed price and preparation schedule. The orders are the same in all members of the family; prices, beliefs, and preparation differ.*

The proof is the four validations above. The statement does not establish uniqueness of this family within $\mathcal{E}(7, r_0)$, and it is confined to lower-interval pools $\{X < c\}$; other measurable pool sets are not searched, and the numerical record says so through `pricing_family_coverage` and `exhaustive_pricing_search = false`. On the benchmark support the family cannot arise: there $c_L < B_r(m)$ gives entry at least ρ at every flow and Proposition A.2 ties the price experiment to the orders.

Output quantities. With S_Z from (OA.59),

$$\begin{aligned} \Pr(P = 0) &= \frac{F_Z(c-1) + F_Z(c+1)}{2}, \\ \mathsf{E}_c &= \frac{\rho}{2}\{S_Z(c-1) + S_Z(c+1)\}, \qquad \Pr(\text{sale}) = \frac{\rho}{2} S_Z(c-1), \\ \mathcal{R}_{T,c} &= \frac{\rho p}{2} S_Z(c-1), \qquad \Pr(\text{two admissible bidders}) = 0. \end{aligned} \tag{OA.64}$$

The admissible-challenger probability equals the sale probability, because only a prepared high-value challenger can meet the reserve and the incumbent never does. Both E_c and $\mathcal{R}_{T,c}$ are strictly decreasing in c. A higher cutoff pools more flows into the uninformative price and suppresses preparation at flows that would have justified it if observed separately. C.6b recomputes (OA.64) in fifty-digit arithmetic and by independent conditional integration, and reports the endpoints $c = -\log 2$ and $c = 0$ through the registry keys `pool_posterior_cutoff_*`, `pool_entry_cutoff_*`, and `pool_revenue_cutoff_*`.

The seventeen-cutoff grid and the negative controls. The declared grid is $c_j = -\log 2\,(1 - j/16)$ for $j = 0, \ldots, 16$, with the endpoints included. Every grid member must pass the price, buyer, and global-investor validations, and preparation and revenue must decrease along the grid. That ordered check is a check of the implementation of (OA.64), not a continuum theorem. Three controls must fail, each for its own economic reason. At $c = -1$ some positive prices have $3\mu_X < 1$, so the prescribed low-cost preparation is not optimal there. At $c = 1$ the pooled zero-price posterior $\bar{\mu}_1$ makes low-cost preparation profitable, so the prescribed nonpreparation fails. Inside the $c = 0$ pool, a buyer who acts on the raw posterior is rejected as inconsistent with the observed-price information structure, whatever its arithmetic. A fourth control is about identity rather than validation: a deduplication key built from orders alone merges the two valid endpoint continuations, which have the same orders and different atoms, and the C.6b ledger must show that merge so that the key is never used. Failed controls are stored as expected failures with their reason, never as accepted rows.

B. Computer-assisted equilibrium certificates

The certificates in Proposition 3 establish an exact equilibrium inside each reported bracket. Outward interval arithmetic encloses a marginal-profit root for the low-value investor and verifies global optimality for the high-value investor. The argument combines analytical regularity and deviation bounds with finite interval calculations.

B.1. Objects to be certified

Fix the benchmark primitives and a strength r. Candidate orders are $(q_H, q_L) = (1, -v)$ with $v \in (0, 1)$. Let $\tau = (c_H - g_L)/(g_H - g_L)$, $c = (1 - v)/2$, and

$$\mu_v(x) = \text{logistic}\left(\frac{|x+v| - |x-1|}{b}\right), \quad x^*(r, v) = \frac{b \operatorname{logit} \tau + 1 - v}{2}. \tag{OA.65}$$

The certificate region has $1/2 < \tau < M_v$, where $M_v = \text{logistic}((1 + v)/b)$. Thus x^* is strictly between $-v$ and 1. Set $e(x) = \rho + (1 - \rho)\mathbf{1}\{x \geq x^*\}$ and define $A_H = e\Delta_T(1 - \mu_v)$, $A_L = e\Delta_T\mu_v$.

For a signed direction $\epsilon \in \{-1, 1\}$ and magnitude $s \in [0, 1]$, the investor's per-unit convolution and marginal profit are

$$F_\epsilon(s) = \int f(x - \epsilon s)A_\epsilon(x)\, dx, \qquad U'_\epsilon(s) = F_\epsilon(s) + sF'_\epsilon(s) - k, \tag{OA.66}$$

where $A_+ = A_H$ and $A_- = A_L$. In the root condition $\Psi(r, v) = U'_-(v; 1, -v)$, the derivative is with respect to the deviating magnitude s, holding the candidate schedule fixed. The outer change of v recomputes that schedule. Confusing these derivatives would solve the wrong equilibrium condition.

B.2. Global low-type optimality using a Stieltjes measure

The Laplace location family under $1 > -v$ gives a nondecreasing posterior, strictly increasing on $(-v, 1)$. Because $e(\mu)$ is nondecreasing and bounded below by ρ, the function A_L is bounded, nondecreasing, and nonconstant. Its right-continuous representative defines a finite nonzero positive measure dA_L by $dA_L((x, y]) = A_L(y) - A_L(x)$.

Integration by parts on finite intervals, followed by passage to the infinite endpoints, gives

$$F'_L(s) = \int f'(x + s)A_L(x)\, dx = -\int f(x + s)\, dA_L(x). \tag{OA.67}$$

The boundary term vanishes because A_L is bounded and the Laplace density tends to zero in both tails. The measure integral is finite because f is bounded and dA_L has finite mass. It is strictly positive before the minus sign because $f > 0$ and the measure is nonzero. An entry jump contributes its positive atom to this measure.

The function f is globally Lipschitz. Therefore $s \mapsto \int f(x+s)dA_L(x)$ is Lipschitz, implying that F_L' is absolutely continuous and differentiable almost everywhere. Except when $-s$ hits an atom of dA_L, its derivative follows by dominated differentiation. The exceptional set of such atoms is at most countable, so

$$F_L''(s) = -\int f'(x+s)\,dA_L(x), \qquad |F_L''(s)| \le \frac{1}{b}\int f(x+s)\,dA_L(x) = -\frac{1}{b}F_L'(s) \tag{OA.68}$$

almost everywhere. Hence

$$U_L''(s) = 2F_L'(s) + sF_L''(s) \le (2 - s/b)F_L'(s) < 0 \tag{OA.69}$$

for $b > 1/2$ on the unit order interval. Since U_L' is absolutely continuous, this strict almost-everywhere inequality makes it strictly decreasing. Any interior root is therefore the unique global maximizer of U_L over all correctly signed magnitudes, including the endpoints.

B.3. Continuity of the equilibrium root equation

Let $[v_-, v_+]$ lie strictly inside $(0,1)$ and preserve $1/2 < \tau < M_v$. The threshold $x^*(r,v)$ is continuous in v. For any convergent sequence of candidate magnitudes, the posterior and the entry residuals converge pointwise except at the limiting entry boundary. The density and its derivative are translated by bounded amounts; they are dominated by constant multiples of $f(x)$, with an additional factor $1/b$ for the derivative. The residuals are bounded by Δ_T. Dominated convergence therefore gives continuity of both $F_L(v;v)$ and its unilateral derivative component, except for individual density kink points that have zero Lebesgue mass and do not change the integral. Thus $\Psi(r,v)$ is continuous.

Outward enclosures satisfying

$$\inf \Psi(r, v_-) > 0, \qquad \sup \Psi(r, v_+) < 0 \tag{OA.70}$$

prove existence of an exact root $v^* \in (v_-, v_+)$. The concavity result proves optimality at that root against its own price schedule. It does not prove that Ψ has only one root across all possible candidate schedules; no such uniqueness is needed for Proposition 3.

B.4. A global cover for favorable-information deviations

The distributional second derivative of the Laplace density is

$$f'' = \frac{f}{b^2} - \frac{\delta_0}{b^2}. \tag{OA.71}$$

Away from zero, ordinary differentiation gives $f'' = f/b^2$. The first derivative jumps from $1/(2b^2)$ to $-1/(2b^2)$ at zero, contributing the atom $-\delta_0/b^2$. This verifies (OA.71), for example by integrating twice against a compactly supported smooth test function.

Convolving with a bounded measurable A_H yields $F_H'' = (F_H - A_H)/b^2$ as distributions and almost everywhere as functions. Since $F_H, A_H \in [0, \Delta_T]$, the weak second derivative is bounded by Δ_T/b^2 in absolute value. The first derivative has an absolutely continuous, Lipschitz representative; combined with $|F'_H| \le \Delta_T/b$, this gives

$$|U''_H(s)| \le L_U = \frac{2\Delta_T}{b} + \frac{\Delta_T}{b^2}, \qquad s \in [0,1] \text{ a.e.} \tag{OA.72}$$

For $s_j = j/n$, evaluate interval enclosures of $U'_H(s_j)$ **uniformly for every v in the root bracket**. Every point of the order interval lies within $1/(2n)$ of a mesh point, so the global bound is

$$\Gamma_H := \min_j \inf U'_H(s_j; [v_-, v_+]) - \frac{L_U}{2n}. \tag{OA.73}$$

A strictly positive lower enclosure for Γ_H proves that buying the maximum amount is uniquely optimal. Checking only the numerical root midpoint does not suffice: the exact root is known only to lie inside the bracket. The interval cover must hold for the whole bracket. All wrong-signed orders are dominated by zero by the residual signs already proved.

B.5. Elementary antiderivatives and infinite tails

Split the integration domain at $-v$, x^*, 1, and the density center $\zeta = \epsilon s$. Within the posterior's central region let $t = e^{(x-c)/b}$, so $\mu_v = t^2/(1+t^2)$ and $dx = b\,dt/t$. I use the following primitives, each verified by differentiation:

$$\begin{array}{c|cc} & \displaystyle\int e^{x/b}(\cdot)\,dx & \displaystyle\int e^{-x/b}(\cdot)\,dx \\ 1-\mu_v & be^{c/b}\arctan t & be^{-c/b}(-t^{-1} - \arctan t) \\ \mu_v & be^{c/b}(t - \arctan t) & be^{-c/b}\arctan t \end{array} \tag{OA.74}$$

For $x < \zeta$, the density multiplier is $e^{-\zeta/b}e^{x/b}/(2b)$; for $x > \zeta$ it is $e^{\zeta/b}e^{-x/b}/(2b)$. Multiply the relevant primitive by this constant, by $e\Delta_T$, and evaluate it at the segment endpoints. On each segment the marginal-density multiplier for F'_ϵ is $\epsilon\,\mathrm{sgn}(x-\zeta)/b$, so the same segment integral gives the derivative. This avoids differentiating the candidate entry threshold during a unilateral deviation.

Outside $[-v, 1]$ the posterior is constant. For a segment $[a_0, a_1]$ left of ζ, its density mass is $[e^{(a_1-\zeta)/b} - e^{(a_0-\zeta)/b}]/2$; right of ζ it is $[e^{-(a_0-\zeta)/b} - e^{-(a_1-\zeta)/b}]/2$. The infinite endpoint terms vanish analytically. Multiplying by the segment's constant residual integrates both infinite tails without truncation.

Endpoint ordering must be certified. When the low trader is evaluated at $s = v$, the center is identically $-v$ and must be coalesced algebraically, not treated as two uncertain endpoints. The same applies at $s = 1$ for the high trader. For distinct uncertain endpoints, prove the ordering by disjoint interval bounds. If an interval center overlaps an entry boundary or another cut, subdivide the parameter box or the order cell until the ordering can be proved; do not select an order using only midpoints.

B.6. Certificate acceptance and what it establishes

I use exact decimal input strings and interval operations rounded outward. In an interval test, the infimum and supremum of a displayed function evaluation mean the endpoints of its outward interval enclosure; the underlying equilibrium function remains scalar-valued. The starting precision and mesh are specified in

Appendix C.2. An enclosure must retain its full endpoint representation; a printed floating-point midpoint is not the certificate. For a bracket to establish the stated result, I require all of the following: the model support inequalities, low-cost participation and high-cost prior exclusion, strict threshold ordering throughout the bracket, opposite enclosed endpoint signs in (OA.70), positive Γ_H in (OA.73), and a positive pooling-existence margin at the same strength. Entry is enclosed by evaluating the exact tail formula throughout the bracket. Pairwise strictly ordered entry intervals establish the cross-economy increase.

Failure of an enclosure to exclude zero is an unresolved certificate, not evidence of a profitable deviation or nonexistence. A failed sign in a claimed successful certificate is an acceptance failure. Raising precision or refining a bracket is legitimate only with the complete failed and successful record retained.

The result is computer-assisted existence for exact equilibria inside specified intervals. It does not prove uniqueness of the informative profile, exclude mixed equilibria, or certify an entire continuous branch between nodes. A continuation curve between certified nodes is a numerical diagnostic until additional interval and continuation arguments establish it. This distinction governs the correspondence figure.

C. Numerical exercises and quantity registry

C.0. Input declarations, output conventions, and acceptance

This section is the complete numerical contract. The literal values in its input declarations specify experiments; the manuscript's quantitative placeholder fields are filled only by validated outputs or these declared inputs. Mathematical constants, equation and result labels, citation metadata, and software identifiers are structural, not estimated quantities. The empirical scaffold contributes no sample size or empirical estimate to the registry.

All parameters below are exact decimal inputs, converted without a binary-floating intermediate when interval certification is required. A numerical result records its full parameter vector, noise law, cost law, information structure, sale institution, tie convention, and branch. Defaults must not silently cross from one exercise to another. Scalar probabilities are stored as probabilities; percentage formatting is a separate presentation step.

Input declaration: benchmark.

```
h = 10; ell = 1; p = 0.5; rho = 0.25;
c_L = 1; c_H = 6; b = 2; k = 0.02;
r_weak = 1.2; r_strong = 3; r_collapse = 3.6;
noise = Laplace; cost = atomic; fundamental_prior_H = 0.5;
entry_at_indifference = yes; bid_equal_reserve_is_admissible = yes;
cost_halfwidth = 0.1; value_band_halfwidth = 0.05;
binary_alternative_reserve = 1.01; atomless_alternative_reserve = 1.1.
```

Input declaration: moderate acquisition values.

```
h = 2; ell = 1; p = 0.5; rho = 0.25;
c_L = 0.3; c_H = 0.89; b = 2; k = 0.002;
r_weak = 1.05; r_strong = 1.5;
noise = Laplace; cost = atomic; fundamental_prior_H = 0.5.
```

Input declaration: complementary signals.

```
h = 10; ell = 1; p = 0.5; rho = 0.85;
c_L = 1; c_H = 7.14; b = 2; k = 0.015;
r_weak = 1.1; r_strong = 2.3; a = 0.70; d = 0.75;
noise = Laplace; cost = atomic; fundamental_prior_H = 0.5;
T_and_Y_independent_given_fundamental = yes.
```

Input declaration: numerical controls.

```
quadrature_absolute_target = 1e-11;
quadrature_relative_target = 1e-11;
probability_acceptance = 1e-8;
price_identity_acceptance = 1e-8;
entry_optimality_acceptance = 1e-8;
deviation_gain_acceptance = 1e-7;
independent_formula_acceptance = 1e-9;
initial_order_intervals = 400;
refined_order_intervals = 800;
initial_flow_halfwidth = 40;
interval_decimal_precision = 50;
certificate_derivative_intervals = 200;
certificate_refined_intervals = 400;
probability_display_decimals = 6;
parameter_display = shortest_exact_decimal;
certificate_display = outward_interval;
```

Input declaration: continuation, event, and outcome controls.

```
pool_reserve = 7;
price_pool_cutoff_grid = -log2 * (1 - j/16), j = 0..16;
price_pool_negative_controls = {c = -1, c = 1, raw posterior inside the c = 0
↪ pool,
                                orders-only deduplication key};
pricing_family_coverage = recorded per row (text);
exhaustive_pricing_search = false unless a separate exhaustive argument is
↪ recorded;
order_identity_tolerance = 1e-6;
atom_identity_tolerance = 1e-8;
event_arithmetic_digits = 50;
event_sign_resolution = 1e-40;
event_offsets = {1e-4, 1e-6, 1e-8};
tie_rule = entry_at_indifference; bid_equal_reserve_admissible;
           upper plateau enters at tau = M; lower plateau enters at B(m) = c_L;
outcome_measures = {E, A, S, C2, O_H} with S = Pr(R >= p) + A - C2;
C2_mixed_search_minimum_starts = 3;
```

These controls govern the continuation records of C.2 and C.6 and the two regression exercises C.6b and

C.6c. A price atom is validated by the rule in (OA.79) applied to its full preimage, followed by the four checks of A.11: the buyer's preparation decision is optimal at its price information set, the conditional target payoff equals the price pointwise in flow, positive-price and zero-price regions do not collide (a collision is one preimage), and every unilateral investor deviation is evaluated against the entire fixed price and preparation schedule. Average price equal to average payoff is a necessary identity and is recorded, but it accepts nothing on its own. An exact declared event is evaluated from its defining identity in fifty-digit arithmetic with the tie rule applied to the identity; a non-event input has its sign evaluated at that precision and is reported unresolved if the sign is closer to zero than the declared resolution. A tolerance for residual acceptance is not a rule for treating nearby inequalities as equalities.

The outcome measures follow A.7 of the paper appendix. With I the preparation indicator and V the realized challenger value, $\mathsf{E} = \Pr(I = 1)$, $\mathsf{A} = \Pr(I = 1, V \geq p)$, $\mathsf{S} = \Pr(R \geq p \text{ or } [I = 1, V \geq p])$, and $\mathsf{C}_2 = \Pr(R \geq p, I = 1, V \geq p)$. Every continuation row must satisfy $\mathsf{S} = \Pr(R \geq p) + \mathsf{A} - \mathsf{C}_2$ by inclusion and exclusion, with S also integrated directly as a union event, and $0 \leq \mathsf{C}_2 \leq \mathsf{A} \leq \mathsf{E} \leq 1$, $0 \leq \mathsf{S} \leq 1$. High-value ownership O_H is computed from the actual allocation event; outside the benchmark support it is not $e_H/2$ by default, since a high-class value can fail the reserve or lose to the incumbent.

The flow halfwidth determines a diagnostic mesh, not an integration truncation. Integrals are evaluated over the full line, or their omitted tails are explicitly bounded. For a bounded residual $A \leq \overline{A}$ and $|q| \leq 1$, a Laplace truncation outside $[-T, T]$, $T > 1$, loses at most $\overline{A}e^{-(T-1)/b}$ in per-unit gross profit. Multiply by the order magnitude for total payoff. Under logistic noise a valid bound is $2\overline{A}/[1 + e^{(T-1)/b}]$. Include these bounds in the error budget; do not compare a tail-truncated integral with a full integral as though both were exact.

I maintain distinct tolerances for equations, quadrature estimates, and strategic deviations, and four separate error budgets: residual approximation, quadrature, tail truncation, and between-grid deviation coverage. Each is recorded in its own column where the exercise produces it (`quadrature_error`, `tail_bound`, `posterior_inversion_error`, `entry_independent_error`, and the deviation-gain estimate with its rigorous upper bound when one exists); the continuation schema's `error_budget` is the sum of the quadrature and tail terms and is never a bound on the other two. A solver residual is never reported as an upper bound on all four. The refinement pass doubles the order resolution and raises the Gauss-Legendre order on every segment; the quadrature error estimate on each pass is compared with the declared target, tightened by a factor of ten on refinement, and a breach is a breach. If machine precision obstructs the tightened target on some integral, the check is marked unresolved for that row rather than passed. Integration warnings are not suppressed. A root or a finite-grid maximum does not establish equilibrium. For each candidate, evaluate

$$\begin{aligned} \epsilon_P &= \sup_{x \text{ tested}} |P(x) - \mathbb{E}[V_T \mid X = x]|, \\ \epsilon_e &= \sup_{(P,C) \text{ tested}} [\text{profit from reversing entry}]_+. \end{aligned} \tag{OA.75}$$

and

$$\epsilon_q = \max_\theta \left\{ \sup_{q\in[-1,1]} U_\theta(q;\sigma) - \int U_\theta(q;\sigma)\, d\sigma_\theta(q) \right\}. \tag{OA.76}$$

The tested price supremum is a numerical diagnostic unless a uniform bound is supplied. The order supremum must be bounded globally for a computer-assisted label; a finite approximation is identified separately. Check all wrong-signed orders as well as correctly signed ones in numerical diagnostics, even when an analytical sign argument excludes them. For mixed profiles, every positive-weight support action must attain the same maximal payoff within the stated tolerance. Recompute the price and entry schedules once per candidate, then hold them fixed during each unilateral-deviation calculation.

Mandatory probability identities are $\int a_\theta = 1$, $\mathbb{E}\mu_X = 1/2$, and $\mathbb{E}P = \mathbb{E}V_T$. Signal exercises add the joint-posterior identities in A.7. Independent auction integration must match closed forms; independent cost-based and flow-based entry integrations must agree. Repeat numerical diagnostics after doubling the order resolution and tightening both integration targets by a factor of `10`, recorded in the run manifest. A claimed accepted row that breaches any acceptance bound terminates validation. Search candidates that fail are retained as rejected rows; inability to resolve a node is recorded as open, not suppressed or interpolated.

CSV is the boundary between numerical calculations and presentation. Files use UTF-8 and quoted headers when a mathematical symbol contains punctuation. Parameter columns use the symbols `h, ell, p, rho, c_L, c_H, b, k, r`; additional columns such as `Delta_T, B_prior, q_H, q_L, e_H, e_L, E, O_H, R_T, tau, x_star` are transliterations of the paper's notation. The data dictionary below each exercise defines every additional symbol. A branch identifier is a data label, not an economic selection rule.

The candidate-level error-budget columns shared by C.2 and C.6 are
`quadrature_error, quadrature_target, quadrature_error_refined, quadrature_target_refined, quadrature_refinement, tail_truncation_bound, tail_truncation_formula, between_grid_spacing, between_grid_lipschitz_bound, between_grid_closed, between_grid_coverage, residual_approximation_error, posterior_inversion_error, entry_independent_error, deviation_argmax_H, deviation_gain_H, deviation_argmax_L, deviation_gain_L, error_budget_breaches`

The initial and refined quadrature estimates have separate targets. `tail_truncation_formula` explains any analytical tail treatment; zero truncation error is permitted when both infinite tails are integrated analytically. `between_grid_lipschitz_bound` measures the untested deviation allowance, and `between_grid_closed` states whether that allowance or a separate argument closes the deviation test. `between_grid_coverage` names that argument. Residual and posterior errors remain separate from quadrature error. The two deviation witnesses identify the best tested order for each investor type and its gain. `error_budget_breaches` records failed targets; an unavailable field is `n/a`.

The quantity registry has columns `name, value, lower, upper, units, display, exercise, parameter_set, branch, status, source_file, source_row, definition`. The `name` is the exact placeholder key without brackets. Each name is unique. A missing or failed quantity remains unresolved; it is never replaced with zero, a cached value from another parameter set, or a manually typed number. Input declarations can be echoed into the registry with status `input`; derived research quantities use the result-status vocabulary and record their validation. This administrative input type is not an additional result status.

C.1. Baseline, controls, extensions, and scalar reproduction

Inputs. Use the benchmark declaration at weak, strong, and collapse strengths. Evaluate the atomic-cost model with Laplace noise first. Then replace the noise law by logistic noise, retaining the scale parameter, and replace cost atoms by a mixture of uniforms on $[c_L - \varepsilon_C, c_L + \varepsilon_C]$ and $[c_H - \varepsilon_C, c_H + \varepsilon_C]$. Their mixture weights remain $\rho, 1-\rho$. These replacements define separate economies.

Acquisition layer. Compute the pointwise auction outcomes from (OA.50). Integrate independently over R, splitting at p and ℓ, and verify $t_0, t_H, t_L, g_H, g_L, \Delta_T$. Compute $B_r(m), B_r(1/2), B_r(M)$. Theorem margins are

$$\begin{aligned} &\zeta_L = B_{r_1}(m) - c_L, \qquad \zeta_{H0} = c_H - B_{r_0}(1/2), \qquad \zeta_{H1} = B_{r_1}(M) - c_H, \\ &\zeta_0 = k - \Delta_T(r_0), \qquad \zeta_1 = (1-1/b)\rho m \Delta_T(r_1) - k. \end{aligned} \tag{OA.77}$$

The collapse node additionally requires $c_H - B_{r_2}(M) > 0$, $B_{r_2}(m) - c_L > 0$, and its full-order margin. For atomless costs replace the appropriate cost endpoints as in Proposition A.6. Verify all margins before assigning an analytical equilibrium label.

Financial and entry layer. Pooling has $q_H = q_L = 0$, $\mu = 1/2$, $e = \rho$, and $P = t_0 + \rho[(t_H + t_L)/2 - t_0]$. Full orders use the conditional flow densities and posterior in A.4 or A.6. Compute τ, the appropriate threshold, conditional entry, ownership, mean price, and revenue. Always handle threshold regions using the cost comparison: a cutoff above the feasible posterior means no high-cost entry; equality at a Laplace plateau uses the specified tie rule. Do not insert an infeasible threshold into an interior tail formula.

For atomless costs, entry at a posterior is the complete mixture CDF $H_C(B_r(\mu))$. Independently compute entry by integrating the tail probability for each cost within each uniform component. The two integrals must agree. Compute expected paid preparation costs as well as entry probability; they are not mean cost multiplied by unconditional entry when the high-cost component is selected.

Controls. The frozen-profile control uses full orders at both strengths and reoptimizes preparation and prices, but explicitly does not require that the weak investor chooses that profile. The price-hidden control reoptimizes investor behavior and pricing, while the buyer uses the prior: pooling at weak strength and full orders at strong strength follow their respective bounds. The matched-dividend control adds exactly $\mathcal{R}_T^{\text{feedback}} - \mathcal{R}_T^{\text{hidden}}$ to the hidden financial payoff and price. Verify that all residuals, orders, and entry probabilities are unchanged. It is not an additional acquisition payment.

Matched-price panel. The panel distinguishes information from the mean financial price at the strong strength, for each noise and cost law in the baseline extension comparison. Its columns are

```
environment, mean_financial_price, seller_revenue, external_dividend,
preparation_probability, high_value_ownership_probability,
↪  net_acquisition_surplus
```

for the three environments `feedback`, `price_hidden`, and `matched_dividend` at the strong strength in each of the four noise and cost-law combinations. The renderer assembles the panel from the accepted rows of `tables/equilibrium_controls.csv` and `numerics/feedback_comparisons.csv`; it solves nothing. The panel must satisfy, row by row,

$$D_0 = \mathcal{R}_T^{\text{feedback}} - \mathcal{R}_T^{\text{hidden}}, \qquad \mathbb{E}[P^{\text{matched}}] = \mathcal{R}_T^{\text{hidden}} + D_0 = \mathbb{E}[P^{\text{feedback}}],$$

The matched row has the same seller revenue, net acquisition surplus, preparation, high-value ownership, and investor residuals as the price-hidden environment. The external dividend is zero in the two equilibrium rows and D_0 in the matched row. Seller revenue in the matched row does not include the dividend. The dividend is attached to the traded claim, is not paid by any bidder, and is not a seller instrument. The panel's note distinguishes the two analytical outcomes from the matched control, which is an invariance diagnostic

Table 1. Matched-price panel: information versus price level at the strong strength.

Environment	$\mathbb{E}[P]$	$\mathcal{R}_T$	D_0	E	O_H	$\mathcal{W}$
Laplace noise, cost atoms, $r = 3$						
Price observed (feedback equilibrium)	0.872392	0.872392	0.000000	0.522757	0.324192	2.382301
Price hidden (reoptimized equilibrium)	0.614583	0.614583	0.000000	0.250000	0.125000	2.302083
Hidden plus dividend (control)	0.872392	0.614583	0.257809	0.250000	0.125000	2.302083
Laplace noise, cost mixture, $r = 3$						
Price observed (feedback equilibrium)	0.872343	0.872343	0.000000	0.522715	0.324148	2.382417
Price hidden (reoptimized equilibrium)	0.614583	0.614583	0.000000	0.250000	0.125000	2.302083
Hidden plus dividend (control)	0.872343	0.614583	0.257760	0.250000	0.125000	2.302083
logistic noise, cost atoms, $r = 3$						
Price observed (feedback equilibrium)	0.662854	0.662854	0.000000	0.301509	0.161994	2.312047
Price hidden (reoptimized equilibrium)	0.614583	0.614583	0.000000	0.250000	0.125000	2.302083
Hidden plus dividend (control)	0.662854	0.614583	0.048271	0.250000	0.125000	2.302083
logistic noise, cost mixture, $r = 3$						
Price observed (feedback equilibrium)	0.662699	0.662699	0.000000	0.301374	0.161854	2.312415
Price hidden (reoptimized equilibrium)	0.614583	0.614583	0.000000	0.250000	0.125000	2.302083
Hidden plus dividend (control)	0.662699	0.614583	0.048116	0.250000	0.125000	2.302083

Strong benchmark strength; the four noise and cost-law combinations of Table 3, Panel A. $\mathbb{E}[P]$ is the mean financial price, $\mathcal{R}_T$ expected target proceeds (seller revenue), $D_0 = \mathcal{R}_T^{\text{feedback}} - \mathcal{R}_T^{\text{hidden}}$ the external dividend, E the preparation probability, O_H high-value ownership, $\mathcal{W}$ net acquisition surplus.
The matched row adds D_0 to the traded claim of the price-hidden economy, so $\mathbb{E}[P^{\text{matched}}] = \mathcal{R}_T^{\text{hidden}} + D_0 = \mathbb{E}[P^{\text{feedback}}]$, while seller revenue, acquisition surplus, preparation, orders, and investor residuals remain those of the hidden-price environment; the renderer verifies each identity to within 10^{-9}. The dividend is attached to the financial claim, is not paid by any bidder, and is not a sale term: the matched row is an analytical invariance diagnostic, not an equilibrium of a new game.

and not a feasible sale mechanism.

Welfare. At strong strength, evaluate (OA.47) or (OA.48), and independently subtract total net allocation surplus in (OA.49). Also compute the revenue difference directly from conditional entry and from the difference in mean prices. Store these as different quantities. A strength comparison is not labeled a welfare effect of access to prices.

Validation. For each equilibrium candidate, integrate residual profits over the full noise line, splitting at density centers, posterior kinks, and entry thresholds. Check all orders on the declared initial and refined grids. Include the special points at which a density center crosses an entry boundary. The analytical margins provide the global result for the specified equilibria; the grid is an implementation check. Verify posterior inversion from the observed price, including atom probabilities in flat tails. In the frozen-profile control retain any profitable investor deviations as expected counterfactual diagnostics rather than mistakenly treating that profile as equilibrium.

Outputs.

```
tables/auction_primitives.csv:
parameter_set, r, p, t_0, t_H, t_L, g_H, g_L, Delta_T, B_m, B_prior, B_M

tables/equilibrium_controls.csv:
parameter_set, noise, cost_law, r, experiment, q_H, q_L,
```

```
e_H, e_L, E, O_H, R_T, W, expected_preparation_cost,
tau, x_star, matched_dividend, status, accepted

tables/extensions.csv:
parameter_set, noise, cost_law, r_weak, r_strong,
E_weak, E_strong, O_H_weak, O_H_strong,
zeta_L, zeta_H0, zeta_H1, zeta_0, zeta_1, status, accepted

numerics/baseline_deviations.csv:
parameter_set, experiment, noise, cost_law, r, state, q,
U(q), U(candidate), deviation_gain, quadrature_error, tail_bound

numerics/feedback_comparisons.csv:
parameter_set, noise, cost_law, r,
W_feedback, W_hidden, W_gain,
R_T_feedback, R_T_hidden, R_T_gain, matched_dividend,
welfare_identity_error, revenue_identity_error,
residual_invariance_error, status, accepted

figures_data/two_returns.csv:
r, mu, Delta_T, B_r(mu), d_Delta_T_dr, d_B_r_mu_dr

tables/matched_price.csv (derived by the table renderer):
parameter_set, noise, cost_law, r, environment, mean_financial_price,
seller_revenue, external_dividend, preparation_probability,
high_value_ownership_probability, net_acquisition_surplus, role, status, accepted
```

W denotes allocation value net of paid preparation costs, not target revenue. The derivatives in the last file are the explicit uniform formulas. Use the correspondence strength grid from C.2 and beliefs $m, 1/2, M$. These outputs feed Tables 1–3 and Figure 1. Scalar keys are defined in C.8; no displayed value is copied directly from a figure.

The complete extension margins accompany the preparation comparisons in the main paper. A negative sufficient-condition margin does not reject a separately validated outcome.

Table 2. Sufficient-condition margins behind Table 3.

	ζ_L	ζ_{H0}	ζ_{H1}	ζ_0	ζ_1	Minimum	Evidence	Unique
Panel A. Noise and preparation costs (benchmark vector)								
Laplace, cost atoms	1.37×10^{0}	1.20×10^{0}	2.17×10^{-1}	3.33×10^{-3}	2.41×10^{-3}	2.41×10^{-3}	analytical	yes
Laplace, cost mixture	1.27×10^{0}	1.10×10^{0}	1.17×10^{-1}	3.33×10^{-3}	2.41×10^{-3}	2.41×10^{-3}	analytical	yes
logistic, cost atoms	1.37×10^{0}	1.20×10^{0}	2.17×10^{-1}	3.33×10^{-3}	2.41×10^{-3}	2.41×10^{-3}	analytical	yes
logistic, cost mixture	1.27×10^{0}	1.10×10^{0}	1.17×10^{-1}	3.33×10^{-3}	2.41×10^{-3}	2.41×10^{-3}	analytical	yes
Panel B. Moderate values								
Moderate values	1.97×10^{-1}	3.35×10^{-2}	3.01×10^{-2}	8.10×10^{-4}	8.01×10^{-4}	8.01×10^{-4}	analytical	yes
Panel C. Complementary private information								
$a = 0.70$, $d = 0.75$	7.73×10^{-1}	5.25×10^{-2}	4.53×10^{-2}	1.45×10^{-3}	1.80×10^{-3}	1.45×10^{-3}	analytical	yes

The five strict sufficient inequalities of the relevant proposition (low cost, high cost at the prior, high cost at the ceiling, weak-economy trade, strong-economy trade), in the order of (OA.77) for Panels A and B and of (OA.29) for Panel C, and their minimum. A positive minimum places the row inside the analytical uniqueness region; a nonpositive minimum would mean the comparison condition is not met, not that a candidate is rejected.

Parameter vectors: Panel A $(h, \ell, p, \rho, c_L, c_H, b, k) = (10, 1, 0.5, 0.25, 1, 6, 2, 0.02)$ with $(r_0, r_1) = (1.2, 3)$; Panel B $(2, 1, 0.5, 0.25, 0.3, 0.89, 2, 0.002)$ with $(1.05, 1.5)$; Panel C $(10, 1, 0.5, 0.85, 1, 7.14, 2, 0.015)$ with $(1.1, 2.3)$ and $(a, d) = (0.70, 0.75)$. The complete records, including the preparation change in percentage points, are in `tables/extensions_margins.csv`.

C.2. The equilibrium correspondence and interval certificates

Inputs and grid. Use the benchmark primitives. The initial strength mesh is the exact decimal arithmetic progression from `1.005` through `3.800` by `0.005`. Add the weak, strong, collapse, and certified strengths and the boundaries $\mathfrak{r}(k), r_N, r_U, r_C$, with offsets `0.0001` and `0.001` on both sides. Boundary formulas are evaluated numerically to select nodes; the declared equality at r_C is evaluated symbolically as $\tau = M$ for preparation on the upper plateau. Adjacent strength intervals whose accepted branch sets differ are refined to spacing `0.001`. The run manifest records the nodes and refinement intervals.

The certificate input brackets are exact decimals:

```
r = 1.55; v_left = 0.46031618; v_right = 0.46031620;
r = 1.60; v_left = 0.70747537; v_right = 0.70747539;
r = 1.65; v_left = 0.90333198; v_right = 0.90333201.
```

Analytically classified outcomes. First verify membership in the maintained floor/prior domain at each node. Pooling existence is given by $k - \rho\Delta_T/2 \geq 0$ there; its sufficient uniqueness bound is $k - \Delta_T > 0$. Full-order uniqueness is established by $(1 - 1/b)\rho m\Delta_T - k > 0$. Record these margins individually. Compute r_N, r_U, r_C from A.5 and verify each defining equation, including support restrictions. Outside a maintained domain, evaluate the actual prior entry rule and continuation instead of retaining a classification based on ρ.

Full-order candidates. At each node, form the full-order posterior and preparation policy. The uniform derivative bound can establish uniqueness, and the sharper candidate-specific J test in (OA.22) can establish existence. A failed sufficient test does not reject the profile. Evaluate both traders on the initial and refined deviation grids, including wrong-signed orders, and record any profitable deviation with its action. For the pure low-type residual, strict concavity reduces boundary optimality to $U_L'(1) \geq 0$. An unenclosed high-type derivative mesh is a diagnostic. Acceptance without an analytical argument or interval cover remains a numerical diagnostic, even if every tested deviation is unprofitable.

Asymmetric continuation. At each r, solve $\Psi(r, v) = 0$ for candidates with $(1, -v)$ over the searched magnitude interval `[0.005, 0.9999]`. Bracket sign changes on a mesh of spacing `0.01`, add the declared certified brackets, and solve bracketed roots by Brent's method. Minimize $|\Psi|$ near observed local minima and subdivide each such neighborhood on a 41-point mesh to seek additional sign changes. A sign change at a discontinuity without a vanishing residual does not establish a root. A validated tangency candidate remains a numerical diagnostic; a failed or unresolved attempt retains its witness. The unilateral derivative in s holds the candidate schedule fixed, whereas changing v changes that schedule.

At each candidate compute $x^*, e_H, e_L, \mathsf{E}, \mathsf{O}_H$, revenue, the pooling-existence flag, and the deviation diagnostics. At the three declared nodes, Appendix B encloses the root, proves its endpoint signs, verifies low-type concavity, and covers high-type deviations uniformly over the root bracket. No other node inherits those certificates. An interior root approaching $v = 1$ is a branch boundary; the boundary profile is evaluated separately.

Other pure and mixed candidates. The complete correspondence remains open. The pure search evaluates $(u, -v)$ on $[0, 1]^2$ from the nine initial pairs in $\{0.2, 0.6, 0.95\}^2$, with forty damped best-response iterations and a local fixed-point solve. Interior pure roots are polished using both traders' held-schedule first-order conditions before revalidation. Every resulting candidate faces the full initial and refined deviation scans. The mixed search is restricted to high-type mixtures over correctly signed orders and a pure low-type magnitude $v \in [0, 1]$. It runs meshes of spacing `0.05` and `0.025`. This restriction defines the search domain;

it is not an exclusion of low-type mixing in every equilibrium. A mesh check of residual monotonicity is recorded, but does not supply the global premise of the Stieltjes lemma for an arbitrary mixed profile.

A mixed candidate with support points $q_{\theta j}$ and weights $\omega_{\theta j}$ has conditional density $a_\theta(x) = \sum_j \omega_{\theta j} f(x - q_{\theta j})$. Recompute its posterior, entry, and competitive price. Its payoff is the weighted average of the support payoffs. All positive-weight points must be best responses within tolerance, and every off-support deviation must be checked. Preserve distinct roots and supports found from different initializations. A candidate search that fails to resolve a node receives an open flag; it does not create an empty equilibrium set.

Mixed-search attempt records. On each mesh, three starts place high-type weights uniformly, near the full purchase, and near zero, with initial low-type magnitudes `0.5`, `0.9`, and `0.1`. Damped multiplicative-weight updates stop on convergence, a detected stall, or the declared iteration budget, whose default is `120`. The final support is then subjected to payoff-indifference and simplex equations, with bounded support additions and removals. A collapsed pure profile is polished by the continuous pure solver and revalidated before deduplication. Each start appears in `numerics/mixed_search_attempts.csv`, including its initial values, domain, budget, stopping reason, removed support points and their residuals, and final support. `numerics/mixed_supports.csv` records each state's final support action, weight, payoff, and gap to the best tested deviation. Any unresolved start gives the node an explicit open search row in `numerics/correspondence.csv`; that row does not invalidate a separately established equilibrium there. None of the recorded search outcomes is a nonexistence theorem.

Continuation identity and deduplication. Two candidates at a node are the same continuation only when their complete strategies (support points and weights within the order identity tolerance) and their induced price information (every positive-mass price atom's value, state masses, and kind within the atom identity tolerance) agree, together with the preparation rule; this is the rule of (OA.53) applied numerically, and C.6 uses the same one. On the benchmark strength grid the low-cost floor $c_L < B_r(m)$ holds at every node, entry is positive at every flow, and no price pool combines distinct raw posteriors. Constant-posterior atoms remain possible, and their prices are fixed by the orders, so identity reduces to the complete order profile there; the code applies the full rule regardless. Strength labels are normalized to their shortest exact decimal before nodes are compared, so a strength written two ways is one node. A root found by $|\Psi|$ minimization is compared with the roots found by sign change at the same node and merged when it is the same continuation; the raw attempt stays in the ledger with the identifier of the row that absorbed it. `multiplicity_found` counts distinct accepted continuations after this merge, never raw candidates. Rejected candidates keep their breach witness; a candidate rejected because a sufficient condition failed is not a rejection, and the status vocabulary of C.0 applies.

Figure discipline. Plot analytical uniqueness regions separately from existence-only regions. Show every accepted observed branch. Certified points have interval bars. Numerical lines are broken at failed validation, discontinuities, branch changes, or unresolved gaps. The multiplicity flag means that distinct equilibria have been established or found at a node, not that all equilibria have been enumerated. The ordered certified entry intervals provide the stated cross-economy comparison; a fitted derivative on an exploratory line is not substituted for that result.

Outputs.

```
numerics/correspondence.csv and numerics/correspondence_attempts.csv:
r, branch, q_H, q_L, v, e_H, e_L, E, O_H, R_T, tau, x_star, pooling_exists,
```

```
pooling_unique_bound, full_unique_bound, existence_status, uniqueness_status,
 ↪  accepted,
multiplicity_found, epsilon_P, epsilon_e, epsilon_q, tail_bound,
 ↪  unresolved_reason,
candidate_id, continuation_id, parameter_set_id, search_path, duplicate_of,
n_distinct_accepted, result_status
plus the shared error-budget columns in C.0

numerics/mixed_supports.csv:
r, branch, state, support_index, q, weight, U(q), support_gap,
 ↪  off_support_gain_mesh,
accepted, status, init_id, mesh_spacing, iterations, stop_reason, final_gap,
gap_to_best_tested, simplex_residual, residual_monotone, quadrature_error,
search_domain, outcome, candidate_id, continuation_id, parameter_set_id

numerics/mixed_search_attempts.csv:
r, mesh, start_id, start_description, v_init, domain, budget, iterations,
 ↪  stop_reason,
final_gap, final_v_residual, support_solve, support_solve_residual,
 ↪  support_rounds,
pruned_points, support, support_payoffs, gap_to_best_tested, best_tested_q, v,
simplex_residual, min_weight, residual_monotone, outcome, witness,
 ↪  low_support_payoff,
low_gap_to_best_tested, validation

numerics/certificates.csv:
r, v_lower, v_upper, Psi_left_lower, Psi_left_upper,
Psi_right_lower, Psi_right_upper, Gamma_H_lower,
L_U_upper, mesh_intervals, interval_digits,
E_lower, E_upper, pooling_margin_lower, threshold_margin_lower, accepted

numerics/thresholds.csv:
boundary, value, lower, upper, defining_residual,
in_support_domain, low_cost_floor_valid, interpretation
```

`correspondence_attempts.csv` keeps raw candidate attempts and their `duplicate_of` links. `correspondence.csv` removes duplicate attempts while preserving rejected and open records. `search_path` describes the candidate construction; `n_distinct_accepted` and `multiplicity_found` count distinct accepted continuations at the node. The support file links each validated strategy to the full candidate and continuation identity; unvalidated starts carry `n/a` for those foreign keys. The attempt ledger identifies `converged-candidate`, `converged-pure`, and `unresolved` search outcomes separately from validation status. `support_payoffs`, `gap_to_best_tested`, `pruned_points`, and `witness` preserve the reasons for retaining, removing, or rejecting a support.

The threshold file uses exact row labels `pooling_unique_sufficient`, `pooling_existence`, `full_orders_unique_sufficient`, and `high_cost_ceiling`. It also carries separately labeled analytical scalars `m`, `M`, and `laplace_entry_left_limit` needed by the registry; these are not mislabeled activation thresholds. Every row stores its defining expression and checks. Fields that do not apply to a probability scalar are marked not applicable, not set to zero.

All general parameter columns are included or keyed to an immutable parameter declaration whose full contents accompany the file. `uniqueness_status` distinguishes analytical uniqueness from not established; it is never inferred from the number of search hits. `off_support_gain_mesh` is a mesh maximum and is named as one; it is not a bound. `quadrature_error` is checked against the declared target on each pass, and `tail_bound` is zero under Laplace noise only because the exterior tails are integrated in closed form against the constant residual. Feed: Figure 2, Proposition 3, the threshold discussion, and the baseline table. Registry keys `cert_a_*`, `cert_b_*`, and `cert_c_*` refer to the ordered declared nodes, not arbitrary roots selected from a larger search.

C.3. Complementary private signals

Inputs. Start with the complementary-signal declaration. Then cross trader accuracies `{0.68, 0.69, 0.70, 0.71, 0.72}` with buyer accuracies `{0.73, 0.74, 0.75, 0.76, 0.77}`, holding other primitives fixed. These are separate equilibria with exact decimal parameters, not a claim that every grid point meets Proposition A.7.

Objects. Compute m, M, public fundamental bounds μ_-, μ_+, the worst joint buyer posterior $\phi_-(\mu_-)$, and the best joint posterior $\phi_+(\mu_+)$. Compute the five strict margins from (OA.29): the low-cost floor, high-cost exclusion using the buyer's favorable private signal at weak strength, high-cost profitability after a favorable strong public price and private signal, weak residual exclusion, and strong global derivative margin. If a margin fails, withhold the analytical theorem label; do not conclude that the economic reversal is absent.

At both strengths construct signal-contingent order laws and the correct state-dependent entry policy. For the theorem-supported weak node public belief is the prior and both high-cost private-signal types stay out. For the strong node use full orders by T, then compute the public posterior from λ_X, each buyer posterior $\phi_Y(\mu_X)$, and both private-signal cutoffs in (OA.39). Always handle both $Y = +$ and $Y = -$, including always/never and plateau-equality cases. The benchmark example's exclusion of one private-signal type is an outcome, not an imposed simplification for the sweep.

Independent verification. First integrate entry using (OA.40)–(OA.41). Independently sum over the finite fundamental, trader-signal, and buyer-signal states and integrate Z and C. Compute the price directly as the conditional expectation of V_T given flow and compare it with (OA.35). Recover the public posterior from the price inversion and then update using Y; compare with the direct joint likelihood. Verify the residual identities (OA.37) by direct conditional trader payoffs. Finally evaluate all unilateral orders using the fixed candidate schedules. The investor must not be given the buyer's private signal during this calculation.

Outputs.

```
numerics/two_signals.csv:
a, d, r, q_plus, q_minus, mu_lower, mu_upper,
phi_minus_mu_lower, phi_plus_mu_upper,
x_star_Yplus, x_star_Yminus, e_H, e_L, E, O_H, R_T,
```

```
low_cost_margin, private_only_exclusion_margin, joint_entry_margin,
weak_order_margin, strong_order_margin,
posterior_error, residual_error, epsilon_P, epsilon_q,
status, accepted

numerics/two_signal_deviations.csv:
a, d, r, trader_signal, q, U(q), U(candidate),
deviation_gain, quadrature_error, tail_bound
```

The same complete primitive vector accompanies each row. The scalar-margin names map to these columns in order: low cost to `low_cost_margin`, high prior to `private_only_exclusion_margin`, high ceiling to `joint_entry_margin`, weak trade to `weak_order_margin`, and strong trade to `strong_order_margin`. These are comparison-level margins; repeat them consistently on both strength rows or store them in a separately keyed comparison record. The display accuracy keys are percentages generated from the raw probability keys, never separately typed. Table 3 in the main paper uses the declared example. The signal table below reports every accuracy combination and distinguishes the sufficient uniqueness region from other validated outcomes. These outputs also supply the example of Proposition A.7.

The full signal comparison is reported below. The declared example is marked; every other accuracy combination remains visible.

Table 3. Complementary private information: complete accuracy grid.

Investor a	Buyer d	E weak	E strong	$\Delta\mathsf{E}$ (pp)	O_H weak	O_H strong	Minimum margin	Evidence	Comparison condition
0.68	0.73	0.850000	0.850000	0.00	0.425000	0.425000	-1.49×10^{-1}	numerical diagnostic	not met
0.68	0.74	0.850000	0.850000	0.00	0.425000	0.425000	-7.71×10^{-2}	numerical diagnostic	not met
0.68	0.75	0.850000	0.850000	0.00	0.425000	0.425000	-5.30×10^{-3}	numerical diagnostic	not met
0.68	0.76	0.925000	0.880132	-4.49	0.482000	0.449575	-3.75×10^{-2}	numerical diagnostic	not met
0.68	0.77	0.925000	0.882492	-4.25	0.482750	0.451741	-1.27×10^{-1}	numerical diagnostic	not met
0.69	0.73	0.850000	0.850000	0.00	0.425000	0.425000	-1.22×10^{-1}	numerical diagnostic	not met
0.69	0.74	0.850000	0.850000	0.00	0.425000	0.425000	-5.08×10^{-2}	numerical diagnostic	not met
0.69	0.75	0.850000	0.878570	2.86	0.425000	0.448157	9.57×10^{-4}	analytical	met
0.69	0.76	0.925000	0.880909	-4.41	0.482000	0.450290	-3.75×10^{-2}	numerical diagnostic	not met
0.69	0.77	0.925000	0.883109	-4.19	0.482750	0.452331	-1.27×10^{-1}	numerical diagnostic	not met
0.70	0.73	0.850000	0.850000	0.00	0.425000	0.425000	-9.54×10^{-2}	numerical diagnostic	not met
0.70	0.74	0.850000	0.850000	0.00	0.425000	0.425000	-2.48×10^{-2}	numerical diagnostic	not met
0.70*	0.75	0.850000	0.879438	2.94	0.425000	0.448942	1.45×10^{-3}	analytical	met
0.70	0.76	0.925000	0.881610	-4.34	0.482000	0.450945	-3.75×10^{-2}	numerical diagnostic	not met
0.70	0.77	0.925000	0.883674	-4.13	0.482750	0.452879	-1.27×10^{-1}	numerical diagnostic	not met
0.71	0.73	0.850000	0.850000	0.00	0.425000	0.425000	-6.88×10^{-2}	numerical diagnostic	not met
0.71	0.74	0.850000	0.878066	2.81	0.425000	0.447687	9.59×10^{-4}	analytical	met
0.71	0.75	0.850000	0.880215	3.02	0.425000	0.449658	7.77×10^{-4}	analytical	met
0.71	0.76	0.925000	0.882247	-4.28	0.482000	0.451551	-3.75×10^{-2}	numerical diagnostic	not met
0.71	0.77	0.925000	0.884195	-4.08	0.482750	0.453392	-1.27×10^{-1}	numerical diagnostic	not met
0.72	0.73	0.850000	0.850000	0.00	0.425000	0.425000	-4.24×10^{-2}	numerical diagnostic	not met
0.72	0.74	0.850000	0.878915	2.89	0.425000	0.448457	6.16×10^{-4}	analytical	met
0.72	0.75	0.850000	0.880919	3.09	0.425000	0.450315	1.00×10^{-4}	analytical	met
0.72	0.76	0.925000	0.882833	-4.22	0.482000	0.452115	-3.75×10^{-2}	numerical diagnostic	not met
0.72	0.77	0.925000	0.884680	-4.03	0.482750	0.453877	-1.27×10^{-1}	numerical diagnostic	not met

Notes. All 25 accuracy combinations in the declaration in Section C.3 are reported; none is dropped. Other parameters are $h = 10$, $\ell = 1$, $p = 0.5$, $\rho = 0.85$, $c_L = 1$, $c_H = 7.14$, $b = 2$, $k = 0.015$, with strengths $r_0 = 1.1$ and $r_1 = 2.3$. E and O_H are probabilities; $\Delta\mathsf{E}$ is the change in preparation in percentage points from the same validated probabilities. The minimum margin is the smallest of the five conditions in (OA.29); the comparison condition is met when it is strictly positive (6 rows), which places the row inside the analytical uniqueness region. A condition that is not met is not a rejection: every row is a validated equilibrium of the finite check, and its evidence is then numerical diagnostic unless a separate analytical or interval argument supports more. The asterisk marks the example used in the main text. All rows pass the declared numerical acceptance checks.

C.4. Moderate values and constructive nonemptiness

Inputs. Reproduce the moderate declaration exactly. Compute the same primitives, theorem margins, complete equilibrium objects, and controls as in C.1. Use independent auction integration and global analytical margins, not only a best-response mesh, to validate the theorem-supported example.

For the nonemptiness construction, set `ell = 1`, `p = 0.5`, `rho = 0.25`, `b = 2`, and use high-to-low ratios `{1.05, 1.10, 1.25, 1.50, 2.00}`. For each ratio let $D = h - \ell$ and examine $\delta_j = D2^{-j}$ for integer j from 1 through `20`. Define $r_1 = \ell + \delta_j$ and $r_0 = \ell + \delta_j^2/D$. Test the conditions needed in A.6. When the intervals are nonempty choose

$$\begin{aligned} k &= \frac{\Delta_T(r_0) + (1 - 1/b)\rho m \Delta_T(r_1)}{2}, \\ c_H &= \frac{B_{r_0}(1/2) + B_{r_1}(M)}{2}, \qquad c_L = \frac{B_{r_1}(m)}{2}. \end{aligned} \tag{OA.78}$$

Every chosen value is an output of the construction, not a fixed calibration imposed on all ratios. Check $c_L < c_H$, support, and all theorem inequalities explicitly. Keep the sequence of failed and successful j values. Small margins require adequate arithmetic precision; floating-point cancellation is not evidence against nonemptiness. Report the first certified feasible member of the declared sequence and the numerical scale of k, not an invented uniform lower bound on its size.

Outputs.

```
numerics/moderate_values.csv:
h, ell, p, rho, c_L, c_H, b, k, r_weak, r_strong,
zeta_L, zeta_H0, zeta_H1, zeta_0, zeta_1,
E_weak, E_strong, O_H_weak, O_H_strong, status, accepted

numerics/nonemptiness_construction.csv:
h_over_ell, j, delta, r_weak, r_strong, k, c_L, c_H,
all_margins_min, feasible, arithmetic_precision, reason
```

Feed: Table 3 and the moderate-value paragraph. The general construction is an analytical argument; the finite sequence is a diagnostic illustration of its scale.

C.5. Noise laws, threshold distance, and posterior upper tails

Inputs. Fix benchmark strong strength, b, and ρ. Use full orders. Form `401` equally spaced points for the normalized threshold distance $d_\tau = (M - \tau)/(M - 1/2)$ on $[0, 1]$, and add the benchmark τ, $M - 10^{-4}$, $M - 10^{-6}$, and $M + 10^{-6}$. At the endpoints use the stated mathematical limits and tie convention rather than unstable logarithms. Compare Laplace and logistic laws at the same scale parameter; do not claim equal variance or a Blackwell ordering between the laws.

Objects. For each threshold, compute the order-flow cutoff, $\alpha_H(\tau), \alpha_L(\tau)$, posterior upper-tail mass $[\alpha_H + \alpha_L]/2$, and the induced entry $\rho + (1 - \rho)[\alpha_H + \alpha_L]/2$. Verify each tail by direct integration of the conditional flow densities. For the baseline logistic example, compute $x^*/(b\pi/\sqrt{3})$ from the derived noise variance, not from the aggregate-flow variance. Store both variances separately if both standardized cutoffs are reported.

A threshold scan is first an information-experiment comparison. To interpret a row as an equilibrium with changed preparation cost, set $c_H(\tau) = g_L + \tau(g_H - g_L)$, keep the other primitives fixed, and recheck the low-cost floor, cost ordering, and global full-order bound. Label only validated rows as such. At the strong benchmark the full-order bound does not depend on the high-cost threshold, but that fact must not be presumed after changing other primitives.

At $\tau = M$, Laplace admits entry in its entire upper posterior plateau under the tie convention; logistic has zero mass at the unattained endpoint. For $\tau > M$, both tail masses are zero. This distinction is part of the plotted result. On the lower endpoint $\tau = 1/2$, the threshold is zero and no artificial singularity arises.

Outputs.

```
figures_data/posterior_tails.csv:
noise, b, tau, M_minus_tau, normalized_distance, x_star,
alpha_H, alpha_L, posterior_upper_tail_mass, E,
implied_c_H, noise_variance, flow_variance, threshold_noise_sd,
full_order_margin, equilibrium_interpretation_valid,
integration_error, status
```

Feed: Figure 3, the logistic magnitude paragraph, and Table 3's noise rows. Infinite thresholds are stored with an explicit `unattainable` label, not converted to an arbitrary finite cutoff.

C.6. Reserve comparisons and exploratory seller continuation

Inputs. Reproduce the binary-value comparison using the benchmark reserves `0.5` and `1.01` at strengths `1.2` and `3`. Reproduce the atomless-value class economy using half-width `0.05` and reserves `0.5` and `1.1` at the same strengths. In the class economy the investor sees the class only and preparation reveals the exact value. Do not let the investor trade on the within-class realization.

Payoff oracle. Integrate the realized outcomes (OA.50) over the independent uniform incumbent and each conditional value distribution. Split at $R = p$, $R = v$, and all reserve/value support boundaries. Verify against (OA.51) and, where applicable, (OA.52). Compute class-conditional t_H, t_L, g_H, g_L for every reserve, including reserves inside a value band. A formula for $p < \ell$ or for exclusion of an entire class is not used inside a partially excluded band. Each reserve's regime on the full domain (below ℓ, at ℓ, between ℓ and r, at r, between r and h, at h, above h, and the band cases for the class economy) is decided by exact comparison of the declared numbers and written to the row.

Complete continuations. A candidate at a reserve is accepted only as a complete continuation in the sense of (OA.53): the investor's state- or class-contingent order distribution, the pricing rule with every positive-mass price pool, the buyer's information at each price (the pooled posterior on an atom, the inverted posterior elsewhere), the preparation rule by observed price and cost, and the validation evidence. For a candidate order profile, compute its raw flow posterior μ_X. When $\Delta_T > 0$ and there is positive entry, the benchmark inversion applies to the positive-entry prices. If entry vanishes on a set of flows, those flows pool at $P = t_0$, and their posterior is the integral over the full preimage of that price:

$$\mu(P = t_0) = \frac{\int_{\{x:P(x)=t_0\}} a_H(x)\,dx}{\int_{\{x:P(x)=t_0\}} [a_H(x) + a_L(x)]\,dx}. \tag{OA.79}$$

The formula is used only when the denominator is positive. The buyer's entry rule at that price uses this pooled posterior and nothing finer; a candidate whose buyer acts on the raw posterior inside the pool is not measurable with respect to the price and is rejected. The convenient construction for independent costs, $P(\mu) = t_0 + H_C(B(\mu))[t_L - t_0 + \Delta_T \mu]$, prices every positive-entry flow by the inversion and pools exactly the zero-entry preimage; it is strictly increasing wherever entry is positive and $\Delta_T > 0$. It is the one pricing family the C.6 node solver constructs, and the row says so: `pricing_family_coverage` records that alternative cutoffs were not scanned at that node and `exhaustive_pricing_search` is false. Section A.11 shows why this matters. In the example of A.11, the same orders admit a family of pools, each a different continuation, and the convenient construction is one member. This example does not establish such a family at every reserve. Every accepted candidate must pass the price-atom validation of C.0: pooled posterior from the full preimage, buyer optimality at the pooled posterior and pointwise on positive prices, competitive pricing pointwise in flow on both regions, no collision between the atom and the positive-price range, and every unilateral deviation against the whole fixed schedule.

At reserves with $\Delta_T = 0$, all entry impossible, or a degenerate payoff, analyze the constant-price candidate directly. Pooling uses the actual prior entry probability $e_0 = H_C(B(1/2))$, and its correctly signed deviation coefficient is $e_0\Delta_T/2 - k$; it is not automatically $\rho\Delta_T/2 - k$. When $e_0 = 0$, no trade and no entry are supported by constant pricing. All accepted candidates must satisfy the original conditional pricing and entry conditions.

Specified comparisons. At the declared reserves, verify the relevant strict no-trade or full-order bounds using the class payoff spread and the actual floor $e(m) = H_C(B_r(m))$. Compute entry and seller revenue and compare each alternative reserve with the original one at the same strength. This reproduces the profitable alternatives, not a global reserve optimum.

Table 4. Reserve comparisons: orders, margins, evidence, and identity behind Table 4.

Economy and reserve	q_H	q_L	E	A	O_H	$\mathcal{R}_T$	Floor margin	Trading margin	Evidence	Unique
Panel A. Binary values										
Weak (r = 1.2), p = 0.5	0	0	0.250000	0.250000	0.125000	0.392708	1.73×10^{0}	3.33×10^{-3}	analytical	yes
Weak (r = 1.2), p = 1.01	1	-1	0.543573	0.338494	0.338494	0.452756	1.41×10^{0}	9.08×10^{-3}	analytical	yes
Strong (r = 3), p = 0.5	1	-1	0.522757	0.522757	0.324192	0.872392	1.37×10^{0}	2.41×10^{-3}	analytical	yes
Strong (r = 3), p = 1.01	1	-1	0.513373	0.317504	0.317504	0.987487	1.24×10^{0}	1.36×10^{-2}	analytical	yes
Panel B. Atomless values with class information										
Weak (r = 1.2), p = 0.5	0	0	0.250000	0.250000	0.125000	0.392665	1.73×10^{0}	2.99×10^{-3}	analytical	yes
Weak (r = 1.2), p = 1.1	1	-1	0.540309	0.336302	0.336302	0.432173	1.39×10^{0}	1.40×10^{-2}	analytical	yes
Strong (r = 3), p = 0.5	1	-1	0.522760	0.522760	0.324195	0.872367	1.37×10^{0}	2.42×10^{-3}	analytical	yes
Strong (r = 3), p = 1.1	1	-1	0.511638	0.316252	0.316252	1.014500	1.23×10^{0}	1.38×10^{-2}	analytical	yes

Complete parameter identity of every node: `h=10|ell=1|rho=0.25|c_L=1|c_H=6|b=2|k=0.02|noise=Laplace|cost=atoms|cost_halfwidth=0|prior_H=0.5` with `r`, `p`, `value_law` ∈ {binary, uniform_classes} and `eps_V` ∈ {0, 0.05} as listed. Institution, information structure, and tie rule: `reserve_auction:uniform_incumbent[0,r];sale_to_highest_admissible_at_max(p,min(R,v));bid_equal_reserve_is_admissible=yes`, `public_price_feedback;investor_observes=fundamental_state;buyer_observes=price_only`, `entry_at_indifference=yes;bid_equal_reserve_is_admissible=yes;upper_plateau_enters_at_tau=M;lower_plateau_enters_at_B(m)=c_L`; `reserve_auction:uniform_incumbent[0,r];sale_to_highest_admissible_at_max(p,min(R,v));bid_equal_reserve_is_admissible=yes`, `public_price_feedback;investor_observes=value_class_not_within_class_value;buyer_observes=price_only`, `entry_at_indifference=yes;bid_equal_reserve_is_admissible=yes;upper_plateau_enters_at_tau=M;lower_plateau_enters_at_B(m)=c_L`.

A is the probability of an admissible prepared challenger. Floor margin is $B_r(m) - c_L$; trading margin is $k - \Delta_T$ for no trade and $(1 - 1/b)\, e(m)\, m\, \Delta_T - k$ for full orders. Evidence is the result status; Unique reports the uniqueness scope recorded with the continuation (unique within all continuations under the stated bound, or not established). Each node is a validated row of `numerics/reserve_events.csv` joined with `tables/reserve_comparisons.csv` on the full parameter identity, the exact reserve, and the branch.

Event grid. The sweep and the event exercise share one grid rule. Beyond the declared reserve mesh, every node set contains the support and participation events $0, \ell, r, h$; in the class economy $\ell \pm \varepsilon_V$ and $h \pm \varepsilon_V$; the declared reserve alternatives; and the two exact events of the paper appendix, the low-cost floor equality $p_L = h - c_L/m$ and the expensive-entry ceiling equality $p_H = \sqrt{2r(h - c_H/M) - r^2}$, each attached to the strength at which its defining regime applies. Every event carries one-sided offsets `1e-4`, `1e-6`, and `1e-8` on both sides where the offset stays in $[0, h]$ (or $[0, h + \varepsilon_V]$), so that the regime classification on each side is tested separately. An exact event is stored by its defining relation and recomputed from it at fifty digits; its decimal expansion is reference only and never decides a tie. The tie rule is applied to the identity: at p_L the low-cost type enters on the lower plateau under $B(m) = c_L$, and at p_H the high-cost type enters on the upper plateau under $\tau = M$ with $x^* = 1$. A node whose inequality sign cannot be resolved at that precision is written as unresolved. The same continuation need not be supported on both sides of an event, and the record does not assume it.

Exploratory sweep. The initial binary-value reserve mesh runs from `0` to `h` in steps of `0.05`; the class-value mesh runs to `h + epsilon_V`, with the event nodes added in both cases. Intervals with changing accepted branch sets, an unresolved endpoint, or a found revenue maximum are refined to spacing `0.002`. At each applicable node the solver checks pooling, full orders, asymmetric roots and tangencies, and pure profiles from the nine starts in $\{0.2, 0.6, 0.95\}^2$, with forty damped best-response iterations and a local fixed-point solve. Its mixed search uses a `0.05` high-type order mesh, initially uniform weights, an initial pure low-type magnitude `0.5`, and at most `150` multiplicative-update iterations with tolerance `1e-7`. This single restricted mixed search differs from the multiple-start C.2 exercise and supplies no claim about unsearched supports. Candidate schedules use the pricing family described above. Initializations and outcomes are retained in the attempt ledger; there are no bidirectional continuation warm starts. Analytical uniqueness can make the remaining candidate searches unnecessary at a node, and the manifest records that reason.

Identity and deduplication. Every candidate is written to the ledger with a `candidate_id` that records how it was obtained (exercise, node, branch, and sequence) and a `continuation_id` that is the canonical economic identity: the parameter-set identity built from the exact declarations, the institution, the information structure, the tie rule, the complete order profile at the order identity tolerance, the pricing family, the preparation rule, and every positive-mass price atom at the atom identity tolerance. Atoms and pools below the numerical probability cutoff are omitted from the numerical identity comparison. This computational cutoff does not make a positive-probability event a mathematical null set; its mass remains subject to the probability error budget. Accepted candidates are merged only when their continuation identities agree; the absorbed candidate keeps its row with `duplicate_of` pointing at the representative. Candidates are never merged on a rounded strength or reserve, on orders or mixed supports alone, on entry probability alone, on mean price or revenue alone, or on a solver's branch name. Two candidates with the same orders and different price atoms stay distinct, and the row's `identity_note` says so. Counts of continuations found, rejected candidates, unresolved candidates, and merged duplicates are computed from this ledger, never from the grid length.

Reporting. For each reserve retain every accepted continuation found. Report the minimum and maximum entry and revenue across those continuations and the search outcome, which is one of: analytically unique continuation; accepted continuations found with uniqueness not established; no accepted candidate found by the declared searches, which is not a nonexistence proof; or unresolved search with the open rows retained.

These are found-continuation ranges, not certified envelopes of all equilibria. Never replace a failed solve with the nearest successful reserve, optimize a frozen information experiment, or label the best sampled reserve globally optimal. Any proposed seller equilibrium must specify its continuation at unchosen reserves as well as chosen ones.

The online summary of the sweep is headed "highest revenue among continuations found", never "optimal reserve". It lists, per economy and strength, the reserve and continuation with the highest accepted revenue, the number of reserve nodes attempted, the number with more than one distinct accepted continuation, and the number of unresolved nodes, all computed from the attempt ledger. It carries its known limits in the same table note: the search covers the declared candidate classes only, the pricing family at each node is the benchmark construction, and no envelope is certified. It is not in the paper.

The coverage counts use the final pass at each reserve. Earlier attempts that triggered continuation-identity refinement remain in the numerical archive, with their inputs and outcomes; they are not counted as additional final-pass nodes. A bounded refinement of near-root candidates is followed by full revalidation before economically identical continuations are merged.

Table 5. Highest revenue among continuations found: exploratory reserve summary with coverage counts.

	Reserve	Node	Found	E range	$\mathcal{R}_T$ range	Search outcome
Panel A. Highest revenue among continuations found						
Binary values, $r = 1.2$	p_L (≈ 6.281718)	event: floor equality p L	1	0.250000	0.785215	analytically unique continuation
Binary values, $r = 3$	p_H (≈ 1.325270)	event: ceiling equality p H	1	0.506477	1.068854	analytically unique continuation
Class values, $r = 1.2$	p_L (≈ 6.281718)	event: floor equality p L	1	0.250000	0.785215	analytically unique continuation
Class values, $r = 3$	p_H (≈ 1.325270)	event: ceiling equality p H	1	0.506477	1.068854	analytically unique continuation

	Nodes	Events	Several found	Unresolved	None found	Candidates evaluated	Rejected	Cand. unresolved	Duplicates merged
Panel B. Search coverage from the final-pass attempt ledger									
Binary values, $r = 1.2$	392	8	148	116	0	2688	288	116	1744
Binary values, $r = 3$	392	8	148	116	0	2721	321	116	1744
Class values, $r = 1.2$	464	12	155	155	0	3249	382	136	2112
Class values, $r = 3$	416	12	148	116	0	2756	332	116	1744

Highest revenue among continuations found, not an optimal reserve: Panel A selects, for each economy, the node (grid point or exact event) whose highest accepted continuation revenue is largest among the continuations the declared searches found, and reports the ranges of preparation and revenue across the accepted continuations at that node. Range endpoints need not belong to the same continuation; a single value denotes a singleton found range. No global envelope is certified; a node without an accepted continuation is not a nonexistence result, and an unresolved node does not imply an empty equilibrium set.

Panel B counts come from the final-pass attempt and validation ledger of `numerics/reserve_ranges.csv`, one record per solved node: nodes attempted, exact-event nodes among them, nodes with several accepted continuations, nodes left unresolved, nodes with no accepted continuation, candidates evaluated, rejected, unresolved, and duplicates merged by the continuation identity. The following exact-event table lists every exact-event node (reserve at zero, at ℓ, at r, at h, at the class-band edges, at the declared reserves, at the published sampled weak reserve, at the weak-incumbent floor equality $p_L = h - c_L/m$, and at the high-cost ceiling equality p_H) with the candidates found there; one-sided offsets from each event are separate nodes counted in Panel B. Earlier attempts that triggered identity refinement are preserved separately in the numerical diagnostics archive. Known limits: the searches cover the declared candidate families and the declared pricing family only (Online Appendix C.6).

Table 6. Exact-event candidates in the exploratory reserve search.

	Event	Reserve	Found	$\mathcal{R}_T$ max	Unresolved	Search outcome
Binary, $r = 1.2$	reserve zero	0	1	0.147917	no	analytically unique continuation
Binary, $r = 1.2$	declared reserve 0.5	0.5	1	0.392708	no	analytically unique continuation
Binary, $r = 1.2$	reserve equals ell	1	1	0.377083	no	analytically unique continuation
Binary, $r = 1.2$	declared reserve 1.01	1.01	1	0.452756	no	analytically unique continuation
Binary, $r = 1.2$	reserve equals r	1.2	1	0.400229	no	analytically unique continuation
Binary, $r = 1.2$	ceiling equality p H	p_H ($\approx$ 1.691903)	1	0.535726	no	analytically unique continuation
Binary, $r = 1.2$	floor equality p L	p_L ($\approx$ 6.281718)	1	0.785215	no	analytically unique continuation
Binary, $r = 1.2$	reserve equals h	10	1	0.000000	no	accepted continuation(s) found; uniqueness not established
Binary, $r = 3$	reserve zero	0	1	0.655633	no	analytically unique continuation
Binary, $r = 3$	declared reserve 0.5	0.5	1	0.872392	no	analytically unique continuation
Binary, $r = 3$	reserve equals ell	1	1	1.049608	no	analytically unique continuation
Binary, $r = 3$	declared reserve 1.01	1.01	1	0.987487	no	analytically unique continuation
Binary, $r = 3$	ceiling equality p H	p_H ($\approx$ 1.325270)	1	1.068854	no	analytically unique continuation
Binary, $r = 3$	reserve equals r	3	1	0.375000	no	analytically unique continuation
Binary, $r = 3$	floor equality p L	p_L ($\approx$ 6.281718)	1	0.785215	no	analytically unique continuation
Binary, $r = 3$	reserve equals h	10	1	0.000000	no	accepted continuation(s) found; uniqueness not established
Classes, $r = 1.2$	reserve zero	0	1	0.147873	no	analytically unique continuation
Classes, $r = 1.2$	declared reserve 0.5	0.5	1	0.392665	no	analytically unique continuation
Classes, $r = 1.2$	reserve equals ell minus epsV	0.95	1	0.390321	no	analytically unique continuation
Classes, $r = 1.2$	reserve equals ell	1	1	0.540618	yes	unresolved search (open nodes or unconverged starts retained)
Classes, $r = 1.2$	reserve equals ell plus epsV	1.05	1	0.444542	no	analytically unique continuation
Classes, $r = 1.2$	declared reserve 1.1	1.1	1	0.432173	no	analytically unique continuation
Classes, $r = 1.2$	reserve equals r	1.2	1	0.400229	no	analytically unique continuation
Classes, $r = 1.2$	ceiling equality p H	p_H ($\approx$ 1.691903)	1	0.535726	no	analytically unique continuation
Classes, $r = 1.2$	floor equality p L	p_L ($\approx$ 6.281718)	1	0.785215	no	analytically unique continuation
Classes, $r = 1.2$	reserve equals h minus epsV	9.95	1	0.000000	no	accepted continuation(s) found; uniqueness not established
Classes, $r = 1.2$	reserve equals h	10	1	0.000000	no	accepted continuation(s) found; uniqueness not established
Classes, $r = 1.2$	reserve equals h plus epsV	10.05	1	0.000000	no	analytically unique continuation
Classes, $r = 3$	reserve zero	0	1	0.655609	no	analytically unique continuation
Classes, $r = 3$	declared reserve 0.5	0.5	1	0.872367	no	analytically unique continuation
Classes, $r = 3$	reserve equals ell minus epsV	0.95	1	1.033719	no	analytically unique continuation
Classes, $r = 3$	reserve equals ell	1	1	1.018643	no	analytically unique continuation

Continued on next page

Table 6. Exact-event candidates in the exploratory reserve search (continued).

	Event	Reserve	Found	$\mathcal{R}_T$ max	Unresolved	Search outcome
Classes, $r = 3$	reserve equals ell plus epsV	1.05	1	0.999860	no	analytically unique continuation
Classes, $r = 3$	declared reserve 1.1	1.1	1	1.014500	no	analytically unique continuation
Classes, $r = 3$	ceiling equality p H	p_H (≈ 1.325270)	1	1.068854	no	analytically unique continuation
Classes, $r = 3$	reserve equals r	3	1	0.375000	no	analytically unique continuation
Classes, $r = 3$	floor equality p L	p_L (≈ 6.281718)	1	0.785215	no	analytically unique continuation
Classes, $r = 3$	reserve equals h minus epsV	9.95	1	0.000000	no	accepted continuation(s) found; uniqueness not established
Classes, $r = 3$	reserve equals h	10	1	0.000000	no	accepted continuation(s) found; uniqueness not established
Classes, $r = 3$	reserve equals h plus epsV	10.05	1	0.000000	no	analytically unique continuation

Outputs.

```
tables/reserve_comparisons.csv:
value_law, signal_information, epsilon_V, r, p,
t_0, t_H, t_L, g_H, g_L, Delta_T,
q_H, q_L, e_H, e_L, E, R_T,
low_cost_floor_margin, trading_margin, payoff_oracle_error,
status, accepted

numerics/reserve_continuations.csv (and numerics/reserve_events.csv):
value_law, r, p, branch, q_H, q_L,
E, O_H, R_T, no_entry_price_mass, posterior_in_no_entry_pool,
epsilon_P, epsilon_e, epsilon_q, status, accepted, unresolved_reason,
p_exact, regime, band_status, low_cost_floor_relation, ceiling_relation,
duplicate_of, identity_note,
candidate_id, continuation_id, parameter_set_id, institution_id,
information_structure_id, order_rule_id, price_rule_id,
pricing_family_coverage, exhaustive_pricing_search,
pool_id, pool_set_definition, pool_cutoff_exact, price_atom_value,
state_H_pool_mass, state_L_pool_mass, pool_probability, pool_posterior,
preparation_rule_id, tie_rule_id, event_id, event_defining_relation,
preparation_probability, admissible_challenger_probability, sale_probability,
two_admissible_bidders_probability, high_value_ownership_probability,
seller_revenue, mean_financial_price,
pricing_error, belief_error, entry_deviation_gain,
investor_deviation_gain_estimate, investor_deviation_gain_upper_bound,
error_budget, result_status, existence_scope, uniqueness_scope,
search_coverage_scope, rejection_reason, run_id
plus the shared error-budget columns in C.0

numerics/reserve_ranges.csv (and numerics/reserve_event_ranges.csv):
value_law, r, p, accepted_continuations_found,
E_min_found, E_max_found, R_T_min_found, R_T_max_found,
search_unresolved, global_envelope_certified,
p_exact, event_id, candidates_evaluated, candidates_accepted_raw,
candidates_rejected, candidates_unresolved, duplicates_merged, search_outcome
```

Data dictionary for the continuation schema. `p` is the exact declared decimal, or the thirty-two-digit decimal of an event value; `p_exact` is the declaration or the event's defining relation with the primitives it uses. `regime` is the reserve's position on the full payoff domain; `band_status` says whether a class reserve lies below, inside, or above each value band. `low_cost_floor_relation` and `ceiling_relation` take the values strict positive, strict negative, equality event, or unresolved, from the fifty-digit evaluation. `candidate_id` is the attempt identity and `continuation_id` the economic identity defined above; `duplicate_of` is empty or the representative's candidate identity;

identity_note records same-orders-different-atoms cases. parameter_set_id is the canonical string of exact declarations (not a hash); institution_id, information_structure_id, and tie_rule_id are the literal definitions of the sale rule, of who observes what, and of the tie conventions. order_rule_id is the canonical order profile; price_rule_id names the pricing family, its pool sets, and the positive-entry price map. pricing_family_coverage and exhaustive_pricing_search are as declared in C.0. pool_id, pool_set_definition, pool_cutoff_exact, price_atom_value, state_H_pool_mass, state_L_pool_mass, pool_probability, and pool_posterior describe the zero-entry atom when one exists and are n/a otherwise; a positive-price plateau atom in a flat posterior tail is a different kind of atom and is not written as a no-entry pool. preparation_rule_id is the buyer's action by price region and cost. event_id and event_defining_relation identify the node on the event grid. preparation_probability, admissible_challenger_probability, sale_probability, two_admissible_bidders_probability, and high_value_ownership_probability are $\mathsf{E}, \mathsf{A}, \mathsf{S}, \mathsf{C}_2, \mathsf{O}_H$ of C.0; seller_revenue is $\mathcal{R}_T$ and mean_financial_price is $\mathbb{E}P$, which must agree with it. pricing_error, belief_error, and entry_deviation_gain are the pointwise price, pooled-belief, and buyer-optimality diagnostics; investor_deviation_gain_estimate is the maximal gain found on the declared and refined order grids, and investor_deviation_gain_upper_bound is a rigorous bound when one exists (zero when an analytical bound covers the whole order interval) and n/a otherwise, never zero by default. error_budget is the quadrature plus tail term. result_status uses the vocabulary of C.0; existence_scope, uniqueness_scope, and search_coverage_scope state, in words, what the row establishes and what was searched. rejection_reason carries the breach witness of a rejected candidate; unresolved_reason the reason an open row is open. run_id names the producing exercise. In the range file, search_outcome is one of the four outcomes listed under Reporting, and global_envelope_certified is false unless an additional exhaustive argument establishes the full envelope.

Feed: Table 4 and the sale-design discussion. The research question about seller-optimal discovery is not assigned an answer by this exploratory file.

C.6b. Price pooling at fixed orders

This exercise is the regression for A.11. It runs the declared cutoff family at the price-pool reserve and the negative controls, and it is the smallest test that the continuation representation above is right.

Inputs. The benchmark declaration with the reserve replaced by the declared pool_reserve, at strength 1.2, full orders $(1, -1)$, atomic costs, Laplace noise, and the cutoff grid declared in C.0. The auction objects are computed in exact rational arithmetic and by direct integration of (OA.50), and both must give $t_0 = t_L = g_L = 0, t_H = 7, g_H = 3$.

Method. For each cutoff, build the complete pooled schedule (OA.58) with its prescribed preparation policy and validate it as a continuation. The pool masses $\int_{-\infty}^{c} a_\theta$ are computed in fifty-digit arithmetic from (OA.59) and by segment-wise Gauss-Legendre integration, and the two must agree within the independent-formula acceptance. The pooled posterior is (OA.60) from the full preimage. Buyer optimality is checked at the zero-price atom against the pooled posterior for both cost types, and pointwise on the positive-price region on a flow mesh that starts at the cutoff and includes every breakpoint; the row stores the minimum slack in each case. Competitive pricing (OA.61) is checked pointwise on a mesh spanning both regions and including the cutoff from both sides, together with the residual identities (OA.62) against direct conditional payoffs; the collision test requires a positive gap between the positive-price range and the zero atom. The upper plateau is recorded as a positive-entry atom with its own masses and posterior. Investor deviations are scanned on the declared and refined order grids for both signs and both types against the fixed schedule. The global bound (OA.63) is evaluated in exact rational arithmetic with an outward-rounded $e^{2/b}$, and the row stores the rational lower bounds on J_c and on U'_θ; the numerical J_c from both types and the minimum of U'_θ on a two-hundred-point mesh are stored beside them as the implementation check, and a numerical value below the rational bound is a breach. The output quantities (OA.64) are computed in closed form and by conditional integration; both are compared with each other and, at the two endpoints, must agree to ten decimals with the values written into the manifest as landmarks.

Acceptance. Every grid member passes every validation. Preparation and revenue strictly decrease along the grid. The rational bounds satisfy $J_c \geq 7/32$ and $U'_\theta \geq 143/1600$. The two endpoints survive deduplication as distinct continuations, a repeated construction of one endpoint from a second start merges into it with a different `candidate_id` and the same `continuation_id`, and the orders-only key merges the two endpoints, which the ledger records as the reason that key is forbidden. The controls $c = -1$, $c = 1$, and the raw-posterior buyer inside the $c = 0$ pool fail, and each fails with the breach named in A.11. A control that does not fail as required stops the exercise.

Outputs.

```
numerics/price_pool_regression.csv:
h, ell, p, rho, c_L, c_H, b, k, r, noise, cost, prior_H, q_H, q_L,
t_0, t_H, t_L, g_H, g_L, cutoff_exact, cutoff_decimal,
pool_probability, state_H_pool_mass, state_L_pool_mass, pool_posterior,
buyer_slack_zero_price_c_L, buyer_slack_zero_price_c_H,
buyer_slack_positive_price_c_L_min, buyer_slack_positive_price_c_H_min,
global_bound_J_lower_rational, global_bound_Uprime_lower_rational,
J_numeric, min_dU_mesh,
preparation_probability, sale_probability, admissible_challenger_probability,
two_admissible_bidders_probability, high_value_ownership_probability,
seller_revenue, mean_financial_price, closed_form_vs_integration_error,
pricing_error, entry_deviation_gain, investor_deviation_gain_estimate,
investor_deviation_gain_upper_bound, validation_status, accepted,
expected_failure_reason, candidate_id, continuation_id, run_id
```

`cutoff_exact` is the symbolic cutoff and `cutoff_decimal` its thirty-two-digit expansion. Negative controls and the duplicate start appear with `accepted = false` and an `expected_failure_reason`; they are not accepted rows. Feed: the registry keys `pool_*` of C.8, Proposition A.10 in the paper appendix, and A.11.

C.6c. Exact reserve events and admissible-bid outcomes

This exercise runs the C.6 node solver, with the continuation identity and event handling above, on the event grid alone. It leaves the full sweep untouched and exists so that the event candidates and the outcome measures can be checked without the hour-long scan.

Inputs. Both economies at both strengths, on the event grid of C.6: 0, ℓ, r, h, the band edges in the class economy, the declared reserve alternatives, p_L at every strength and p_H where $\ell < p_H < r$, each with the three one-sided offsets on both sides, and the reserve `6.280` used as the sampled weak maximum in the previous sweep.

Method and acceptance. At p_L in the weak economy, the row must show the equality-event classification of the floor, a single accepted full-order continuation whose status names the equality and the tie rule rather than the strict inequality, $\mathsf{E} = \rho$, sale probability $\rho/2$, $\mathsf{C}_2 = 0$, and $\mathcal{R}_T = \rho p_L/2$ from both the row and a fifty-digit evaluation; the three offsets below the event must classify the floor as strictly positive and the three above as strictly negative. At p_H in the strong economy, the row must show the equality-event classification of the ceiling with a strictly positive floor and a strictly positive full-order bound, $\tau = M$, $x^* = 1$, $\alpha_H = 1/2$, $\alpha_L = e^{-2/b}/2$, and entry and revenue agreeing with the fifty-digit evaluation; where p_H falls outside its defining regime the node is an ordinary reserve node and must not be classified as an event. At the sampled reserve `6.280`, $\mathsf{E} = \rho$, sale $\rho/2$, $\mathsf{C}_2 = 0$, and $\mathcal{R}_T = \rho\,(6.280)/2$: preparation there does not mean two admissible bidders, because the incumbent is excluded. On every row the union identity for S and the ordering $0 \leq \mathsf{C}_2 \leq \mathsf{A} \leq \mathsf{E} \leq 1$ hold, and the payoff oracle error is within tolerance. Deduplication follows the C.6 rule, and every count in the manifest comes from the ledger.

Outputs. `numerics/reserve_events.csv`, with the continuation schema of C.6 and an added `event_classification` text column, and `numerics/reserve_event_ranges.csv` with the range schema of C.6. Feed: the exact-event paragraphs of the paper appendix and the online summary of C.6. Event candidates improve sampled maxima; they prove neither global optimality nor envelope completeness, and a node without an accepted candidate is a node where the declared searches found none.

C.7. Bargaining weights and the payment property

Inputs. Use the benchmark h, ℓ and weak/strong incumbent distributions, but set the bargaining institution's reserve to zero as specified in Proposition A.8. Evaluate seller weights from `0` to `0.99` by `0.01`, including the exact midpoint `0.5`. The endpoint $\eta = 1$ can be stored separately as a payment-stage limit, not as a positive-cost entry equilibrium.

Objects and verification. Compute $T_\eta(R, v)$ and the winning challenger's profit pointwise, then integrate over the incumbent. Independently compare the results with (OA.44), including no-entry revenue $\eta r/2$. Verify feasibility of each transfer between fallback and winning value and the Nash first-order condition for interior weights. Compute differences between strong and weak distributions for Δ_η and $G_{\theta,\eta}$. The spread difference must change sign at the midpoint, with zero at that exact weight. This is an acquisition-stage exercise; do not append a trader or entry prediction without solving the new continuation game.

Outputs.

```
figures_data/bargaining.csv:
eta, r, t_0_eta, t_H_eta, t_L_eta, Delta_eta,
G_H_eta, G_L_eta, spread_strength_difference,
profit_H_strength_difference, profit_L_strength_difference,
transfer_feasibility_error, integration_error, status
```

Feed: Figure 4 and Proposition A.8. Parameter values and the change of institution are stated in the figure caption.

C.8. Quantity definitions and provenance

Every quantitative placeholder in the two manuscripts resolves to one row of the quantity registry, generated from the quantity manifest and the validated exercise outputs in C.1 to C.7. The manifest specifies each key's definition, producing exercise, source file, row selector, display rule, units, and exact value for a declared input. The registry records the resulting value, outward bounds for an enclosure, evidence status, and the source row used. The complete key list and definitions are in the distributed quantity dictionary, generated from the same manifest. This subsection states their common rules.

Rules. A declared input is echoed with status `input` and printed as its exact decimal. A derived quantity is filled only from an accepted source row selected by the full parameter declaration, information regime, cost and noise law, continuation identity, and branch label, never by a rounded scalar. The selector must match exactly one row. A difference records both rows used to form it. The row's result status must permit the statement the manuscript makes with it: an analytical claim needs an analytical row, a certified interval needs every certificate predicate, and a numerical diagnostic is printed as one. An unresolved source leaves the placeholder unfilled with its reason, and the substitution step refuses to produce a filled manuscript in that state. Mathematical constants, equation and result labels, bibliographic identifiers, and schema labels are literal text, not registry quantities. No renderer replaces a registry value with a cached or typed number.

Units and rounding. Probabilities are stored as probabilities and displayed by `decimal_6`; `percent_integer` converts a declared accuracy to a percentage at display time only. Margins and bounds use `scientific_10`; thresholds and standardized distances use `decimal_9`. Enclosures use `outward_interval_8` or `outward_interval_10`, rounding the lower endpoint down and the upper endpoint up; one-sided bounds use `lower_bound_10`, `lower_bound_12`, or `upper_bound_12`, rounding conservatively. Rounding is applied after every validation on the unrounded value. A displayed lower bound that supports a strict inequality must remain strictly positive after rounding; if outward rounding would print a zero, more digits are printed rather than rounding inward. Model units are the per-share normalization of A.1.

Quantity families.

Quantity group	Exercise	Units
Declared inputs	C.0	exact input
Auction payoffs	C.1	model units
Equilibrium and control outcomes	C.1	probability; model units
Welfare and revenue differences	C.1	model units
Sufficient-condition margins	C.1, C.3, C.4	model units
Thresholds and posterior bounds	C.2	model units; probability
Distributional and moderate examples	C.1, C.4, C.5	probability; normalized distance
Complementary signals	C.3	probability; percentage
Certified asymmetric nodes	C.2	order; probability; model units
Class-economy reserve comparison	C.6	probability; model units
Price-pool endpoints	C.6b	probability; model units

The basic definitions the families share are $\mathsf{E} = (\bar{e}_H + \bar{e}_L)/2$, $\mathsf{O}_H = \bar{e}_H/2$ on the benchmark support and the allocation-event value elsewhere, $\mathcal{R}_T = t_0 + [\bar{e}_H(t_H - t_0) + \bar{e}_L(t_L - t_0)]/2$, the margins (OA.77) and (OA.29), the thresholds of A.5, the certificate predicates of B.6, and the outcome measures of C.0. Within a parameter set, `r_weak`, `r_strong`, and `r_collapse` select the declared strengths; the certified keys refer to the three declared nodes in order, not to roots chosen from a larger search; the reserve keys select on `value_law`, `epsilon_V`, `r`, and `p` exactly; the price-pool keys select on the symbolic cutoff. The minimum-margin keys are a convenience for acceptance and do not replace the recorded components. A figure is never the source of a scalar, and an apparent curve is never used to fill a missing certificate.

Production order. The benchmark, moderate-value, and complementary-signal inputs and margins come first, then the threshold boundaries, the benchmark entry, control, ownership, and welfare quantities, and the certified brackets; these determine which theorem regions the text can populate. The correspondence and reserve sweeps, the price-pool regression, and the event exercise follow. `numerics/check_registry.py` regenerates the registry in fresh processes under several ambient decimal precisions and import orders and requires identical output, so the printed values do not depend on the environment that produced them.

D. Institutional pilot and empirical design

The proposed pilot identifies publicly visible sale opportunities in which a prospective buyer can still decide whether to investigate while target shares trade. This appendix specifies eligibility, event coding, and the evidence needed to establish the decision interval. It reports a research design, without an extracted sample or estimated effects.

D.1. The decision interval and sample-selection rule

The pilot establishes the institution required by the model before attempting to estimate its causal mechanism. The unit of observation is a prospective buyer's participation decision within a publicly visible sale process for a listed target. A process is eligible when contemporaneous public information identifies a plausible acquisition opportunity, the target continues to trade, and at least one economically meaningful participation decision by a potential challenger remains open. Public visibility and the buyer's later decision must be separately dated.

A public strategic review, disclosed approach, or open bidding contest is a candidate starting point, not automatic inclusion. The public record must show that a prospective buyer could still decide whether to investigate or submit a substantive proposal. A transaction announced only after the buyer set was fixed does not qualify merely because the target traded during its earlier confidential negotiations. Likewise, a named firm mentioned by an adviser is a potential buyer, not an entrant; an NDA alone need not establish costly acquisition preparation. I retain each of these events but distinguish their economic content.

Selection proceeds without conditioning on whether a later bidder wins, whether a target return has a particular sign, or whether the process appears to support the theory. Construct a process-level screening log from the eligible filing and announcement universe, record inclusion and exclusion decisions and their evidence, and retain uncertain cases for adjudication. Multiple processes for the same target receive separate process identifiers when the record documents a termination and restart. No cases are selected and no sample count or estimate is reported in this design.

D.2. Information available at the time, not only at the filing date

I distinguish the date an event occurred from the date its content first became public. A definitive proxy can describe a private meeting retrospectively. That description supports the occurrence of the meeting, but does not put it in the stock market's information set on the meeting date. For every economically relevant event, record the occurrence date or bounded date range, the first verified disclosure date, the public source, and the specificity of the disclosed information.

The candidate modeled interval begins when the opportunity and relevant competitive threat are publicly interpretable and ends at the challenger's material participation decision. If the rival's identity, financing position, or proposed terms become public later, those facts enter the information set only then. An unverified rumor is coded as a rumor with its source and uncertainty, not as an established public bid. If the chronology provides only a month or a relative ordering, preserve that precision rather than inventing a daily date.

A process can contain more than one candidate interval. The pilot retains the chronology and identifies each interval explicitly; it does not select the interval producing the strongest return–entry association. Confidential diligence by the challenger may follow public visibility even when its identity remains private to the market. In that case the record must support an interpretation of what match-related information a public investor could hold, rather than presuming that the market knew the eventual entrant.

D.3. Sources and event coding

The primary chronology is the target's merger-background discussion in definitive proxy or tender-offer filings, supplemented by current-report filings, transaction exhibits, contemporaneous company releases, and time-stamped public reporting. SEC accession numbers and exact source locations accompany every coded factual claim. A filing's narrative is evidence, not an instruction to treat every party as having the same information. Contradictory dates or actor descriptions remain visible until resolved.

I use the following event categories in the pilot. An initial approach identifies who initiated contact and whether it was solicited. Public visibility identifies the first disclosure of the acquisition opportunity and the competitive facts revealed. Buyer contact records outreach and replies without equating them with entry. Confidentiality and data access record when information became available to the buyer. Diligence entry records a substantiated commitment to evaluation, including its stated scope. A first substantive proposal records price, consideration, financing, and diligence conditions. A revised proposal records what changed and which new information preceded it. Withdrawal records the actor making the decision and the reason when stated. Final selection records the board's choice, agreement, or process termination. Announced and actual deadlines are separate events.

The event ledger has a row for each event–actor link. The following schema is a data specification, not an extracted dataset:

```
pilot/processes.csv:
process_id, target_id, target_name, listing_status,
process_start_date_lower, process_start_date_upper,
public_visibility_date_lower, public_visibility_date_upper,
process_end_date_lower, process_end_date_upper,
lead_buyer_id, process_status, eligible, exclusion_reason,
information_interval_established, unresolved_issue
```

```
pilot/events.csv:
process_id, event_id, actor_id, actor_role, event_type,
event_date_lower, event_date_upper, event_date_precision,
first_public_date_lower, first_public_date_upper,
public_information_content, contemporaneous_or_retrospective,
price_or_range, consideration, financing_condition,
diligence_condition, binding_status, initiation_actor,
withdrawal_decision_actor, reported_reason,
accession_or_source_id, source_locator, verbatim_evidence,
source_publication_time, evidence_confidence, adjudication_status

pilot/intervals.csv:
process_id, interval_id, challenger_id,
interval_start_lower, interval_start_upper,
participation_decision_lower, participation_decision_upper,
public_opportunity_evidence, public_incumbent_evidence,
participation_evidence, information_complementarity_interpretation,
ordering_verified, eligibility_decision, unresolved_issue

pilot/adjudication.csv:
record_id, field, initial_value, alternative_value,
conflict_type, supporting_sources, resolution, resolver,
resolution_date, unresolved_reason
```

An absent field means missing, not applicable, or not stated, with a corresponding reason; it is not a zero price, an absence of a buyer, or a voluntary withdrawal. For price ranges, store both endpoints rather than the midpoint alone. A proposal's classification as formal or informal follows the source's description and the coded contractual features; no uniform price threshold or arbitrary document label is used to manufacture a distinction across processes. Quotations are checked against the source location, and summaries are stored separately from the evidence text.

D.4. What would establish institutional relevance

The pilot supports the timing if it documents a public opportunity, trading during that opportunity, and a subsequent participation decision that remained open. It supports the preparation margin if a buyer had to commit resources or obtain additional target information before submitting an executable bid. It supports an incremental-information interpretation if the relevant uncertainty was not already fully resolved by the buyer's public statements and the record leaves room for distinct target-side and buyer-side knowledge.

None of these observations alone establishes that prices caused entry. A target price rising before a competing bid can anticipate that bid. A private process described after announcement does not show that the public market knew it beforehand. A buyer's later success does not prove that it entered after receiving a favorable price signal. Such cases may illuminate institutions while remaining outside the mechanism's direct empirical evidence.

The first output is a sourced chronology and an accounting of eligible decision intervals, with exclusions and ambiguous cases retained. The next output is a short institutional section describing which elements of the modeled game occur in the verified intervals. The existing process illustration in the manuscript is not a substitute for this selection and coding exercise.

D.5. Measurement and eventual empirical tests

The participation outcome is commitment to investigation or substantive bidding, measured at the most precise date the evidence supports. Public bid counts are a separate, narrower outcome. High-quality ownership in the theory is a structural value concept; a winning challenger is observable, but winning alone is not evidence that its underlying acquisition value was high. Any empirical proxy for match quality must be separately justified.

Incumbent strength is measured using information available before the challenger's decision. A final offer reflects subsequent entry and cannot serve as a predetermined initial threat. Candidate descriptive measures include an initial disclosed proposal, pre-existing financing capacity, and publicly measurable buyer–target complementarity. Their interpretation must respect whether the measure changes competitive strength, common acquisition value, financing constraints, or several objects at once.

Returns, trading activity, and other information measures are timestamped before the outcome. Define the public-information cutoff before constructing financial windows, and retain alternative date bounds when the event time is interval-censored. A stock return is not itself a measure of how much information was revealed. The empirical design must distinguish a change in the expected level of sale proceeds from a change in the market's information about challenger quality.

A targeted causal design would shift information transmission without independently changing acquisition values, preparation costs, or financing. No such instrument is asserted here. Alternatively, a structural exercise could jointly model anticipated entry and learning from prices, using the final institution-specific theorem to identify discriminating restrictions. The pilot determines which route is institutionally credible before either exercise is undertaken.

E. Reproducibility and the numerical boundary

The replication package separates independent node verification, the numerical exercises, and presentation. This section describes their inputs, recorded environments, and reproducibility checks.

E.1. Distributed verification and execution environment

The verification directory contains independent node calculations, exploratory asymmetric searches, and interval certificates. The documented execution environment is Python 3.13.5, NumPy 2.3.5, SciPy 1.17.0, and mpmath 1.3.0. These are metadata for the distributed results, not a claim that every other compatible environment fails. Each new execution records its actual interpreter, package versions, operating system, arithmetic precision, parameter declarations, tolerances, and source-file hashes.

The numerical layer under `numerics/` implements the exercises in Appendix C, including the searches over incumbent strengths and reserves, with the coverage limits stated there. Accepted output rows supply the quantity registry, tables, and figures, and each exercise writes a run manifest under `numerics/manifests/` with its inputs, method, tolerances, software versions, the SHA-256 of every output, and its pass or fail result. The verification gate refuses any output whose hash no longer matches its manifest, so a rerun of the producing exercise is the only way to change a number. The empirical pilot in Appendix D remains a coding specification and produces nothing in this build.

Each displayed exercise has one producer, run from the repository root with the project interpreter. The output schemas in Appendix C identify its files, including both C.2 attempt ledgers. Commands are:

```
source .venv/bin/activate
export OPENBLAS_NUM_THREADS=1 OMP_NUM_THREADS=1 MKL_NUM_THREADS=1
export NUMEXPR_NUM_THREADS=1
PY=.venv/bin/python
$PY numerics/exercises/c1_baseline.py
$PY numerics/exercises/c2_correspondence.py --workers=20
$PY numerics/exercises/c3_signals.py
$PY numerics/exercises/c4_moderate.py
$PY numerics/exercises/c5_noise.py
$PY numerics/exercises/c6_reserve.py --workers=20
$PY numerics/exercises/c6b_price_pools.py
$PY numerics/exercises/c6c_reserve_events.py --workers=8
$PY numerics/exercises/c7_bargaining.py
```

The broad exercises use twenty workers on the replication VM, and the event exercise uses eight. The thread settings prevent numerical libraries from multiplying process concurrency. A smaller worker count changes runtime without changing the scientific grid or tolerances. Reduced diagnostic runs use `--quick` and a separate `--out` directory where supported. They do not replace the declared release corpus. The presentation layer then runs in this order:

```
$PY numerics/check_registry.py
$PY numerics/registry.py
$PY numerics/substitute.py
$PY numerics/render/render_all.py
$PY numerics/render/latex.py
$PY numerics/verify.py --final
```

Two entry points wrap these sequences. `make peer-release` validates the existing completed calculation outputs, runs the tests, regenerates presentation, typesets, runs the final gate, and assembles the deliverables after checking a matching visual-inspection record. For a new revision, the release script's `--build-only` option produces the files for inspection and `--package-only` validates and packages them afterward. `make peer-reproduce` runs the same declared sequence in a fresh working and output directory, without any accepted result carried over from a previous run, and compares the canonical numerical outputs, exact input declarations, certificate validity, resolved references, and extracted manuscript text with the release. The release runner fixes creation-date metadata with `SOURCE_DATE_EPOCH`; the comparison requires identical substantive content rather than relying on byte-identical PDFs. `replication/README.md` lists every producer command and the role of every file.

Registry calculations use a local decimal context with precision 60, so their serialized values do not depend on ambient arithmetic precision or import order; the registry check runs the generator in fresh processes at several ambient precisions and import orders and requires identical output. The final verification includes the registry manifest, numerical acceptance checks, and strict substitution. Source manuscripts remain editable; filled Markdown, LaTeX, tables, and figures are regenerated.

The `verification/` directory is the independent seed delivered before the numerical layer existed. The numerical exercises do not import its calculations. The independent certificate audit compares its separately rerun results with the local interval port. To rerun its checks in isolation, use:

```
python -m pip install -r requirements.txt
python review_checks.py --out results --mode core
python review_checks.py --out results --mode signals
python explore_asymmetric.py
```

Run these commands from an isolated copy of `verification/` and retain the distributed result files separately from newly generated results. For the certificate, run the following from the repository root:

```
bash audit/peer_polish/reference_seed/run_seed.sh
```

This guarded runner executes a byte-identical seed copy and refuses optimized Python or `PYTHONOPTIMIZE`, because the seed uses assertions for acceptance. Record a failure as a failure; do not change a tolerance or catch an exception solely to obtain an accepted output. The interval certificate requires outward interval arithmetic and its analytical cover. Replacing it with ordinary floating-point quadrature changes the evidentiary status.

E.2. Contents and scope of the distributed files

The core verifier computes auction expectations, theorem margins, full-order objects, threshold distinctions, moderate values, noise-law comparisons, and the specified class-value reserve alternatives. Its signal mode uses the conditional laws of Appendix A.7. The exploratory asymmetric search identifies candidate roots and checks numerical deviations. The certificate evaluator uses the antiderivatives and derivative cover of Appendix B to establish selected continuous-action equilibria.

The distributed directory contains:

```
verification/
  review_checks.py
  explore_asymmetric.py
  certify_asymmetric.py
  requirements.txt
  results/
    core_results.json
    deviations.csv
    two_signal_results.json
    two_signal_deviations.csv
    asymmetric_candidates.json
    asymmetric_interval_certificates.json
    certificate_run_log.txt
```

The core results hold primitive vectors, inequality margins, threshold calculations, candidate outcomes, and reserve comparisons. The deviation ledger records tested state–order alternatives and their payoffs. The signal results hold private-information margins, joint-probability checks, and the associated deviation ledger. The asymmetric candidate file is exploratory: its root and local-search output is not an exhaustive equilibrium enumeration. The interval file contains the root-sign enclosures, order-cover bounds, and entry intervals that complete the selected computer-assisted existence arguments. The certificate log provides the readable arithmetic output; the complete interval endpoints, not their displayed midpoints, are the numerical proof objects.

These files support the statements at their documented parameter points. They do not supply an optimizer over sale terms, a first-price bidding equilibrium, or the full set of mixed trading equilibria. The exercises in Appendix C extend numerical coverage while preserving those limits. Their results are reported separately from the independent node checks.

E.3. Implementation contract for the full exercises

The numerical layer is written as pure functions with explicit parameter arguments and immutable returned records. Its modules keep the layers apart: `auction.py` (payoffs on the full reserve domain, with the direct integration oracle), `noise.py` and `information.py` (densities, posteriors, price schedules), `deviations.py` and `quadrature.py` (unilateral-deviation payoffs with analytic Laplace tails), `search.py` (candidate search), `validation.py` and `continuations.py` (independent validation, price atoms, continuation identity, and the wide schema), `reserve_events.py` (exact events and regime classification), `certificates.py` (the interval port of Appendix B), `thresholds.py`, `signals.py`, and `render/` (figures and tables). `params.py` holds the exact decimal declarations of C.0 as strings and the `Controls` record of tolerances; exercises parse them as decimals and never through a binary float where a certificate depends on them. A search routine returns candidates and unresolved nodes; only the validation layer assigns `accepted`. No renderer solves an equilibrium, changes a parameter, or drops a branch.

The only boundary to the presentation layer is the validated CSV output specified in Appendix C, written through one writer that prints floats by their shortest round-trip representation and marks a missing value `n/a`. Tables are joined on the complete parameter declaration and branch or continuation label, never on a rounded strength or price. The extensions table is assembled from the distributional, moderate-value, and signal outputs, with fields that do not apply left as not applicable; a missing signal bound or an uncomputed welfare statistic is never filled with zero. The reserve table uses the declared comparisons; the found-continuation ranges and the event candidates are reported online under the heading fixed in C.6.

The scalar registry is generated from validated rows and exact input declarations. The substitution stage checks that every placeholder key is defined once, that its source row exists and is unique under the declared selector, and that its result status permits the intended statement. It fails before creating a filled manuscript if a required value remains open. Certificates are substituted as outward intervals and conservative lower bounds. Ordinary diagnostics carry their specified rounding without being promoted to certified enclosures.

Every exercise writes a manifest with inputs, method, tolerances, software versions, output hashes, and pass or fail outcomes; `numerics/verify.py` reads every manifest and fails on a missing pass, a changed hash, a failed revalidation of the declared nodes, a failed certificate, or an open placeholder. Failed searches, negative controls, duplicate starts, and arithmetic retries remain in the record with their reasons. A reader can therefore distinguish a rejected candidate from a numerical failure, an unresolved search from a nonexistence argument, and an accepted node from an exhaustive characterization.

E.4. Figure and table rendering

Figures have no internal titles. Captions state the economic comparison, parameter specification, information regime, units, and result-status interpretation. Multi-panel figures use the panel labels specified in the manuscript. The correspondence figure shows separate branches and breaks lines at unresolved nodes. Certified nodes retain interval bars even where the intervals are visually small. A uniqueness boundary is labeled by the theorem it satisfies, not by the first root found by a solver.

Tables preserve the difference between equilibrium comparisons and fixed-profile controls. A frozen-profile row can contain profitable investor deviations without contradicting its purpose; it is not labeled an equilibrium. Welfare columns distinguish target revenue from allocation surplus net of preparation costs. Reserve ranges are described as ranges across continuations found unless an exhaustive argument establishes more. No selected branch, matched mean price, or apparent hump replaces the underlying equilibrium checks.

Dow, James, Itay Goldstein, and Alexander Guembel. 2017. "Incentives for Information Production in Markets Where Prices Affect Real Investment." *Journal of the European Economic Association* 15 (4): 877–909. https://doi.org/10.1093/jeea/jvw023.