

OCTOBER 2026

“Reputation, Disclosure, and the Scope of Entry”

Rubik Khachatryan and Georgy Lukyanov

Reputation, Disclosure, and the Scope of Entry*

Rubik Khachatryan[†]

Georgy Lukyanov[‡]

Abstract

This paper studies how learning about an incumbent affects the scope of entry when competitive responses use resources shared across markets. An entrant chooses whether to launch in neither, one, or both of two markets. Entry into the second market reduces the incumbent's cost-reducing response in the first and can make one-market entry unattractive. The entrant learns about the incumbent's capability from a record of its response to an earlier rival. More frequent publication encourages a less capable incumbent to imitate a more capable one. An observed response then becomes less informative, and entry after that record expands. We compare publication of conduct with a public audit of capability. For an open set of parameters with uniform setup costs, full publication maximizes total surplus within the specified policy class when publication costs are low. Removing the interaction between response costs across markets reverses this choice, while preserving all early and singleton-market payoffs. A capability audit is dominated in both economies. Expected entry scope is constant across the considered policies within each technology, although productive investment and the allocation of entry change.

Keywords: reputation; entry scope; information disclosure; competitive response; multimarket competition.

JEL classification: D82, L13, L15.

*We thank Emiliano Catonini, Alla Friedman, Markus Gebauer, Vitalijs Jascisens, Ella Khromova, Steven Kivinen, Ekaterina Mitskevich, Konstantinos Shamruk, and Alexey Verenikin for helpful comments and conversations. Georgy Lukyanov acknowledges support from the French National Research Agency (ANR) under grant ANR-17-EURE-0010 (Investissements d'Avenir program). All remaining errors are our own.

[†]International College of Economics and Finance, HSE University, Moscow, Russia. Email: rakhachatryan@edu.hse.ru.

[‡]Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France. Corresponding author. Email: georgy.lukyanov@tse-fr.eu.

1 Introduction

A firm entering an established industry has to decide how widely to launch. Entry into one market requires a smaller initial expenditure, but allows the incumbent to concentrate its response. A broader launch is more costly and may also make that response more difficult to organize. For example, a transport operator opening services in several regions may face an incumbent whose operations managers, scheduling specialists, and technical staff are shared across regions. The return to entering an additional market then depends partly on what happens in the markets already entered.

The entrant may know the incumbent's fleet and facilities without knowing how costly it is to adapt its operations. Earlier competitive episodes provide information about this capability. However, an incumbent that expects future entry may respond to an earlier rival partly to influence what the next entrant believes. The interpretation of past conduct therefore depends on the incentive to produce it. In particular, making responses easier to observe can encourage less capable incumbents to imitate more capable ones.

This paper studies the interaction between these two considerations. We develop a two-period Cournot model in which an incumbent's cost of organizing a competitive response is private information and persists across periods. An early rival has already entered a separate, temporary market. The incumbent chooses whether to make a standardized investment that reduces its marginal production cost. A different firm subsequently chooses whether to enter neither, one, or both of two later markets. Before doing so, it observes an accurate record of the early action with a given probability. Its own private information concerns a sunk setup cost per market.

The central comparison concerns the technology connecting markets. With shared response costs, entry into the second market reduces the incumbent's investment in the first. The entrant gains both from serving an additional market and from facing a weaker response in the initial one. This can make a broad launch preferable to a narrow launch even when the second market is smaller. We compare this economy with one in which response costs are separable, preserving every early-response and singleton-market payoff. We show that removing the cross-market interaction can reverse the welfare-maximizing publication rule. Thus the way an incumbent organizes its response matters for the value of disclosing its conduct.

The argument proceeds in three steps. First, we solve the response and entry decisions. Both incumbent types invest after entry because investment earns current operating profits. Shared costs nevertheless reduce investment per market when the entrant opens both markets. If the incremental return to the second market exceeds the return to the first, no setup cost makes one-market entry strictly optimal. Firms choose either a broad launch or no entry. A fall in the perceived capability of the incumbent can eliminate the interval of costs supporting

a narrow launch. Individual choices may skip a scale, while aggregate entry probabilities remain continuous under atomless setup costs.

Second, we derive learning from the early episode. The capable incumbent benefits from the early response on current-profit grounds. The less capable incumbent incurs a current loss and responds only if the improvement in its reputation is sufficiently valuable. More frequent publication raises this incentive. Where the low type randomizes, an observed response becomes less informative about capability, and the later entrant chooses a broader launch after that record. The record is still accurate: the change comes from which types choose the recorded action. This is a conditional comparison and does not imply that disclosure raises aggregate entry.

Third, we evaluate disclosure in the full game. An early investment undertaken at a private loss can increase total surplus by expanding output. Publication affects this productive investment as well as the information available to the later entrant. We compare publication of conduct with an independent public audit that reveals capability before the early action. For an open set of primitives with uniform setup costs, full publication is optimal when publication costs are low. In the matched separable economy, no publication is optimal. A capability audit is dominated in both cases, even if it is costless. Under shared costs, the stronger continuation incentive generates enough additional early investment to outweigh the later allocation loss from publication. Under separable costs, it does not.

Within each technology, the expected number of markets entered is identical under all the policies considered. Welfare changes through early investment and the allocation of later entry across incumbent types and setup costs. The comparison is therefore not captured by mean entry scope. The theorem establishes a local result within a specified class of audit and publication rules; it does not characterize unrestricted information design or a firm's voluntary disclosure policy.

The transport example helps explain the timing. The early encounter could concern a seasonal service, and the later entrant learns from that separate experience. The early rival remains active after the incumbent responds; its market subsequently closes because the episode is temporary. The later firm is consequently neither a surviving early entrant nor a firm experimenting through a small launch of its own.¹

1.1 Related literature

The reputational link between competitive encounters originates in Kreps and Wilson (1982) and Milgrom and Roberts (1982); Fudenberg and Kreps (1987) study reputation against

¹The shared resource is the ability to organize and implement operating adjustments. Public physical capacity can be conditioned on throughout. The private information concerns the cost of using managerial and technical resources, rather than an unobserved fleet size.

multiple opponents. In a particularly relevant contribution, Srinivasan (1991) analyzes an incumbent’s cost signaling across markets, including entry undertaken for its signaling value. Here the later rival chooses its own scope, anticipating a profitable response by the incumbent after entry. Other connections between markets have different strategic effects: Feng et al. (2023) study uniform pricing as a barrier to entry, while Jost (2023) examines the commitment needed to make a fighter-brand launch effective. Our incumbent’s final response requires no prior commitment. This distinguishes the mechanism from pre-entry capacity investment (Spence, 1977; Dixit, 1980) and from complete-information dynamic predation (Rey et al., 2025).

Disclosure can affect entry and welfare even when the expected amount of entry is unchanged. Hwang and Kirby (2000) establish this point in a Cournot model with two incumbents and a privately informed potential entrant. More reliable cost disclosures change entry conditional on the signals, while leaving its ex ante probability constant under their uniform-cost and interiority restrictions (Proposition 1). Their welfare comparisons can favor either full or no disclosure. Our mean-scope identity has the same implication for the limits of entry counts as a welfare measure. The additional mechanism is productive reputational imitation followed by endogenous multimarket entry. In Hwang and Kirby, entry also reveals the entrant’s private production cost to incumbents. Our private setup cost is sunk before competition, so entry reveals no payoff-relevant information about the entrant’s operating technology.²

Endogenous investment and disclosure are closely connected in several related models. Jiang and Xiong (2025) study disclosure of costs determined by investment; Arya et al. (2025) examine disclosure with product-quality investment and investment spillovers. Ishinagi and Oh (2024) already combine investment observability with congestion costs across markets. In their model both firms operate in both markets, whereas ours makes the later rival’s participation and scope endogenous and links its beliefs to an earlier response. Atanasov (2024) studies disclosure incentives created by information spillovers across a firm’s segments. Our matched comparison instead varies the response-cost interaction while preserving early and singleton-market fundamentals. This identifies a particular source of a reversal in optimal publication, rather than a general claim that investment interactions can change the value of disclosure.

Finally, observability changes incentives to build reputation. This distinction between information and behavior is central to Prat (2005), Bar-Isaac (2012), and Bar-Isaac and Deb (2021); Li (2025) studies its role in policymaking and information provision. In a competitive-entry setting, Hoang and Ruckes (2017) link disclosure, risk management, and

²Hwang and Kirby evaluate domestic welfare, excluding the foreign entrant’s profit. We include all firms’ profits and consumer surplus. Their disclosure parameter changes signal precision about exogenous costs; ours changes the frequency of an accurate record of a chosen action.

the information available to a potential entrant. Our continuation incentive comes from the scope of entry against a persistent incumbent type. Work on airline competition likewise motivates attention to competitive history and anticipated entry (Goolsbee and Syverson, 2008; Kim, 2009; Sweeting et al., 2020). It does not establish the shared-resource mechanism studied here, which remains a theoretical explanation to be assessed in an application.

Section 2 presents the environment. Section 3 solves the competitive response and entry-scope decisions. Section 4 derives the reputation equilibrium, and Section 5 compares disclosure policies. Section 6 discusses implications and limits. Appendices A and B establish the policy comparison and verify the example. Appendices C–E report additional numerical results, the finite-market extension, and aggregation under nonuniform setup costs.

2 The model

There are three firms: an incumbent I , an early rival P , and a later entrant E . The early rival operates in a temporary market; the later entrant chooses its own launch in two other markets. The incumbent’s type $\theta \in \{H, L\}$ determines the cost of organizing a competitive response and remains the same across these encounters. Type H has the lower cost. The incumbent knows its type, and the other firms assign prior probability $\mu \in (0, 1)$ to H . Demand, base production costs, and physical capacity are public. The later entrant learns from a public record of the early action when that record is available. We first describe the later markets, whose entry decisions determine the incumbent’s incentive to build reputation.

2.1 The later markets

The later entrant can serve two markets, indexed by $i = 1, 2$. Inverse demand in market i is

$$P_i = c_0 + A_i - q_{Ii} - q_{Ei}, \quad A_1 > A_2 > 0, \quad (2.1)$$

and its marginal production cost is $c_0 > 0$. The entrant privately draws a per-market setup cost z from an atomless distribution F on $[0, Z]$, independently of the incumbent’s type and all observations. It chooses an entered set $S \subseteq \{1, 2\}$ and incurs real cost $|S|z$.

After observing S , the incumbent chooses investment $x_i \in [0, c_0]$ on each entered market. Investment lowers its marginal cost from c_0 to $c_0 - x_i$. These investments are observable before the firms choose quantities simultaneously. Their real cost is

$$C_\theta(x; S) = \frac{\chi}{2} \sum_{i \in S} x_i^2 + \frac{\gamma_\theta}{2} \left(\sum_{i \in S} x_i \right)^2, \quad \chi > \frac{8}{9}, \quad 0 < \gamma_H < \gamma_L. \quad (2.2)$$

The first term is the cost of implementing investment locally. The second is the cost of drawing on resources used in both markets: an increase in x_1 raises the marginal cost of x_2 , and conversely. We write $X_\theta(S) = \sum_{i \in S} x_{\theta i}(S)$ for total investment.

In the transport example, x_i is the cost reduction obtained by reorganizing service on a contested route. The same staff may implement adjustments on several routes even when vehicles are already allocated to them. A lower γ_θ means that organizing these simultaneous adjustments is less costly.³

Investment is an adaptation to competition and is available only on entered markets. An unentered market retains marginal cost c_0 and monopoly output $A_i/2$. This restriction rules out a prior monopoly investment decision. If the same cost reduction were available before entry, that additional decision would have to be included in the game.

The entrant's setup cost is sunk before the incumbent invests. It changes the willingness to enter but has no effect on demand or marginal production cost. Learning about z from the chosen scope therefore does not change the incumbent's operating response. This assumption separates the analysis from signaling by an entrant with private information about productivity or demand.

2.2 The early market and public records

The early rival has already entered a separate, one-period market with demand intercept A_P . The incumbent chooses between a standardized response R , reducing its marginal cost by y , and no response A . We assume

$$0 < y < \min\{A_P, c_0\}. \quad (2.3)$$

The response costs $(\chi + \gamma_\theta)y^2/2$, as implied by the one-market version of (2.2). As in the later markets, the rival's marginal cost is c_0 ; it observes the investment before Cournot competition. This observation within the early market is distinct from the public record available to the later entrant. The early rival's activity ends with this market episode. Its entry is given, and the binary response menu is a maintained restriction.⁴

With probability $\tau \in [0, 1]$, an independent verification process publishes an accurate record of the early action. Otherwise the record is empty. Publication is independent of the incumbent's type, its action, and the entrant's later cost draw. The later entrant observes this record and no other information about the early episode. An increase in τ therefore makes an accurate record more frequent; it does not improve the precision of a noisy signal.

A record may document whether a specified operating adjustment was implemented in a previous competitive episode. An empty record conveys no information because the same

³The quadratic specification yields an explicit solution. The relevant interaction is $\partial^2 C_\theta / \partial x_i \partial x_j = \gamma_\theta > 0$ for $i \neq j$. This increasing marginal cost is distinct from a binding physical-capacity constraint.

⁴A standardized response may correspond to implementing a specified operating package rather than choosing an arbitrary level of investment. The subsequent allocation across entered markets is continuous. We do not claim that the binary early choice solves an unrestricted continuous signaling game.

publication lottery applies to responses and absences of response. This interpretation requires that the later entrant cannot already infer the action from public prices, quantities, or other operating details. Independently verified records of private service contracts provide one possible setting. Freely observed and revealing prices would require a different information structure.

Resources become available again after the early episode. Only the organizational cost γ_θ persists. Early conduct thus affects later entry through beliefs, without using up resources needed in the later markets. The incumbent discounts later profit by $\delta \in (0, 1]$, and we use the same weight for welfare. All firms maximize expected profit. We study perfect Bayesian equilibrium: investment and quantity choices are sequentially optimal, entry is optimal after each record, and beliefs satisfy Bayes' rule at records with positive probability.

The timing is as follows. The publication rule is fixed before the incumbent learns its type. The incumbent then chooses its early response, and the firms in that market compete in quantities. Verification produces a public record. After observing the record and its own setup cost, the later entrant chooses its markets. The incumbent allocates investment across the entered markets, and quantity competition follows. In Section 5, we compare publication rules and allow an independent public capability audit before the early action. Table 1 collects the notation.

Table 1: Main notation

Symbol	Role and domain	Meaning
θ, μ, p	Type; beliefs in $[0, 1]$	Capability; prior and posterior probability of H
A_i, c_0	Positive primitives	Demand intercept above base marginal cost; base cost
χ, γ_θ	Positive primitives	Local and shared response-cost coefficients
A_P, y, δ	Early primitives; discount	Early demand, standardized response, and later payoff weight
z, F, Z	Private draw; distribution	Per-market setup cost and its support bound
S, k	Entry choices	Entered set and number of markets
$x_{\theta i}, X_\theta$	Response choices	Market investment and total investment
τ, ρ, w	Probabilities in $[0, 1]$	Verification; low-type response; overall response
$B_{\theta k}, \ell_{\theta k}$	Derived payoffs	Entrant operating profit; incumbent loss from entry
V_k, T_k, G	Derived choice terms	Expected profit; type exposure; one-market choice gap
$P_k, K, \bar{K}$	Derived quantities	Scope probability; conditional and unconditional mean scope
g_θ, κ	Derived early payoffs	Private response gain; low-type response cost
Λ_θ, D	Continuation values	Expected entry loss; low-type benefit from reputation
$h_\theta, \omega_\theta, W$	Derived welfare	Early response gain; terminal gain; expected terminal gain
α, a, c	Policy probability and costs	Capability-audit probability; unit audit and publication costs

3 Competitive responses and entry scope

3.1 The incumbent's response

We solve the game backwards. Suppose the later entrant has chosen a set S . For $x_i < A_i$, both firms produce positive quantities and Cournot competition gives

$$q_{Ii} = \frac{A_i + 2x_i}{3}, \quad q_{Ei} = \frac{A_i - x_i}{3}, \quad \Pi_{Ii} = \frac{(A_i + 2x_i)^2}{9}, \quad \Pi_{Ei} = \frac{(A_i - x_i)^2}{9}. \quad (3.1)$$

Investment raises the incumbent's output and lowers its rival's output. It also raises total output. Let $b_i = 4A_i/9$ and $d = \chi - 8/9 > 0$. After omitting the constant zero-investment profit, the incumbent's objective on the interior region is

$$\sum_{i \in S} b_i x_i - \frac{d}{2} \sum_{i \in S} x_i^2 - \frac{\gamma_\theta}{2} \left(\sum_{i \in S} x_i \right)^2. \quad (3.2)$$

Proposition 1. *For every entered set, the incumbent has a unique globally optimal response. If the following candidate satisfies $0 < x_{\theta i}(S) < \min\{A_i, c_0\}$ for all $i \in S$, it is that response:*

$$X_\theta(S) = \frac{\sum_{i \in S} b_i}{d + |S|\gamma_\theta}, \quad x_{\theta i}(S) = \frac{b_i - \gamma_\theta X_\theta(S)}{d}. \quad (3.3)$$

When the relevant candidates are interior, adding the other market strictly lowers investment in the initial market and strictly raises total investment. If both types' candidates are interior, the type with lower γ_θ invests more in each entered market.

Proof. The incumbent's operating profit has second derivative $8/9$ for $x_i < A_i$. If investment excludes the entrant, operating profit is $(A_i + x_i)^2/4$, with second derivative $1/2$. At $x_i = A_i$, profit is continuous and its first derivative falls from $4A_i/3$ to A_i . Subtracting local cost with $\chi > 8/9$ makes each market's net return strictly concave over the whole feasible interval. Subtracting the convex shared cost preserves strict concavity in the vector. Continuity on the compact investment set gives existence; strict concavity gives uniqueness. An interior stationary candidate is therefore globally optimal, including against exclusionary deviations.

The first-order conditions are $b_i - dx_i - \gamma_\theta X = 0$. Summing them gives (3.3). For $\gamma = \gamma_\theta$ and $i \neq j$,

$$\begin{aligned} x_{\theta i}(\{i\}) - x_{\theta i}(\{i, j\}) &= \frac{\gamma[(d + \gamma)b_j - \gamma b_i]}{d(d + \gamma)(d + 2\gamma)}, \\ X_\theta(\{i, j\}) - X_\theta(\{i\}) &= \frac{(d + \gamma)b_j - \gamma b_i}{(d + \gamma)(d + 2\gamma)}. \end{aligned} \quad (3.4)$$

The numerator is positive precisely because the pair response in market j is positive. Finally,

$$\frac{\partial x_{\theta i}(S)}{\partial \gamma_\theta} = -\frac{\sum_{j \in S} b_j}{(d + |S|\gamma_\theta)^2} < 0. \quad (3.5)$$

□

The response has a current-profit motive. At zero investment, its marginal return is $b_i > 0$, so the incumbent invests even in the final period. An entrant anticipates this profitable investment when deciding where to launch. Shared implementation costs determine how the response is spread across markets.

Define the entrant's total operating profit $B_\theta(S)$ and the incumbent's loss from entry $\ell_\theta(S)$ relative to the unentered monopoly markets. Under interiority,

$$B_\theta(S) = \sum_{i \in S} \frac{(A_i - x_{\theta i}(S))^2}{9}, \quad (3.6)$$

$$\ell_\theta(S) = \sum_{i \in S} \frac{5A_i^2}{36} - \frac{1}{2} \sum_{i \in S} b_i x_{\theta i}(S). \quad (3.7)$$

To obtain the second expression, monopoly profit is $A_i^2/4$, zero-investment duopoly profit is $A_i^2/9$, and the optimized value of (3.2) is $\sum_i b_i x_i/2$. Thus a change in market demand affects entrant profit, incumbent loss, and the response incentive jointly. Losses need not be positive for arbitrary primitives; the conditions used below require their positivity and ordering explicitly.

3.2 The entrant's scope decision

Throughout this section, maintain the interiority condition in Proposition 1 for both types and every nonempty entered set. Suppose the later entrant's posterior probability of high capability is p . Its first question is which market to choose if it enters only one. From (3.3),

$$B_\theta(\{i\}) = \frac{A_i^2}{9} \left(1 - \frac{4}{9(d + \gamma_\theta)} \right)^2. \quad (3.8)$$

Market 1 is therefore the preferred singleton for both types and at every posterior. We can write $S_0 = \emptyset$, $S_1 = \{1\}$, and $S_2 = \{1, 2\}$, abbreviating the associated profits and losses by $B_{\theta k}$ and $\ell_{\theta k}$. Quantities for $k = 0$ are zero.

The entrant chooses the largest maximizer of

$$U_k(p, z) = V_k(p) - kz, \quad V_k(p) = B_{Lk} - pT_k, \quad T_k = B_{Lk} - B_{Hk}. \quad (3.9)$$

The tie rule matters only for a zero-probability set of setup costs. The term T_k measures how much less the entrant earns against a high-capability incumbent when it chooses scale k .

Define the return to the first market, the incremental return to the second, and their difference:

$$a_1(p) = V_1(p), \quad a_2(p) = V_2(p) - V_1(p), \quad G(p) = a_1(p) - a_2(p). \quad (3.10)$$

The entrant chooses one market only if its setup cost is low enough to make the first market worthwhile and high enough to make the second unattractive. The next proposition states this comparison and its implications for changes in beliefs.

Proposition 2. *When $G(p) > 0$, the entrant chooses two markets for $z \leq a_2(p)$, one market for $a_2(p) < z \leq a_1(p)$, and no entry for $z > a_1(p)$. When $G(p) \leq 0$, it chooses two markets for $z \leq V_2(p)/2$ and no entry otherwise. All intervals are intersected with the cost support.*

If $0 = T_0 \leq T_1 \leq T_2$, chosen scope weakly increases as p falls, for every setup cost. For atomless F , each scope probability $P_k(p)$ and mean scope $K(p) = \sum_k kP_k(p)$ are continuous in p .

Proof. One-market entry beats no entry if $z \leq V_1$, and it beats two-market entry if $z \geq V_2 - V_1$. The larger-scale tie rule makes its selected interval $(a_2, a_1]$. This interval is empty when $G \leq 0$, in which case comparison of the two endpoints gives the cutoff $V_2/2$. At $G = 0$ the sole three-way tie is resolved in favor of two markets.

For the monotonicity claim, take $p' < p$ and two scales $j < k$. Moving to p' increases $U_k - U_j$ by $(p - p')(T_k - T_j) \geq 0$. If k was chosen at p , the smaller scale cannot become strictly preferable at p' . If equality remains, the tie rule also excludes a switch to j .

When $G > 0$, probabilities are $P_2 = F(a_2)$, $P_1 = F(a_1) - F(a_2)$, and $P_0 = 1 - F(a_1)$. When $G \leq 0$, they are $P_2 = F(V_2/2)$, $P_1 = 0$, and $P_0 = 1 - P_2$. Extend F by zero and one outside its support. At $G = 0$, both cutoffs equal $V_2/2$. Continuity of F and the intercepts then proves the claim, including at the change of regime. $\square$

The type ordering in Proposition 1 implies $T_1, T_2 > 0$. It does not imply $T_2 \geq T_1$, which is a separate condition for monotone scope. The example below satisfies this condition strictly. The fixed best singleton follows from (3.8); with more general market heterogeneity, the identity of the preferred portfolio can itself change with beliefs.

Write $\Gamma_\theta = 2B_{\theta 1} - B_{\theta 2}$. If

$$\Gamma_L < 0 < \Gamma_H, \quad (3.11)$$

then the one-market gap crosses zero at

$$p^* = \frac{-\Gamma_L}{\Gamma_H - \Gamma_L} \in (0, 1). \quad (3.12)$$

Above p^* , there is a nonempty one-market cost interval, provided the cutoffs lie in the support. Below it, one-market entry is never chosen. The gap shrinks because $G'(p) = T_2 - 2T_1 = \Gamma_H - \Gamma_L > 0$: a decline in perceived capability raises the value of two-market entry by more than twice the increase in the singleton value.

The economic reason is the return that a broader launch creates in the first market. A low-capability incumbent finds it particularly costly to respond in both markets. Entering the second market can therefore raise the entrant's profit in the first by enough to outweigh the smaller size of the second. Against a high-capability incumbent, this benefit may be weaker. Beliefs then determine whether a narrow launch remains attractive.

3.3 A matched separable benchmark

To isolate the role of the shared resource, replace (2.2) by

$$C_{\theta}^{\text{sep}}(x; S) = \frac{\chi + \gamma_{\theta}}{2} \sum_{i \in S} x_i^2. \quad (3.13)$$

This changes only the interaction between investments in different markets. Every singleton response and every early-market payoff remain unchanged. With two-market entry, each investment is now chosen as if the other market were absent. Entrant operating profits are additive, and

$$2B_{\theta}(\{1\}) - B_{\theta}^{\text{sep}}(\{1, 2\}) = B_{\theta}(\{1\}) - B_{\theta}(\{2\}) > 0. \quad (3.14)$$

The one-market interval is therefore nonempty at every posterior when the cutoffs are interior. Reputation can still shift both thresholds. What disappears is the additional return to broad entry from reducing the response in the first market. We will use exactly this comparison in the welfare analysis, recomputing the reputation equilibrium under each technology.

3.4 A running example

The same primitives will be used for the reputation and policy calculations:

$$\begin{aligned} A_1 = A_P = c_0 = d = 1, \quad A_2 = \frac{17}{20}, \quad \chi = \frac{17}{9}, \quad \gamma_H = \frac{1}{10}, \quad \gamma_L = \frac{1}{2}, \\ \mu = \frac{1}{4}, \quad \delta = \frac{19}{20}, \quad y = \frac{5}{8}, \quad z \sim U\left[0, \frac{1}{10}\right]. \end{aligned} \quad (3.15)$$

Market 2 has a demand intercept 15 percent below that of market 1. This difference determines both firms' operating payoffs and the allocation of investment. The example illustrates the mechanism and supplies a feasible economy for the disclosure comparison; it is not a calibration to a particular industry.

Table 2 shows the operating outcomes. All responses are interior. In either market, the high type invests more than the low type. For both types, entering the second market reduces investment in the first. In this economy, $\Gamma_L < 0 < \Gamma_H$, and $p^* \simeq 0.629059$.

Figure 1 relates the operating-profit comparison to entry. In its upper panels, the one-market value is initially above the chord between no entry and two-market entry. As perceived capability falls, that advantage shrinks and then disappears. The lower panels show the resulting intervals of setup costs. Once one-market entry is no longer optimal, firms either enter both markets or stay out.

Type	Markets	Investment	$B_{\theta k}$	$\ell_{\theta k}$	$\omega_{\theta k}$
H	1	0.4040	0.039463	0.049102	0.186439
H	2	0.3759, 0.3093	0.075763	0.097281	0.301066
L	1	0.2963	0.055022	0.073045	0.149920
L	2	0.2389, 0.1722	0.115408	0.153619	0.219690

Table 2: Operating outcomes under shared costs. Investment is listed in market order. Profits and incumbent losses follow from the same Cournot response; ω is the later surplus gain before setup costs, defined in Section 5.




Fig. 1 Returns to entry and the resulting scope choice. The upper panels plot operating profit $V_k(p)$ against the number of markets; the dashed chord joins zero and two-market entry. Labels report $G = 2V_1 - V_2$: only $G > 0$ permits one-market entry. The lower panels show the selected scope, enlarging the relevant setup-cost interval from the full support $[0, 0.1]$. Filled endpoints implement the largest-scale tie rule. All three posteriors are reached along the same verification equilibrium

For a finite belief change, some firms may switch directly from no entry to two markets. This requires comparison of choices against every alternative at both endpoints. The precise switching set for a finite market menu is given in Appendix D. No positive mass has to move at a single common posterior: the cost intervals expand and contract continuously as the belief changes.

4 Reputation and disclosure

The later entrant would like to know which incumbent type it faces. The incumbent would like to be regarded as capable, since that reputation discourages entry. We now ask when the low type is willing to pay for an early response to obtain this reputational benefit, and how its incentive changes with publication.

Assumption 1. For every nonempty entered set and both types, the candidate in (3.3) is interior. The derived payoffs satisfy $0 < \ell_{\theta_1} < \ell_{\theta_2}$ and $0 < T_1 < T_2$. The cost distribution has a continuous, strictly positive density on $[0, Z]$. For every $p \in [0, 1]$, the pairwise cutoffs $V_1(p)$, $V_2(p) - V_1(p)$, and $V_2(p)/2$ lie in $(0, Z)$.

The assumptions give reputation a continuation value. Entry reduces incumbent profit, a broader launch causes a larger loss, and higher perceived capability deters entry. The support restriction ensures a positive mass of firms changes its entry decision when beliefs change. These conditions are restrictions on the product-market primitives, rather than additional free continuation payoffs; Appendix B verifies them for the example.

Lemma 1. *Under Assumption 1, the expected continuation loss*

$$\Lambda_\theta(p) = \sum_{k=0}^2 \ell_{\theta_k} P_k(p) \quad (4.1)$$

is continuous and strictly decreasing on $[0, 1]$.

Proof. Continuity follows from Proposition 2. A lower p weakly raises entry scope at every cost, and losses are strictly ordered by scope. The largest cost at which any entry occurs is $\max\{V_1(p), V_2(p)/2\}$. Both terms are strictly decreasing in p , so this cutoff strictly rises when p falls. The interval between the old and new cutoffs has positive probability and changes from no entry to positive entry. Expected losses therefore strictly rise. $\square$

The incumbent's current gain from choosing the standardized response in the early market is

$$g_\theta = \frac{4A_P y}{9} - \frac{d + \gamma_\theta}{2} y^2. \quad (4.2)$$

Suppose $g_H > 0 > g_L$, and write $\kappa = -g_L > 0$. The high type benefits from the response immediately. The low type would prefer not to respond in the absence of a reputational gain. In either case, the early rival continues to produce a positive quantity: after a response it supplies $(A_P - y)/3$ and earns $(A_P - y)^2/9$.

Consider strategies in which H responds and L responds with probability ρ . The overall response probability is $w = \mu + (1 - \mu)\rho$. Bayes' rule gives

$$p_R = \frac{\mu}{w}, \quad p_A = 0, \quad p_\emptyset = \mu. \quad (4.3)$$

The no-response record reveals the low type when it occurs. An empty record does not change beliefs because verification is independent of the action and type. At $\tau = 0$ the two published records are impossible; at $\tau = 1$ the empty record is impossible. Beliefs at such records do not affect the outcome.

Define $D(p) = \Lambda_L(0) - \Lambda_L(p)$. This is the low type's reduction in expected entry loss when its reputation is p instead of zero. Its total gain from responding is

$$-\kappa + \tau \delta D(p_R). \quad (4.4)$$

The low type pays κ to make the early response. If the action is published, a response prevents the later entrant from identifying it as the low type, saving expected entry losses $\delta D(p_R)$. If no record appears, the entrant keeps the prior after either action. Only the published-record branch therefore matters for the incentive comparison.

Proposition 3. *Maintain Assumption 1, $g_H > 0 > g_L$, and*

$$\delta D(\mu) < \kappa < \delta D(1). \quad (4.5)$$

Set $\tau_0 = \kappa/[\delta D(1)]$ and $\tau_1 = \kappa/[\delta D(\mu)]$, so that $0 < \tau_0 < 1 < \tau_1$. There is an equilibrium with H responding surely and

$$\rho(\tau) = 0 \quad \text{for } \tau \leq \tau_0, \quad \tau \delta D(p_R(\tau)) = \kappa \quad \text{for } \tau > \tau_0. \quad (4.6)$$

Above τ_0 , the low type mixes with a unique $\rho \in (0, 1)$. Its response probability is continuous and strictly increasing in τ , and the response posterior is continuous and strictly decreasing. If, in addition,

$$g_H > \delta[\Lambda_H(0) - \Lambda_H(1)], \quad (4.7)$$

this equilibrium outcome is unique, apart from zero-probability entry ties and beliefs at impossible records.

Proof. As ρ rises from zero to one, p_R falls continuously from 1 to μ . By Lemma 1, the gain $D(p_R)$ strictly falls. At $\rho = 0$, (4.4) is nonpositive precisely when $\tau \leq \tau_0$. At $\rho = 1$, it is strictly negative even for $\tau = 1$, by the left inequality in (4.5). Hence there is a unique interior indifference point when $\tau > \tau_0$, and no positive mixing probability at $\tau = \tau_0$. The high type prefers response on this branch because its current gain is positive and the response record induces lower continuation losses than the no-response record. These strategies and Bayes' rule support the stated equilibrium.

The strict monotonicity and continuity of D give a continuous inverse. An increase in τ requires a decrease in $D(p_R)$ to preserve indifference. Thus p_R falls and ρ rises.

For uniqueness, the greatest possible continuation disadvantage to H from responding is at most $\delta[\Lambda_H(0) - \Lambda_H(1)]$. Condition (4.7) therefore makes response strictly optimal for H

under any continuation beliefs. Pooling on response cannot be an equilibrium: the response posterior would be μ , and even a deviation belief $p_A = 0$ would provide the low type a reputation gain below κ . Pooling is thus excluded, and the no-response record reveals L whenever verification has positive probability. The scalar argument then determines the low type's strategy. At $\tau = 0$, the current payoff signs directly determine both actions. $\square$

The upper bound on κ allows sufficiently frequent publication to induce imitation. The lower bound prevents the low type from always responding, even at $\tau = 1$. Hence publication first leaves the types separated and then induces mixing. The formal boundary $\tau_1 > 1$ records that pooling lies outside the feasible policy range.

Condition (4.7) is a sufficient uniqueness condition. We impose it in the running example so that the welfare comparison does not depend on choosing among equilibria. Without this bound, the proposition establishes the equilibrium in which the high type always responds but need not exclude other outcomes.

Corollary 1. *Under the conditions of Proposition 3, raising verification within the mixing range weakly increases entry scope after a verified response for every fixed z , and strictly raises its expectation. Each conditional scope probability remains continuous. If (3.11) holds and $p_R(1) < p^* < 1$, one-market entry disappears at*

$$\tau^* = \frac{\kappa}{\delta D(p^*)} \in (\tau_0, 1). \quad (4.8)$$

Proof. The posterior falls strictly along the mixing range. Proposition 2 then gives weakly greater scope at every cost, while the strictly moving participation cutoff in Lemma 1 gives a positive-mass set of strict increases. Continuity follows from composition. The crossing is unique because p_R is strictly decreasing, and substitution into the mixing equation gives (4.8). $\square$

More frequent publication changes the composition of incumbents behind a response record: a larger fraction are low types. The response consequently becomes less effective at deterring a broad launch. This comparison is across publication regimes. Within any one regime, a response still indicates higher capability than an observed absence of response.

For a fixed setup cost that leads to entry under both policies, expected investment per entered market also falls in the mixing range. At a fixed scope, a lower probability of H reduces expected investment. If scope rises from one market to two, investment in market 1 falls for each type; investment in market 2 is lower still because $A_2 < A_1$. This is a comparison for the same entrant cost, rather than an average over a changing set of entering firms.




Fig. 2 Publication, imitation, and entry after a response. The left panel shows the response posterior and the low type's response probability. The right panel shows the probabilities of entering zero, one, and two markets conditional on a verified response. Dotted lines mark $\tau_0 \simeq 0.523$ and $\tau^* \simeq 0.835$. The values shown at $\tau = 0$ are limits; a verified response has zero probability at that policy

4.1 Publication in the example

At (3.15), the private early gains are

$$g_H = \frac{145}{2304}, \quad \kappa = \frac{35}{2304}. \quad (4.9)$$

Both equilibrium restrictions and the high-type uniqueness bound hold strictly. The separation threshold is $\tau_0 \simeq 0.523077$. One-market entry disappears after a response at $\tau^* \simeq 0.834775$. At full verification, the response posterior is approximately 0.525122 and the low type responds with probability 0.301439.

The change is economically modest in this example. Between publication probabilities of approximately 0.655 and 1, mean entry scope after a response rises from 0.837 to 0.946 markets. The one-market probability starts at only about 1.46 percent and falls to zero. Figure 2 shows the whole path; Appendix C reports the numerical values and the costs that switch directly from no entry to two markets.

Both thresholds in Figure 2 change behavior continuously. At τ_0 , the low type begins to imitate. At τ^* , the one-market entry interval disappears. Neither creates a jump in the conditional probabilities.

4.2 Entry across all records

The conditional expansion need not imply an increase in entry across all histories. In the uniform-cost case, when the pairwise cutoffs are interior,

$$K(p) = \frac{V_2(p)}{Z}. \quad (4.10)$$

Indeed, if one-market entry is available, mean scope is $[G + 2(V_2 - V_1)]/Z = V_2/Z$. If it is unavailable, mean scope is $2(V_2/2)/Z$, giving the same expression. Since V_2 is affine in p and $w p_R = \mu$,

$$\bar{K}(\tau) = (1 - \tau)K(\mu) + \tau[(1 - w)K(0) + wK(p_R)] = K(\mu). \quad (4.11)$$

For the shared-cost example, mean scope is $138433/131220 \simeq 1.054969$ markets at every publication rate. The expansion after a response is offset by changes in the frequencies and composition of records. This is an affine-payoff cancellation, not a general consequence of Bayes' rule.⁵ The distribution of scope and the allocation of entry can still change. Appendix E gives examples in which nonuniform setup costs make the unconditional effect positive or negative.

5 Disclosure and welfare

We now compare two ways of making information available to the entrant. Publishing conduct preserves the incumbent's incentive to build reputation. Revealing capability before conduct removes that incentive, since the early response can no longer change what is known about the type. We evaluate these instruments by total surplus, taking account of both early investment and subsequent competition.

An information policy is chosen before the incumbent learns its type. With probability $\alpha \in [0, 1]$, an independent audit publicly reveals the exact type before the early response. Audit selection is independent of type, and whether an audit occurs and its result are public before conduct. If no audit occurs, the prior remains μ and the early action is published with probability τ . Expected real policy costs are $a\alpha + (1 - \alpha)c\tau$, with $a, c \geq 0$. We optimize over these two probabilities.⁶

The objective includes consumer surplus and all firms' profits, net of real investment and setup costs. A welfare-maximizing disclosure authority is one interpretation. We do not

⁵Hwang and Kirby (2000, Proposition 1) similarly obtain an invariant ex ante entry probability with uniform private production costs and interior entry cutoffs, despite changes in entry conditional on disclosure. Here the invariant object is the number of markets entered, and the distribution of records also responds to reputational imitation.

⁶An audit requires verifiable information about organizational capability beyond the action record. Partial signals, type-dependent audit selection, an audit announced only after conduct, and selectively edited action records are outside this policy class.

model its institutional incentives or the incumbent's voluntary choice of what to disclose. A consumer-surplus objective or different distributional weights could select a different rule.

5.1 Private and social response incentives

In a contested market with an interior response, consumer surplus is $(2A_i + x_i)^2/18$. Adding operating profits gives

$$CS_i + \Pi_{Ii} + \Pi_{Ei} = \frac{8A_i^2 + 8A_ix_i + 11x_i^2}{18}. \quad (5.1)$$

In the early market, moving from zero investment to y increases consumer surplus by $2A_P y/9 + y^2/18$ and changes the rival's profit by $-2A_P y/9 + y^2/9$. Their sum is $y^2/6$. The total-surplus gain from response is therefore

$$h_\theta = g_\theta + \frac{y^2}{6}. \quad (5.2)$$

The low type's privately costly response raises early total surplus whenever $\kappa < y^2/6$. In the example, $h_L = 115/2304 > 0$. The incumbent pays for investment but does not capture its full output benefit. Reputational imitation can therefore induce socially useful investment even though it lowers the incumbent's current profit.

For later entry, the surplus gain relative to the unentered monopoly markets, before setup costs, is

$$\omega_\theta(S) = \sum_{i \in S} \left[\frac{5A_i^2}{72} + \frac{4A_ix_{\theta i}}{9} + \left(\frac{11}{18} - \frac{\chi}{2} \right) x_{\theta i}^2 \right] - \frac{\gamma_\theta}{2} X_\theta^2. \quad (5.3)$$

Here monopoly surplus is $3A_i^2/8$, and response costs are already subtracted. Write

$$W_\theta(p) = \int [\omega_\theta(S^*(p, z)) - |S^*(p, z)|z] dF(z), \quad W(p) = pW_H(p) + (1-p)W_L(p). \quad (5.4)$$

The posterior determines the entry choice; the actual type determines the response and surplus. This distinction matters when the low type imitates the high type.

Up to constants common to the disclosure policies, gross welfare without an audit is

$$\begin{aligned} \mathcal{W}_C(\tau) &= \mu h_H + (1-\mu)\rho(\tau)h_L \\ &\quad + \delta\{(1-\tau)W(\mu) + \tau[(1-w)W(0) + wW(p_R)]\}. \end{aligned} \quad (5.5)$$

Under a full capability audit, each type acts on its current incentive: H responds and L does not. An action cannot change a type that is already known. Gross welfare is

$$\mathcal{W}_F = \mu h_H + \delta[(1-\mu)W(0) + \mu W(1)]. \quad (5.6)$$

Net expected welfare is

$$\mathcal{W}(\alpha, \tau) = \alpha(\mathcal{W}_F - a) + (1-\alpha)[\mathcal{W}_C(\tau) - c\tau]. \quad (5.7)$$

The affine dependence on α follows from the timing: audit selection is independent and publicly resolved before conduct.

5.2 A policy comparison

The benchmark in (3.13) leaves early incentives and every singleton-market outcome intact. Only the response to two-market entry changes. Nevertheless, that change affects expected entry losses and hence the value of reputation. The continuation game and the early mixing probability must therefore be solved separately for the two technologies.

Uniform setup costs make the policy comparison tractable. Within each entry region, $W(p)$ is quadratic and $D(p)$ is affine. The low type's indifference condition then gives a rational expression for welfare along the mixing path. Its curvature has a constant sign within each region, so endpoint derivative checks suffice to bound the derivative throughout that region. Appendix A gives the expressions, including both one-sided derivatives at a boundary.

The following proposition makes the sufficient conditions transparent. Put

$$A_F = \mathcal{W}_F - \mathcal{W}_C(0), \quad \Delta = \mathcal{W}_C(1) - \mathcal{W}_C(0). \quad (5.8)$$

In the separating range, publication changes information without changing early conduct, and

$$\mathcal{W}_C(\tau) = \mathcal{W}_C(0) + \tau A_F, \quad 0 \leq \tau \leq \tau_0. \quad (5.9)$$

Proposition 4. *Suppose the equilibrium in Proposition 3 is unique, $A_F < 0$, and welfare is continuous and piecewise continuously differentiable.*

- (i) *If $\Delta > 0$ and every mixing-interval derivative satisfies $\mathcal{W}'_C(\tau) > \Delta$, full publication is uniquely optimal for $0 \leq c < \Delta$ and no publication is uniquely optimal for $c > \Delta$. At $c = \Delta$, only the two endpoints tie.*
- (ii) *If the whole mixing segment is affine and $\Delta < 0$, no publication is uniquely optimal for every $c \geq 0$.*

In either case, the optimal audit probability is uniquely $\alpha = 0$ for every $a \geq 0$.

Proof. The separating segment is strictly decreasing, also after subtracting publication costs. In case (i), if $c < \Delta$, net welfare rises on every mixing interval. Full publication exceeds no publication by $\Delta - c > 0$, so it is the unique optimum. If $c \geq \Delta$, integrating the strict derivative bound from any mixing $\tau < 1$ to 1 gives

$$\mathcal{W}_C(\tau) - \mathcal{W}_C(0) < \tau \Delta.$$

Every interior mixing choice is therefore worse than no publication after costs, as is every interior separating choice. Comparing the endpoints gives the stated cutoff and tie. In case (ii), net welfare is affine on each of the two segments. Its possible maximizing vertices are 0, τ_0 , 1; the latter two yield strictly lower net welfare than $\tau = 0$. Finally, $\mathcal{W}_F - a < \mathcal{W}_C(0)$, while the no-audit branch can always select $\tau = 0$. Equation (5.7) excludes every $\alpha > 0$ at an optimum. $\square$

Welfare change	Shared costs	Separable costs
Full conduct gain	0.009973596	-0.000660155
Early surplus component	0.011284351	0.001020148
Later surplus component	-0.001310755	-0.001680303
Full capability audit gain	-0.003309410	-0.001868596

Table 3: Welfare relative to no disclosure within each technology, before policy costs. The later component includes discounting. The comparison changes only the cross-market response interaction; it preserves early and singleton-market payoffs.

Theorem 1. *At the primitives in (3.15), and throughout an open neighborhood within the uniform-cost family, the following policies maximize total surplus over $(\alpha, \tau) \in [0, 1]^2$:*

- (i) *With shared costs, $\alpha = 0$ and $\tau = 1$ when $0 \leq c < c^*$, while $\alpha = 0$ and $\tau = 0$ when $c > c^*$. Only these endpoints tie at $c = c^*$.*
- (ii) *With the matched separable costs, $\alpha = 0$ and $\tau = 0$ for every $c \geq 0$.*

The policies are unique except for the stated endpoint tie, for any $a \geq 0$. At the displayed primitives, $c^* = \mathcal{W}_C(1) - \mathcal{W}_C(0) \simeq 0.009973596$ in the shared economy. The threshold varies with primitives in the neighborhood.

Proof. The continuation and equilibrium conditions hold strictly in both technologies, including (4.7); Appendix B reports the primitive calculations. Table 3 gives the welfare comparisons. Under shared costs, the derivative on the first mixing interval decreases from approximately 0.024725 to 0.023924. On the second mixing interval it is constant at approximately 0.024954. Each strictly exceeds the full-publication gain 0.009974, and $A_F < 0$. Proposition 4(i) applies.

Under separable costs, the one-market interval exists at every posterior. Consequently W is globally quadratic and D is linear through zero. Substitution into the mixing condition makes the welfare path affine on the entire mixing segment. The full-publication gain and the audit gain are both negative, so Proposition 4(ii) applies.

The exact expressions and derivative checks are in Appendix A. Interiority, incentive bounds, the ordering of region boundaries, and the finite derivative inequalities are all strict. They vary continuously with primitives while the region ordering is preserved. The separable affine property is structural. The same policy comparison therefore holds in an open neighborhood of the displayed economy. $\square$

The reversal comes from different returns to reputation under the same early investment incentives. With shared costs, low-type imitation begins when publication reaches about 0.523; with separable costs, it begins only at about 0.924. At full publication the low type




Fig. 3 Gross welfare effects of conduct disclosure. Low publication probabilities change information but induce no additional response, reducing welfare here. Once imitation begins, welfare rises. Full disclosure recovers the initial loss under shared costs but not under separable costs. Publication costs would subtract $c\tau$ from each curve. The optimum is established analytically, not by the plotted grid

responds with probabilities 0.3014 and 0.0273, respectively. The same productive early action is therefore induced much more often when later responses use a shared resource.

Table 3 separates this early gain from the later allocation effect. Full publication lowers later surplus in both economies. Under shared costs, the additional early investment more than compensates for that loss; under separable costs, it does not. Figure 3 also shows why an endpoint comparison alone would be insufficient. Welfare initially falls, before imitation starts. Proposition 4 establishes the global ranking of publication rules by bounding the entire mixing path.

A capability audit improves the entrant's information but need not improve total surplus. The entrant does not internalize its effects on consumer surplus and incumbent profit. Relative to no disclosure, the audit changes later entry while leaving early behavior unchanged; that reallocation reduces welfare in the example. Relative to full conduct disclosure, it also removes productive low-type imitation.⁷

Expected scope remains $K(\mu)$ for every (α, τ) within each technology. On both audit and no-audit branches, the posterior has mean μ , and (4.10) is affine. This does not make

⁷The later entrant itself benefits from the audit. Its expected net profit is highest under the capability audit, intermediate under full conduct disclosure, and lowest under no disclosure in both displayed economies. Appendix B explains this private value of information and gives the values.

mean scope equal across the two technologies. It shows why, within each one, the welfare comparison requires information about investment and who enters against which incumbent.

The theorem identifies a strict local comparison. It does not isolate the disappearance of one-market entry as a necessary condition for a policy reversal: the matched change in costs affects both imitation incentives and later allocation. The ranking also depends on the early response size. At $y = 0.620$, full publication is better than no publication under both technologies; at $y = 0.645$, it is worse under both. Appendix C reports these comparisons and the sensitivity checks.

6 Implications and limitations

The model suggests two considerations for a firm planning its launch. First, entering an additional market can improve profitability in a market already entered by weakening the incumbent's cost-reducing response there. A market-by-market calculation misses this benefit when implementation resources are shared. Second, past responses have to be interpreted in light of the environment in which they were observed. Within a fixed publication regime, a response is evidence of higher capability than an observed absence of response. Across regimes, more frequent publication encourages imitation and can lower the reputation attached to the same action.

These comparisons also guide an empirical application. Holding demand and capability fixed, the response in an existing market should fall when another contested market draws on the same resources. The disclosure comparison is harder to identify because the composition of incumbents behind a response record changes with the reporting regime. A correlation between publicity and entry scope would not, by itself, identify that effect. Variation in access to records would have to be distinguished from selection into the records.⁸

Several restrictions determine where these implications apply. The early encounter is temporary, resources become available again, and only organizational capability persists. A rival that remains active and uses the same resources would change the later investment problem. Durable early cost reductions would create a technological link in addition to learning. The transport interpretation is therefore closest to separate service episodes, with operating adjustments that are not already inferable from public prices and quantities. It does not identify a particular existing disclosure mandate.

The entrant's private information is also specific. Setup costs affect its willingness to launch but are sunk before subsequent competition. If a broad launch instead signals low production costs or high demand, the incumbent would learn from the entrant's choice and

⁸In the model, responses and absences of response are sampled at the same rate. Selective reporting of unusually aggressive behavior would generally make an empty record informative and change the equilibrium.

adjust its response. Hwang and Kirby’s entry model illustrates this additional learning channel. Our analysis isolates learning about the incumbent; it does not rule out entrant signaling as another explanation for a relation between scope and competitive response.

Some results extend beyond two markets. For a finite menu, the entrant can compare its best operating profit at each scale and subtract the per-market setup cost. Atomless costs still imply continuous scale probabilities. But the preferred portfolio at a given scale may change with beliefs. A switch between equally large portfolios can change incumbent losses for a positive mass of cost types, even if scale probabilities are continuous. Monotone scope also requires an ordering of belief effects across scales. Appendix D states these distinctions. The two-market specification avoids composition switches because market 1 is always the preferred singleton.

The equilibrium and conditional-entry results allow general positive setup-cost densities under Assumption 1. Uniformity is used for the explicit policy theorem: it gives piecewise quadratic welfare and affine mean scope. With nonuniform costs, unconditional scope can rise or fall with publication, as Appendix E demonstrates. Strict rankings of two fixed policies survive sufficiently small distributional perturbations, but the global endpoint optimum has been established here within the uniform family.

The main lesson is that the technology connecting competitive responses across markets affects both the return to a broad launch and the value of disclosing past conduct. Shared resources can make entry into the second market weaken the incumbent’s response in the first. They also change how much a less capable incumbent will invest to preserve its reputation. The matched comparison shows that these effects can reverse the optimal publication rule while leaving all early and singleton-market payoffs unchanged. Evaluating disclosure therefore requires more than measuring how much information is released or counting the expected number of entered markets.

A Welfare calculations and the policy certificate

This appendix gives the calculations underlying Theorem 1. All quantities follow from product-market primitives; the accompanying program evaluates the displayed expressions with rational arithmetic.

A.1 A common representation of the two technologies

For this calculation, introduce a parameter $\lambda \in [0, 1]$ and write

$$C_\theta^\lambda(x; S) = \frac{\chi + (1 - \lambda)\gamma_\theta}{2} \sum_{i \in S} x_i^2 + \frac{\lambda\gamma_\theta}{2} \left(\sum_{i \in S} x_i \right)^2. \quad (\text{A.1})$$

The shared and matched separable technologies correspond to $\lambda = 1$ and $\lambda = 0$. Every singleton cost equals $(\chi + \gamma_\theta)x^2/2$, independently of λ . Thus the interpolation preserves both singleton responses and the early gains. No monotonicity in λ is asserted.

Let $d_\theta^\lambda = d + (1 - \lambda)\gamma_\theta$ and $r_\theta^\lambda = \lambda\gamma_\theta$. On an interior entered set of size k , investment is

$$X_\theta^\lambda(S) = \frac{\sum_{i \in S} b_i}{d_\theta^\lambda + kr_\theta^\lambda}, \quad x_{\theta i}^\lambda(S) = \frac{b_i - r_\theta^\lambda X_\theta^\lambda(S)}{d_\theta^\lambda}. \quad (\text{A.2})$$

The formulas for B_θ and ℓ_θ remain those in (3.6) and (3.7). To compute ω_θ , replace χ in (5.3) by $\chi + (1 - \lambda)\gamma_\theta$ and replace its final γ_θ by $\lambda\gamma_\theta$.

A.2 Integration over setup costs

Write $\omega_k(p) = p\omega_{Hk} + (1 - p)\omega_{Lk}$ and maintain interior entry cutoffs with $F(z) = z/Z$. If $G(p) > 0$, entry occurs in both markets for $z \leq a_2 = V_2 - V_1$ and in one market for $a_2 < z \leq a_1 = V_1$. Direct integration gives

$$W(p) = \frac{\omega_1(p)G(p) + \omega_2(p)[V_2(p) - V_1(p)] - \frac{1}{2}\{V_1(p)^2 + [V_2(p) - V_1(p)]^2\}}{Z}, \quad (\text{A.3})$$

$$\Lambda_L(p) = \frac{\ell_{L1}G(p) + \ell_{L2}[V_2(p) - V_1(p)]}{Z}. \quad (\text{A.4})$$

If $G(p) \leq 0$, only zero and two-market entry are chosen, and

$$W(p) = \frac{\omega_2(p)V_2(p)}{2Z} - \frac{V_2(p)^2}{4Z}, \quad \Lambda_L(p) = \frac{\ell_{L2}V_2(p)}{2Z}. \quad (\text{A.5})$$

The expressions agree at $G = 0$. The welfare function is continuous and quadratic on each region, while the loss is continuous and affine. Global concavity of W is neither assumed nor needed.

A.3 Derivatives along the mixing path

Within one entry region, let $W(p) = w_0 + w_1p + w_2p^2$ and $D(p) = u + vp$, where $v > 0$. Define

$$A = \frac{\kappa}{\delta v}, \quad B = \frac{u}{v}, \quad M = \mu h_L, \quad N = \delta\mu[w_0 - W(0)]. \quad (\text{A.6})$$

The mixing condition gives $p_R = A/\tau - B$, so $A - B\tau = \tau p_R > 0$. Using $w = \mu/p_R$ in (5.5) yields

$$\mathcal{W}_C(\tau) = C + L\tau + \frac{M\tau + N\tau^2}{A - B\tau}, \quad (\text{A.7})$$

where

$$C = \mu(h_H - h_L) + \delta W(\mu) + \delta\mu w_2 A, \quad (\text{A.8})$$

$$L = \delta[W(0) - W(\mu) + \mu w_1 - \mu w_2 B]. \quad (\text{A.9})$$

Consequently,

$$\mathcal{W}'_C(\tau) = L + \frac{MA + 2NA\tau - NB\tau^2}{(A - B\tau)^2}, \quad (\text{A.10})$$

$$\mathcal{W}''_C(\tau) = \frac{2A(AN + BM)}{(A - B\tau)^3}. \quad (\text{A.11})$$

The sign of the curvature is constant within the interval. Its smallest derivative therefore occurs at one of the endpoints. The one-sided derivatives on both sides of a region boundary must be checked.

If the same quadratic welfare formula holds over all posteriors and $D(p) = vp$, then $B = N = 0$. Equation (A.7) is affine in τ . This is the case under separable costs at the stated primitives, because one-market entry remains available throughout $[0, 1]$. The same property holds in a neighborhood preserving that strict ordering. Under shared costs, it holds on the lower-posterior region only.

A.4 Exact checks

Under shared costs, the entry threshold and verification boundaries are

$$p^* = \frac{1703317}{2707724}, \quad \tau_0 = \frac{7502175450000}{14342386622903}, \quad (\text{A.12})$$

$$\tau^* = \frac{1998604930950000}{2394184073893711}, \quad \tau_1 = \frac{20667150000}{9839206981}. \quad (\text{A.13})$$

At full verification,

$$p_R(1) = \frac{5166787500}{9839206981}, \quad \rho(1) = \frac{4672419481}{15500362500}. \quad (\text{A.14})$$

The full-publication gain is

$$c^* = \frac{153171970047678689}{15357747657530880000}. \quad (\text{A.15})$$

The two welfare polynomials are

$$W(p) = \begin{cases} \frac{12716829293}{136048896000} + \frac{64371749}{2448880128}p - \frac{3454078723}{172186884000}p^2, & p \leq p^*, \\ \frac{502791135133}{5509980288000} + \frac{201049876073}{6667076148480}p - \frac{25962769156243}{1260494084322000}p^2, & p \geq p^*. \end{cases} \quad (\text{A.16})$$

The accompanying program constructs these polynomials by integration, rather than fitting them to sampled values. Substituting their coefficients and the affine losses into (A.10) gives the following strict rational bounds, with endpoint derivatives taken from within the indicated mixing interval:

$$\begin{aligned} 0.00997 < c^* < 0.00998, \quad -0.003310 < A_F < -0.003309, \\ 0.023924 < \min_{\tau \in [\tau_0, \tau^*]} \mathcal{W}'_C(\tau) < 0.023925, \\ 0.024953 < \mathcal{W}'_C(\tau) < 0.024954 \quad (\tau^* < \tau \leq 1). \end{aligned} \quad (\text{A.17})$$

Decimal endpoints in (A.17) denote terminating rationals. Thus these bounds prove the strict derivative comparison with c^* without reliance on floating-point tolerances or a policy grid.

Under separable costs,

$$\tau_0 = \frac{468885965625}{507219175976}, \quad \tau_1 = 4\tau_0, \quad \Delta = -\frac{132340599141833363}{200468979170570880000}. \quad (\text{A.18})$$

Here $-0.001869 < A_F < -0.001868$ and $-0.000661 < \Delta < -0.000660$. The globally valid welfare polynomial is

$$W(p) = \frac{22593147817}{220399211520} + \frac{84484984051}{10417306482000}p - \frac{826440139952}{78780880270125}p^2. \quad (\text{A.19})$$

Its mixing segment is affine, as shown above. These exact checks establish all conditions used in Proposition 4.

B Feasibility and the entrant's value of information

The primitives in (3.15) give interior, strictly positive responses below both the demand intercept and the base production cost. For example, under shared costs the two-market responses are $(203/540, 167/540)$ for H and $(43/180, 31/180)$ for L . The singleton responses in market 1 are $40/99$ and $8/27$. All quantities in Table 2 use these same primitives, including $y = 5/8$ for the early episode.

For each technology, $0 < \ell_{\theta_1} < \ell_{\theta_2}$ and $0 < T_1 < T_2$. The three pairwise entry cutoffs are interior at $p = 0$ and $p = 1$. Since each is affine, this verifies the support condition over the whole posterior range. Table 4 reports the remaining equilibrium bounds. The common $g_H = 145/2304 \simeq 0.062934$ exceeds the largest possible discounted reputational loss in both technologies; the stated equilibrium is therefore unique under Proposition 3.

Quantity	Shared costs	Separable costs
$\delta D(\mu)$	0.007232111	0.004108223
κ	0.015190972	0.015190972
$\delta D(1)$	0.029041549	0.016432892
$\delta[\Lambda_H(0) - \Lambda_H(1)]$	0.018305496	0.011046438
g_H	0.062934028	0.062934028

Table 4: Equilibrium incentive bounds at the displayed primitives. The first three rows verify that the low type separates at low verification and mixes at full verification. The final two verify the sufficient uniqueness condition.

The net profit of the later entrant, before learning its setup cost, is

$$U(p) = \int \max_{k \in \{0,1,2\}} \{V_k(p) - kz\} dF(z). \quad (\text{B.1})$$

This function is convex because it is an integral of maxima of affine functions. Full capability disclosure yields weakly greater expected U than any Bayesian signal about the same type. In the example the inequalities are strict, as shown in Table 5. This private value of information does not imply a corresponding ranking of total surplus, which also includes consumers and the incumbent.

Information policy	Shared costs	Separable costs
No disclosure	0.027823977	0.019896557
Full conduct disclosure	0.028094236	0.020207173
Full capability audit	0.028566971	0.020241980

Table 5: Expected net profit of the later entrant before its setup cost is drawn. Values are not discounted. These private information gains need not coincide with total-surplus gains.

Appendix C reports sensitivity to the early response size and numerical perturbations of the primitives. The open-neighborhood claim follows from the strict analytical inequalities above, rather than from those numerical diagnostics.

C Additional numerical results

This appendix reports the numerical detail accompanying the three figures above. The parameters, equilibrium, and policy comparison are those of the main example; these calculations add no new assumptions to the main example.

C.1 Entry after a published response

Table C.1 follows the shared-cost equilibrium at three publication rates. The initial one-market probability is only about 1.46 percent. Over the displayed change in publication, all setup costs in $(0.0425749, 0.0472948]$ switch from no entry to two markets, a set with probability about 4.72 percent. These are properties of the example, not estimates of market effects.

C.2 Sensitivity of the policy comparison

The neighborhood claim in Theorem 1 follows analytically from strict inequalities and continuity, together with the structural affine property under separable costs. As a numerical diagnostic, the program also changes each of nine primitives $(A_1, A_2, d, \gamma_H, \gamma_L, \mu, \delta, y, Z)$ by plus or minus 0.01 percent. All 512 corner economies pass the same certificate. At plus or

τ	p_R	ρ	P_0	P_1	P_2	K
0.654928	0.800000	0.083333	0.574251	0.014576	0.411173	0.836922
0.834775	0.629059	0.196559	0.547654	0.000000	0.452346	0.904691
1.000000	0.525122	0.301439	0.527052	0.000000	0.472948	0.945897

Table C.1: Conditional entry along the shared-cost equilibrium. All scope probabilities and mean scope refer to a verified response. Values are rounded to six decimals.

minus 0.1 percent, 482 of 512 pass and 30 fail. Corner checks do not prove that an entire parameter box satisfies the result. In particular, the larger exercise must not be read as evidence of a broad parameter region.

Table C.2 varies the standardized early response while leaving the other primitives fixed. Each row reports the gain from full conduct disclosure over no disclosure. The table compares those two fixed policies; it does not independently certify a global optimum for every row. At larger values of y , the low type’s current loss can be too large to induce imitation under separable costs even at full verification, so disclosure changes only information on that branch.

Early response y	Shared costs	Separable costs
0.620	0.015385517	0.002403995
0.625	0.009973596	-0.000660155
0.630	0.005995998	-0.001868596
0.640	0.000714675	-0.001868596
0.645	-0.001179954	-0.001868596

Table C.2: Sensitivity of the full-conduct-disclosure gain over no disclosure. Other primitives are fixed. Each row is a comparison of the two endpoint policies, rather than a separate global-optimality certificate.

The effect of y is economically relevant. Both technologies have exactly the same early private and social gains at any common value of y , but different continuation incentives. The policy reversal occurs when those differences create enough additional productive imitation under shared costs, and not under separable costs. If the private loss is very small, disclosure can generate a positive full-publication gain in both; if it is too large, neither can recover the informational loss. The displayed reversal is a strict local comparison between these possibilities.

D Entry scope with finitely many markets

This appendix extends the scale results to finitely many markets and states the additional conditions needed for monotonicity and the reputation equilibrium.

Let the available markets be a finite set $\mathcal{N}$. For any portfolio $S \subseteq \mathcal{N}$, let $B_\theta(S)$ denote the entrant's operating profit under type θ , after the incumbent's optimal response. Its best operating profit at scale k is

$$v_k(p) = \max_{|S|=k} \{pB_H(S) + (1-p)B_L(S)\}, \quad v_0(p) = 0. \quad (\text{D.1})$$

Each v_k is continuous and piecewise affine, but need not be globally affine because the best portfolio can change. Choose the largest maximizing scale in $v_k(p) - kz$. Define

$$l_k(p) = \max_{j>k} \frac{v_j(p) - v_k(p)}{j - k}, \quad u_k(p) = \min_{j<k} \frac{v_k(p) - v_j(p)}{k - j}, \quad (\text{D.2})$$

with the empty maximum equal to $-\infty$ and the empty minimum equal to $+\infty$. The selected-scale cell is

$$I_k(p) = (l_k(p), u_k(p)] \cap [0, Z]. \quad (\text{D.3})$$

The half-open cell implements largest-scale selection: a tie with a larger scale excludes k , whereas a tie with a smaller scale includes it. A cell with $l_k \geq u_k$ is empty. Scales unsupported by the upper boundary of (k, v_k) are never uniquely optimal; (D.3) also restricts costs to $[0, Z]$.

Proposition D.1. *For an atomless distribution of setup costs, every scale probability induced by (D.1) is continuous in p .*

Proof. Fix p . Distinct scales have different payoff slopes in z , so they can tie only at finitely many costs. Outside those costs, the selected scale is the unique maximizing scale and remains so in a neighborhood of p , by continuity of all v_k . The indicators of each selected scale therefore converge almost surely as posteriors approach p . Dominated convergence gives continuity of their probabilities. Ties between different portfolios of the same scale do not affect this argument. $\square$

A sufficient condition for entry scope to increase as p falls is that for every $k > j$ and $p' < p$,

$$v_k(p') - v_j(p') \geq v_k(p) - v_j(p). \quad (\text{D.4})$$

To verify it, suppose scale k is selected at p but a smaller scale j is selected at p' . The optimality inequalities and (D.4) imply equality throughout; largest-scale selection then contradicts the choice of j at p' . When each scale uses a fixed portfolio and $v_k(p) = B_{Lk} - pT_k$, nondecreasing T_k is sufficient. No such ordering follows from portfolio optimization alone.

For completeness, consider a finite reduction from p_1 to $p_2 < p_1$ and two scales $0 \leq h < h' \leq |\mathcal{N}|$. Their pairwise intersection is

$$\zeta_{h'h}(p) = \frac{v_{h'}(p) - v_h(p)}{h' - h}. \quad (\text{D.5})$$

Up to finitely many tie costs, the costs switching exactly from h to h' are

$$\left(\zeta_{h'h}(p_1), \zeta_{h'h}(p_2)\right) \cap I_h(p_1)^\circ \cap I_{h'}(p_2)^\circ. \quad (\text{D.6})$$

Both optimality cells are needed to exclude a third scale. Under (D.4), positive-probability switches strictly increase mean scope, since no costs switch downward. Yet a scale disappearing at a single posterior cannot create an aggregate jump, by Proposition D.1.

Continuity of scale probabilities does not imply continuity of incumbent losses. If equally large portfolios exchange optimality, all costs choosing that scale can switch composition together; different losses across those portfolios can then cause a jump. The two-market model avoids this issue because market 1 is the strictly best singleton. More generally, the reputation equilibrium needs additional conditions on composition switches.

E Aggregation across published and unpublished records

Let $K(p)$ be expected entry scope under an atomless cost distribution satisfying the continuation assumptions. Along the equilibrium with H responding surely, write $w = \mu + (1 - \mu)\rho$, so $p_R = \mu/w$. The difference between mean scope with a verified record and mean scope without a record is

$$J(w) = (1 - w)K(0) + wK(\mu/w) - K(\mu). \quad (\text{E.1})$$

It can also be written as

$$J(w) = (1 - w)[K(0) - K(\mu)] - w[K(\mu) - K(p_R)]. \quad (\text{E.2})$$

If K is decreasing, both bracketed differences are nonnegative; their difference need not be. At a differentiability point,

$$J'(w) = K(p_R) - K(0) - p_R K'(p_R). \quad (\text{E.3})$$

The level difference is nonpositive and the derivative term nonnegative: imitation changes both the frequency and content of response records.

Unconditional mean scope and its derivative are

$$\bar{K}(\tau) = K(\mu) + \tau J(w(\tau)), \quad (\text{E.4})$$

$$\bar{K}'(\tau) = J(w(\tau)) + \tau J'(w(\tau))w'(\tau). \quad (\text{E.5})$$

The first term in (E.5) holds the experiment fixed; the second captures its change through imitation. Neither has a general sign.

Both signs occur with smooth, positive densities. Keep the shared-cost primitives in (3.15) and replace the uniform distribution by

$$F_\varepsilon(z) = (1 - \varepsilon)\frac{z}{Z} + \varepsilon\left(\frac{z}{Z}\right)^2, \quad \varepsilon \in \{-1/10, 1/10\}. \quad (\text{E.6})$$

The density is everywhere at least 9 on $[0, 1/10]$, and the equilibrium and uniqueness inequalities continue to hold. The respective separation thresholds are approximately 0.52067 and 0.52550. On either separating interval, $w = \mu$ and direct integration gives

$$\overline{K}'_\varepsilon(\tau) = \mu K_\varepsilon(1) + (1 - \mu)K_\varepsilon(0) - K_\varepsilon(\mu) = \varepsilon \frac{119877142504381}{8067162139660800}. \quad (\text{E.7})$$

The derivative is approximately -0.001486 for $\varepsilon = -1/10$ and 0.001486 for $\varepsilon = 1/10$. These examples do not extend the uniform-cost policy theorem.

At an entry-region boundary, p_R falls as τ rises. A right derivative in τ therefore uses a left derivative in p . The mixing condition gives

$$p'_{R,+}(\tau) = -\frac{D(p_R)}{\tau D'_-(p_R)}, \quad p'_{R,-}(\tau) = -\frac{D(p_R)}{\tau D'_+(p_R)}, \quad (\text{E.8})$$

provided the one-sided derivatives of D exist and are positive. Apply the same chain rule to (E.5). At the separation boundary, the derivative from below has $w' = 0$. Levels remain continuous.

Under uniform costs and interior cutoffs, affine $K(p) = V_2(p)/Z$ implies $J(w) = 0$. Equations (E.4)–(E.5) then give the exact zero effect in the main text, including at kinks. Both the cost distribution and the entry geometry are needed.

References

- Arya, A., Frimor, H., Mittendorf, B., and Pfeiffer, T. (2025). Disclosure to competitors in light of endogenous firm investments. *Contemporary Accounting Research*, 42(3), 1960–1986. <https://doi.org/10.1111/1911-3846.13050>.
- Atanasov, T. (2024). Strategic disclosure incentives in a multisegment firm. *The Accounting Review*, 99(6), 27–50. <https://doi.org/10.2308/TAR-2023-0155>.
- Bar-Isaac, H. (2012). Transparency, career concerns, and incentives for acquiring expertise. *The B.E. Journal of Theoretical Economics*, 12(1), Article 4. <https://doi.org/10.1515/1935-1704.1796>.

- Bar-Isaac, H., and Deb, J. (2021). Reputation with opportunities for coasting. *Journal of the European Economic Association*, 19(1), 200–236. <https://doi.org/10.1093/jeea/jvaa008>.
- Dixit, A. (1980). The role of investment in entry-deterrence. *Economic Journal*, 90(357), 95–106. <https://doi.org/10.2307/2231658>.
- Feng, H., Li, Y., and Shuai, J. (2023). Uniform pricing as a barrier to entry. *Journal of Industrial Economics*, 71(1), 176–191. <https://doi.org/10.1111/joie.12320>.
- Fudenberg, D., and Kreps, D. M. (1987). Reputation in the simultaneous play of multiple opponents. *Review of Economic Studies*, 54(4), 541–568. <https://doi.org/10.2307/2297482>.
- Goolsbee, A., and Syverson, C. (2008). How do incumbents respond to the threat of entry? Evidence from the major airlines. *Quarterly Journal of Economics*, 123(4), 1611–1633. <https://doi.org/10.1162/qjec.2008.123.4.1611>.
- Hoang, D., and Ruckes, M. (2017). Corporate risk management, product market competition, and disclosure. *Journal of Financial Intermediation*, 30, 107–121. <https://doi.org/10.1016/j.jfi.2016.07.003>.
- Hwang, Y., and Kirby, A. J. (2000). Competitive effects of disclosure in a strategic entry model. *Review of Accounting Studies*, 5, 57–85. <https://doi.org/10.1023/A:1009664918749>.
- Ishinagi, Y., and Oh, J. (2024). Observable and unobservable investment information in multiple markets. *Economic Analysis Letters*, 3(3), Article 60. <https://doi.org/10.58567/ea103030001>.
- Jiang, X., and Xiong, Y. (2025). Disclosing endogenous cost information. *The Accounting Review*, 100(2), 249–268. <https://doi.org/10.2308/TAR-2023-0296>.
- Jost, P.-J. (2023). Price commitment and the strategic launch of a fighter brand. *Quantitative Marketing and Economics*, 21(3), 381–435. <https://doi.org/10.1007/s11129-023-09266-9>.
- Kim, S. H. (2009). Predatory reputation in US airline markets. *International Journal of Industrial Organization*, 27(5), 592–604. <https://doi.org/10.1016/j.ijindorg.2009.02.002>.
- Kreps, D. M., and Wilson, R. (1982). Reputation and imperfect information. *Journal of Economic Theory*, 27(2), 253–279. [https://doi.org/10.1016/0022-0531\(82\)90030-8](https://doi.org/10.1016/0022-0531(82)90030-8).

- Li, H. (2025). Transparency and policymaking with endogenous information provision. *RAND Journal of Economics*, 56(1), 91–105. <https://doi.org/10.1111/1756-2171.12495>.
- Milgrom, P., and Roberts, J. (1982). Predation, reputation, and entry deterrence. *Journal of Economic Theory*, 27(2), 280–312. [https://doi.org/10.1016/0022-0531\(82\)90031-X](https://doi.org/10.1016/0022-0531(82)90031-X).
- Prat, A. (2005). The wrong kind of transparency. *American Economic Review*, 95(3), 862–877. <https://doi.org/10.1257/0002828054201297>.
- Rey, P., Spiegel, Y., and Stahl, K. (2025). A dynamic model of predation. *RAND Journal of Economics*, 56, 738–755. <https://doi.org/10.1111/1756-2171.70027>.
- Spence, A. M. (1977). Entry, capacity, investment and oligopolistic pricing. *Bell Journal of Economics*, 8(2), 534–544. <https://doi.org/10.2307/3003302>.
- Srinivasan, K. (1991). Multiple market entry, cost signalling and entry deterrence. *Management Science*, 37(12), 1539–1555. <https://doi.org/10.1287/mnsc.37.12.1539>.
- Sweeting, A., Roberts, J. W., and Gedge, C. (2020). A model of dynamic limit pricing with an application to the airline industry. *Journal of Political Economy*, 128(3), 1148–1193. <https://doi.org/10.1086/704760>.