# Inference with an Incomplete Model of Firm Exit

Guillermo Marshall*

August 2, 2026

**Abstract**

In a declining industry, the firm that outlasts its rivals inherits the market, so each firm delays its own exit in the hope that a rival concedes first, and private information about who can hold on the longest sustains the standoff. We study what exit data reveal about the resulting strategic delay using an incomplete model that accommodates multiple equilibria. The analysis is based on two inequalities that hold in every equilibrium—a firm does not exit while its profit covers its cost, and it does not stay when even monopoly profit would not. Together, these inequalities bound each firm's cost, and, with it, the strategic delay and its decomposition into a competition component (oligopolistic free-riding) and an information component (learning under asymmetric information). Assuming firms play symmetric strategies bounds the delay and its competition component away from zero; the decomposition between the components is equilibrium-dependent. We apply the bounds to movie exhibition, bowling, book stores, and florists in US county data between 1986 and 2016. Exiting firms survive 7.6 to 11.9 years past the industry's peak; the bounds attribute at most a seventh to about two fifths of that survival to strategic delay, with the largest share in bowling.

**Keywords:** Exit; war of attrition; contest; industry dynamics; partial identification

# 1 Introduction

Firms in declining industries exit one at a time. How long each waits before leaving determines how quickly capacity is withdrawn and which firms remain once the industry settles at its smaller size, and strategic behavior distorts both margins. Capacity leaves slowly, because each firm delays its own exit in the hope that a rival concedes first. And the firms left standing need not be the most efficient: the contest admits multiple equilibria, and which firm outlasts the others depends on the equilibrium played, not only on costs. The exit process, therefore, governs both the speed of adjustment in a declining industry and the efficiency of the configuration that survives it.

Consider a firm deciding when to exit as demand falls. While active with $n-1$ rivals, it earns a flow $\pi_n(t)$ that declines in calendar time, against an opportunity cost $\theta$, the value

*University of British Columbia, Sauder School of Business; email: guillermo.marshall@sauder.ubc.ca

of its assets outside the industry. Once decline pushes $\pi_n(t)$ below $\theta$, every instant the firm stays costs it $\theta - \pi_n(t)$. Waiting is nonetheless valuable for two reasons. First, if a rival exits, each survivor's flow steps up from $\pi_n(t)$ to $\pi_{n-1}(t)$: by absorbing losses for longer, the firm buys the chance to inherit the thinner market. Second, when costs are private, a rival that has not yet exited is informative; the lower its cost, the longer it holds on, so the firm reads a rival's continued presence as evidence that the contest will last. The war-of-attrition model of Fudenberg and Tirole (1986) formalizes this decision.

Taking the model to data requires solving it, and solving it requires assumptions on the primitives. The classic war of attrition has a continuum of equilibria (Riley, 1980; Hendricks et al., 1988). Fudenberg and Tirole (1986) consider a two-firm model and obtain a unique equilibrium by restricting the cost distribution and the profit flows: with positive probability, a firm's cost is low enough that staying in forever is dominant, and the market may keep both firms forever. Takahashi (2015a) preserves the restriction in the $N$-firm game and builds an estimator on the uniqueness it delivers. The structural approach also assumes that firms play a common symmetric strategy, so that exits occur in cost order and each exit time inverts into a cost through the equilibrium indifference condition.

The payoff from the inversion is a decomposition of the strategic delay in exit, which is defined as each firm's extra years past the exit date a profit-maximizing planner would choose. The *competition* component is the delay that results from each firm free-riding on the chance that a rival goes first. A firm times its exit against its own profit flow rather than the marginal social flow, so it holds on past the point at which the industry would gain from its exit. This component exists even when the firms' costs are common knowledge. The *information* component is the extra delay that private costs add on top: the standoff during which a firm absorbs a loss while it waits to learn whether a rival concedes first. Takahashi (2015a) runs this program on the US movie-theater shakeout and assigns 96% of the average 2.68-year delay to competition and 4% to information. Both numbers rest on the uniqueness restrictions, on the equilibrium they deliver, and on symmetry. Symmetry is restrictive when firms differ in ways the econometrician does not observe, and it rules out inefficient survival: under a common strategy, the survivors are the lowest-cost firms by construction.

This paper asks two questions. First, what do exit data reveal about the strategic delay under minimal assumptions? Second, which assumptions carry identifying power? To answer them, we impose neither the uniqueness restrictions nor an equilibrium selection, and we solve for no equilibrium. Instead, we work with two dominance inequalities that hold in every equilibrium of the game. First, a firm does not quit while its current flow still covers its cost of continuing, so an exiting firm's cost is at least its flow at the moment it leaves. Second, a firm does not stay active when even the most favorable continuation, becoming a monopolist, fails to cover its cost, so its cost is at most the monopoly flow. Together, the two inequalities bound each firm's cost, market by market, with no assumption on how rivals play. The bounds are therefore robust to equilibrium multiplicity: they hold when several equilibria coexist, and when different markets play different ones.

The dominance inequalities on their own leave the total and information delays unsigned and put only a nonnegative floor under the competition delay. Adding the assumption of symmetric strategies, used only through its implication that exits occur in cost order, signs the decomposition: the total delay and its competition component are positive and bounded away from zero, and a within-market chain of bounds makes the information component become strictly positive where the chain binds. What symmetry leaves open is the split between the two components. That is, the size of the delay rests on symmetry, not on the equilibrium; the split between the components rests on the equilibrium.

Our main empirical exercise bounds the strategic delay in four declining local-oligopoly industries in the County Business Patterns county panel, 1986–2016: motion picture exhibition, bowling centers, book stores, and florists. The bounds take the profit function as an input, so we measure every component of it from decisions taken outside the war itself. On average, an exiting firm in these industries survives 7.6 to 11.9 years past the industry's peak. The bounds attribute 0.1 to 1.5 of those years to the war at least and 1.2 to 3.9 at most, between a seventh and about two fifths of observed survival; the information component is capped at 1.1 to 2.4 years, and the competition component is bounded away from zero at the point estimates in all four industries. The war's share of survival is largest in bowling, the industry with the strongest oligopolistic free-riding incentive.

The second application returns to the shakeout that Takahashi (2015a) studies: 2,606 county markets losing indoor movie theaters to television between 1949 and 1955. His structural estimates provide a point-identified benchmark. We keep his market definition and his estimated profit function and drop everything else his estimator imposes: the equilibrium selection and the parametric cost distribution. The total delay is bounded between 2.4 and 4.2 years, and his 2.7-year equilibrium-selected estimate lies inside: the bounds recover the size of the delay without fully specifying the game, selecting an equilibrium, or solving the game. The split between the delay components remains open, competition at 0.1 to 3.3 years and information at 0.2 to 3.9, so his exit record is consistent with a war sustained mostly by business stealing and with one sustained mostly by private information.

Our contribution is twofold. First, we characterize what exit data identify about the strategic-delay decomposition under each layer of assumptions: dominance alone, dominance plus symmetry, and the full equilibrium. Second, we quantify the answer in two applications, providing new evidence on exit inefficiency in declining industries.

## 1.1 Related literature

Our paper contributes to several strands of the literature. First, we contribute to the empirical war-of-attrition literature. Riley (1980) introduces the war of attrition with privately known types; Fudenberg and Tirole (1986) develop the incomplete-information duopoly exit model that we take as our data-generating process, and Ghemawat and Nalebuff (1985) characterize the complete-information benchmark in which the contest collapses and exit is immediate. Hendricks et al. (1988) document the equilibrium multiplicity of the war of attrition; the support restriction on the cost distribution that removes it in Fudenberg and Tirole (1986) is one of the assumptions we drop. Bulow and Klemperer (1999) generalize the war of attrition to $N$ prizes, and show that the lowest-cost firms never exit. Takahashi (2015a) estimates an $N$-player extension of Fudenberg and Tirole (1986) on US movie-theater data and decomposes the resulting strategic delay into competition and information components. The econometrics of interdependent timing decisions outside the attrition framework is developed by Honoré and de Paula (2010), whose treatment of simultaneous durations shares our concern with what joint exit timing can identify.

Second, we contribute to the partial-identification literature. Our work relates to Haile and Tamer (2003), who bound the valuation distribution in English auctions using two dominance restrictions. Ciliberto and Tamer (2009) estimate an entry model using a method that does not require equilibrium selection. Armstrong (2013) extends the Haile and Tamer (2003) auction bounds to allow for unobserved auction-level heterogeneity. Aradillas-López et al. (2013) relax the button-auction equilibrium assumption in a related way. Gentry and Li (2014) partially identify auction models with selective entry.

Within dynamic games, we build most directly on Abito and Chen (2023), who develop a partial-identification framework for dynamic games via continuation-payoff bounds derived from the repeated-game equilibrium payoff set. Our object is different: they bound continuation-value primitives via a Folk-Theorem argument in a collusion setting; we bound firms' costs and exit delays via dominance restrictions. Décamps et al. (2024) derive robust testable implications for the war of attrition in continuous time; their focus is on signing and testing predictions of the theory under an arbitrary equilibrium, while ours is the set identification of the strategic-delay decomposition.

## 1.2 Outline

The rest of the paper is organized as follows. Section 2 sets up the environment. Section 3 develops the dominance restrictions and bounds each firm's realized cost. Section 4 states the symmetry assumption and bounds the strategic-delay decomposition market by market, with and without it. Section 5 presents the main application using the County Business Patterns county panel. Section 6 applies the bounds to the US movie-theater shakeout of 1949–55. Lastly, Section 7 concludes.

# 2 Environment

Our model of exit is based on the declining-industry version of Fudenberg and Tirole (1986).[1]

There are $N \geq 2$ firms, all active at $t = 0$, discounting payoffs at rate $r$. Firm $i$ privately draws an opportunity cost $\theta_i \geq 0$ from a distribution $G_i$ with density $g_i$; draws are independent across firms, and they need not be identically distributed (unless otherwise noted).

The only action is the choice of an exit time. When $k$ firms are active, each active firm earns instantaneous flow $\pi_k(t)$, gross of $\theta_i$, with

$$\pi_N(t) \leq \pi_{N-1}(t) \leq \cdots \leq \pi_1(t),$$

so $\pi_k$ is decreasing in the number of active rivals; $\pi_1(t)$, the flow of a firm with no active rivals, is the monopoly flow. A firm that has exited cannot re-enter.

The profit flow has two arguments. The subscript $k$ is the *endogenous* market structure: it steps down by one at each exit. The argument $t$ is *exogenous*. The industry is in secular decline, so each $\pi_k(\cdot)$ is decreasing in calendar time and drifts toward a floor as $t \to \infty$. The

[1] In the growing-industry case the cost bound developed below has the constant monopoly break-even cost $r\int_0^\infty e^{-rt}\pi_1(t)\,dt$ as its upper bound rather than the current monopoly flow $\pi_1(t)$; the analysis goes through with this substitution.

decline is deterministic and commonly observed. Time is a public state, and the only private information is the vector of costs $\{\theta_i\}$, so each $\pi_k(t)$ is a known function of time.

A pure strategy maps a firm's cost to its exit time, $\theta \mapsto T_i(\theta)$, applied while rivals remain active. We do not require firms to share a common strategy (unless otherwise noted). The econometrician observes a panel of markets and, in each, the exit times $t_1, \ldots, t_R$ of the $R$ exiters with $R \leq N$ together with covariates $X$.

We place no further structure on the $G_i$ or on the flows: no parametric family for costs, and no condition linking the support of costs to the floors the flows decline toward. Fudenberg and Tirole (1986) assume that with positive probability a firm's cost is low enough that staying in forever is dominant: the support of the cost distribution reaches below the level that $\pi_N$ declines toward, a joint restriction on the cost distribution and the profit flows. The restriction implies that the war may never resolve: with positive probability every firm draws such a cost, no firm ever exits, and the market accommodates all $N$ incumbents forever. That never-exit probability supplies the boundary condition at infinity that selects a single solution of the equilibrium differential equation; without it, the game has a continuum of equilibria (Riley, 1980). Takahashi (2015a) applies the restriction to the $N$-firm game, where it delivers a unique symmetric equilibrium and makes a full-solution estimator feasible. We impose no such restriction, so the equilibrium set may be large, symmetric, and asymmetric alike, and different markets in a panel may sit on different equilibria. Everything below is robust to this multiplicity, because the restrictions we use hold in every equilibrium separately.

The standard empirical approach to this game imposes a symmetric equilibrium and recovers each firm's cost by inverting the equilibrium indifference condition, which defines a one-to-one map between costs and exit times, single-valued only under the uniqueness restrictions above; Appendix A states it. As mentioned, our analysis does not rely on the equilibrium indifference condition or on choosing a particular equilibrium.

# 3 Dominance restrictions

Everything in the paper rests on three facts that hold in every equilibrium of the exit game, whatever selects the equilibrium and whatever firms believe about one another. The three facts are statements about a single firm's own behavior, so they hold for any strategy profile.

**Lemma 1** (Own-cost monotone exit)**.** *Fix any behavior of firm $i$'s rivals. Then firm $i$'s best-response exit time $T_i(\theta_i)$ is nonincreasing in its own cost $\theta_i$, in every equilibrium and against any beliefs.*

*Proof.* Fix the rivals' behavior. Rivals' types are independent of $\theta_i$, and their strategies condition only on their own types and the public history, which does not reveal $\theta_i$ while $i$ remains active; the law of the active-count process $\{n(t)\}$ that firm $i$ faces therefore does not depend on $\theta_i$. Writing a plan as a stopping time $\tau$ adapted to the public history, the firm's payoff is

$$U(\tau, \theta_i) \;=\; \mathbb{E}\Big[\int_0^\tau e^{-rt}\pi_{n(t)}(t)\,dt\Big] \;+\; \theta_i\,\mathbb{E}\big[e^{-r\tau}\big]/r,$$

the discounted flow while active plus the present value of the exit payoff. The coefficient on $\theta_i$ is decreasing in $\tau$, so for $\tau \leq \tau'$, the difference $U(\tau', \theta_i) - U(\tau, \theta_i)$ is decreasing in $\theta_i$:

$U$ has decreasing differences in $(\tau, \theta_i)$. By Topkis's monotonicity theorem, if $\theta_i' > \theta_i$, then $T_i(\theta_i) \geq T_i(\theta_i')$. That is, the best-response exit time is nonincreasing in $\theta$. The argument nowhere uses the rivals' strategies or firm $i$'s beliefs beyond the independence of $\{n(t)\}$ from $\theta_i$, hence it holds whatever rivals do, and whether or not strategies are common. $\square$

**Lemma 2** (Individual-rationality bounds)**.** *A firm observed to exit at $t$ with $k$ firms active just before has a cost satisfying*

$$\pi_k(t) \ < \ \theta \ < \ \pi_1(t).$$

The lower bound is immediate: a firm earning a positive flow net of its cost in the current state ($\theta < \pi_k(t)$) strictly prefers to stay, so an exiting firm must have $\theta > \pi_k(t)$. A firm cannot be sure its rivals will exit, so its willingness to remain requires that even the most optimistic continuation (all rivals leave and it becomes a monopolist) be worth it, so $\theta < \pi_1(t)$. Lemma 2 bounds the latent type in cost units using only the profit function.

**Lemma 3** (Active-firm bound)**.** *A firm active at time $t$ has cost $\theta < \pi_1(t)$, in every equilibrium and for any strategy profile, when payoffs are decreasing in $t$.*

The argument is the upper arm of Lemma 2 applied to a firm that has not yet left: if $\theta \geq \pi_1(t)$, then every continuation flow from $t$ onward is at most $\pi_1(s) \leq \pi_1(t) \leq \theta$, so remaining active is dominated and the firm would already have exited. Therefore, a firm still active when the sample ends at $h$ reveals $\theta < \pi_1(h)$, an upper arm for every censored cost.

**Proposition 1** (Realization bounds)**.** *The cost of firm $i$, observed to exit at $t_i$ with $k_i$ firms active just before, satisfies*

$$\pi^i_{k_i}(t_i) \ < \ \theta_i \ < \ \pi^i_1(t_i),$$

*and a firm still active at the observation horizon $h$ satisfies $\theta_i < \pi^i_1(h)$ under non-increasing payoffs.*

This bounds each incumbent's realized cost in the market where it operates, often the object of direct interest. The width $\pi^i_1(t_i) - \pi^i_{k_i}(t_i)$ is narrowest for the last exit (width $\pi^i_1 - \pi^i_2$) and widest for the first ($\pi^i_N$ far below $\pi^i_1$). It holds whatever each firm's strategy may be.

**Monopolist exits are point-identified.** At $k_i = 1$ the bounds $(\pi_1(t), \pi_1(t))$ degenerate. This is not lost information but the opposite: a monopolist faces a single-agent stopping problem with no rival to wait for, exits at its myopic boundary, and its exit therefore reveals $\theta_i = \pi^i_1(t_i)$ exactly.

# 4 Bounding the strategic-delay decomposition

The central use of a structural war-of-attrition model is to measure how much exit is delayed, and to say why. The model compares the equilibrium exit path to an efficient benchmark and splits the gap into a piece due to oligopolistic free-riding and a piece due to asymmetric information learning. Takahashi (2015a) runs exactly this exercise for the US movie-theater shakeout. We show that the same decomposition can be bounded from the dominance restrictions already in hand, plus the symmetry assumption, and we report alongside it a companion set of bounds that drops the symmetry assumption; the distance between the two sets of bounds measures what symmetry contributes.

**Assumption 1** (Symmetric strategies)**.** All firms play a common strategy, $T_i(\cdot) = T(\cdot)$ for all $i$. Hence, a single decreasing boundary $\Phi(t) \equiv T^{-1}(t)$ governs every firm.

Symmetry is the assumption that the structural approach maintains. It restricts strategies, not primitives, but the two are linked: a symmetric equilibrium is the natural object when costs are drawn from a common distribution, $G_i = G$ for all $i$, and when the $G_i$ differ, an equilibrium in a common strategy need not exist. That is, Assumption 1 carries ex-ante symmetry with it, and it is the only place where a common cost distribution enters. Our results use one implication of it.

**Lemma 4** (Ordered exits)**.** *Under Assumption 1, exits occur in descending order of cost in every market: the $r$-th firm to exit has the $r$-th highest cost, and any firm active at time $t$ has a weakly lower cost than every firm that has exited by $t$.*

*Proof.* All firms play the same $T(\cdot)$, and Lemma 1 makes $T$ nonincreasing in cost, so the chronological exit order is the descending cost order. □

Every result that invokes Assumption 1 uses it only through Lemma 4. The ordering is logically weaker than a common strategy, so each bound below remains valid under any strategy profile, symmetric or not, whose exits occur in cost order; the proofs cite the lemma to keep this visible. Under the ordering, the last firm standing is the lowest-cost firm in its market, and Lemma 3 bounds that cost by $\pi_1(h)$ at the observation horizon.

Symmetry does not restore uniqueness. Without the support condition of Section 2, several symmetric equilibria may coexist, each with its own boundary $\Phi$. Every one of them orders exits by cost, so the bounds below hold in all of them, and in different ones across markets. What the assumption buys is the ordering, not a selection.

## 4.1 Three exit paths

Fix a market with $N$ firms active at $t = 0$, of which $R \leq N$ are observed to exit, at dates $0 < T_{(1)} \leq \cdots \leq T_{(R)}$, while $N - R$ survive to the horizon $h$; write $\pi_n \equiv \pi_n(t)$ for the flow when $n$ firms are active.

We index an exit by its *exit rank* $r$, its place in the observed exit sequence, so the $r$-th exit occurs at $T_{(r)}$. The *active count* at that exit is $n_r = N - r + 1$, the number of firms active just before it, the exiter included; it is the subscript of the flow the exiter earns at its exit, $\pi_{n_r}$, and it is the $k$ of Lemma 2. We index the realized costs by *cost rank*, $\theta_{(1)} < \cdots < \theta_{(N)}$, so $\theta_{(\nu)}$ is the $\nu$-th lowest cost in the market. Exit rank and cost rank are distinct indices in general, and Lemma 4 is what ties them: under cost-ordered exits the $r$-th exiter holds cost rank $n_r$, so its cost is $\theta_{(n_r)}$ and the active count doubles as its cost rank. Section 4.4 drops the tie and lets an exiter's cost rank be set-valued.

The *war of attrition* is the equilibrium itself, with the $r$-th exit time the data point $t^*_r = T_{(r)}$. We compare these observed (equilibrium) exit times against two benchmarks with the goal of bounding the strategic exit delay as well as decomposing it into two components: a competition and an information component. We follow the decomposition in Takahashi (2015a).

The first benchmark is the *coordination benchmark*, which has each firm leave the instant its own flow falls to its cost, with no learning and no ex post regret. Along the coordination path exits occur in descending cost order whatever the equilibrium did, so its $r$-th exit is the firm holding cost rank $n_r$ and occurs at

$$t^C_r = \pi_n^{-1}\big(\theta_{(n)}\big), \qquad n = N{-}r{+}1. \tag{1}$$

Under the ordering of Lemma 4 the observed $r$-th exiter is the same firm that would be the $r$-th exiter in the coordination benchmark.

The second benchmark is the *regulator benchmark* in which a regulator chooses the whole schedule to maximize discounted industry profit,

$$\max_{\{t_1,\dots,t_N\}} \int \sum_{k \text{ active}} \big[\pi_{n_t}(t) - \theta_k\big] e^{-rt}\, dt,$$

where $n_t$ is the number of firms active at $t$ and the inner sum runs over those firms; the schedule internalizes the business-stealing externality each survivor imposes. The benchmark maximizes industry profit rather than welfare: the model has no consumer side, so the schedule weighs the externality among firms only and is silent on any consumer value of the capacity that exits.

**Lemma 5** (Planner decoupling)**.** *Suppose the flow depends only on the number of active firms, $\pi_N \le \cdots \le \pi_1$, and that each $\pi_n$ and the marginal social flow*

$$\Psi_n(t) \equiv n\,\pi_n(t) - (n{-}1)\,\pi_{n-1}(t)$$

*is continuous, positive, and strictly decreasing in $t$, with $\Psi_n(t)$ nonincreasing in $n$ at each $t$. Then the industry-profit-maximizing schedule drops firms in decreasing order of cost, and the firm dropped while $n$ are active leaves at*

$$t^R = \Psi_n^{-1}\big(\theta_{(n)}\big),$$

*the inverse of the marginal social flow of the $n$-th firm. In particular, $t^R$ depends on the dropped firm's own cost only, and $\Psi_n \le \pi_n$ implies $t^R \le t^C$.*

*Proof.* Write $C_n \equiv \sum_{k\le n} \theta_{(k)}$ for the cost-sum when $n$ firms are active, let $s_n$ be the date ending the $n$-firm phase, and set $s_{N+1} \equiv 0$, $\pi_0 \equiv 0$, and $C_0 \equiv 0$. Since $n\pi_n(t)$ does not depend on which firms are active, the planner drops the highest-cost active firm at every boundary, and $C_n = \sum_{k\le n} \theta_{(k)}$. The planner seeks to maximize

$$W(s_1,\dots,s_N) \;=\; \sum_{n=1}^{N} \int_{s_{n+1}}^{s_n} e^{-rt}\big[n\,\pi_n(t) - C_n\big]\, dt.$$

Each $s_n$ enters exactly two terms, as the upper limit of the $n$-firm integral and the lower limit of the $(n-1)$-firm integral, so $W$ is additively separable in the boundaries and

$$\frac{\partial W}{\partial s_n} = e^{-rs_n}\big[n\pi_n(s_n) - C_n\big] - e^{-rs_n}\big[(n-1)\pi_{n-1}(s_n) - C_{n-1}\big] = e^{-rs_n}\big[\Psi_n(s_n) - \theta_{(n)}\big],$$

where the second equality collects the flow terms into $\Psi_n$ and uses $C_n - C_{n-1} = \theta_{(n)}$. Since $\Psi_n$ is strictly decreasing in $t$, $\partial W/\partial s_n$ has the sign of $\Psi_n(s_n) - \theta_{(n)}$, positive below the root and negative above it, so $W$ is strictly quasiconcave in $s_n$ and the root is the global maximizer in that coordinate; the roots satisfy $s_N \leq \cdots \leq s_1$ because $\Psi_{n-1} \geq \Psi_n$ and $\theta_{(n-1)} < \theta_{(n)}$, so the separable optimum is feasible. Finally $\pi_n - \Psi_n = (n-1)(\pi_{n-1} - \pi_n) \geq 0$ is the business-stealing externality, so $\Psi_n \leq \pi_n$ and $t^R \leq t^C$, with strict inequality for $n \geq 2$ wherever the flow strictly declines in the active count. □

Following Takahashi (2015a), the total strategic-interaction delay of the $r$-th exit, $t^*_r - t^R_r$, may be decomposed as

$$\underbrace{t^*_r - t^R_r}_{\text{total}} = \underbrace{\left(t^C_r - t^R_r\right)}_{\text{oligopolistic competition}} + \underbrace{\left(t^*_r - t^C_r\right)}_{\text{strategic (informational) behavior}}.$$

The two components isolate different frictions, and we describe the economics of each in turn.

The competition delay compares two exit rules for the same firm. The coordination rule has the firm leave when its *private* flow $\pi_n$ falls to its cost; the regulator has it leave when its *marginal social* flow $\Psi_n$ does. The gap between the two rules, $\pi_n - \Psi_n = (n-1)(\pi_{n-1} - \pi_n)$, is the business-stealing externality: while the firm stays, each of its $n-1$ rivals earns $\pi_n$ rather than $\pi_{n-1}$, and the firm does not internalize their loss. That is, the firm times its exit against its own flow rather than the marginal social flow, and holds on past the point at which the industry gains from its departure. The friction is present even when every cost is common knowledge, and its size involves no equilibrium object: the wedge is a functional of the profit function.

The informational delay compares equilibrium play with the firm's own myopic optimum. The coordination time $t^C$ is the *myopic* exit point: the instant the firm's current flow stops covering its cost, $\pi_n(t^C) = \theta$. In the war of attrition the firm leaves strictly later, at $t^*$: past $t^C$ it absorbs a loss of $\theta - \pi_n(t)$ per instant, and it does so only because a rival may concede first, after which the market thins and its flow steps up toward the monopoly value $\pi_1(t) > \theta$. Private costs sustain the standoff: were the cost vector common knowledge, the contest would collapse, the higher-cost firm would concede at once, and the equilibrium path would coincide with the coordination path (Ghemawat and Nalebuff, 1985). How long the standoff lasts depends on where the equilibrium places the exiting type between its myopic point and the monopoly ceiling, so the length of the informational delay is a property of the equilibrium, not of the profit function.

## 4.2 Bounds on the three delays

At each observed exit, Lemma 2 brackets the exiter's latent cost, $\pi_{n_r}(T_{(r)}) < \theta_{(n_r)} < \pi_1(T_{(r)})$. Since (1) and Lemma 5 make all three benchmark times decreasing functions of $\theta_{(n_r)}$, the bounds map to an interval on each delay.

Evaluating the three benchmark maps at the endpoints of the exiter's own bounds gives the *direct* bounds. Write $w_n(\theta) \equiv \pi_n^{-1}(\theta) - \Psi_n^{-1}(\theta)$ for the competition wedge at active count $n$, and $\mathcal{I}_r \equiv [\pi_{n_r}(T_{(r)}), \pi_1(T_{(r)})]$. For every market and observed exit rank $r$, under Assumption 1,

$$t_r^* - t_r^R \in \Big[\, T_{(r)} - \Psi_{n_r}^{-1}\big(\pi_{n_r}(T_{(r)})\big), \quad T_{(r)} - \Psi_{n_r}^{-1}\big(\pi_1(T_{(r)})\big) \,\Big], \tag{2}$$

$$t_r^C - t_r^R \in \Big[\, \min_{\theta \in \mathcal{I}_r} w_{n_r}(\theta), \quad \max_{\theta \in \mathcal{I}_r} w_{n_r}(\theta) \,\Big], \tag{3}$$

$$t_r^* - t_r^C \in \Big[\, 0, \quad T_{(r)} - \pi_{n_r}^{-1}\big(\pi_1(T_{(r)})\big) \,\Big]. \tag{4}$$

Each benchmark time is a decreasing function of $\theta_{(n_r)}$, so its image over $\mathcal{I}_r$ is the interval between its endpoint values, and for the difference $w_{n_r}$ of two decreasing functions the image is the min/max over the one-dimensional range. Throughout, the inverse maps are generalized inverses, $\pi_n^{-1}(\theta) \equiv \inf\{t \geq 0 : \pi_n(t) \leq \theta\}$ and $\Psi_n^{-1}(\theta) \equiv \inf\{t \geq 0 : \Psi_n(t) \leq \theta\}$ with $\inf \emptyset = +\infty$.

The competition delay is signed and bounded away from zero. Intervals (2) and (3) are strictly positive, because $\Psi_{n_r} < \pi_{n_r}$ for $n_r \geq 2$ puts $\Psi_{n_r}^{-1}$ strictly below $\pi_{n_r}^{-1}$: the business-stealing wedge is a known functional of the profit function, so the data bound the competition delay, and with it the total delay, away from zero on both sides with no equilibrium input.

The informational delay is signed but not bounded away from zero. Dominance identifies the sign: strict individual rationality forces $\theta_{(n_r)} > \pi_{n_r}(T_{(r)})$, so $t^C = \pi_{n_r}^{-1}(\theta_{(n_r)}) < T_{(r)} = t^*$, and the realized informational delay is strictly positive in every market. Dominance does not identify the size: the data never reveal how far the cost sits above the floor $\pi_{n_r}(T_{(r)})$, and as $\theta_{(n_r)} \downarrow \pi_{n_r}(T_{(r)})$ the implied delay falls to zero, so the identified set in (4) has infimum zero. That is, every cost in the bounds implies a strictly positive delay, but the bounds place no positive floor under its size: any positive magnitude a structural estimate reports for it is supplied by the equilibrium indifference condition, not by the data.

The direct bounds do not use the full content of cost-ordered exits: they invoke Lemma 4 only to give the $r$-th exiter cost rank $n_r$. The lemma carries more than that, because it also ranks each exiter against every firm that leaves after it. Consider a market with two observed exits. The first exiter's own bounds say its cost exceeds its flow at exit, $\pi_N(T_{(1)})$. The second exiter reveals, at its own later exit, a cost above $\pi_{N-1}(T_{(2)})$, and because exits are cost-ordered, the first exiter's cost exceeds the second's. If the second floor is the higher of the two (i.e., $\pi_{N-1}(T_{(2)}) > \pi_N(T_{(1)})$), which happens when the flow gained from facing one fewer rival exceeds what the clock has taken between the two dates, it becomes the binding floor on the first exiter as well. The lower bounds of later exiters can thus be used to sharpen the bounds associated with each exit.

Three objects summarize a market's exits. The *own floor* $f_r \equiv \pi_{n_r}(T_{(r)})$ is the $r$-th exiter's flow at its exit, the lower arm of Lemma 2. The *ceiling* $u_r \equiv \pi_1(T_{(r)})$ is the monopoly flow at the same date, the upper arm. The *chain floor*

$$g_r \;\equiv\; \max_{r \leq r' \leq R} f_{r'}$$

is the highest own floor among the exiter itself and every later exit in its market.

**Proposition 2** (Delay bounds under symmetry). *Fix a market with observed exit ranks $r = 1, \dots, R$ at times $0 < T_{(1)} \leq \cdots \leq T_{(R)}$ and active counts $n_r = N - r + 1$, impose Assumption 1, and suppose $g_r < u_r$ for every $r$ (a crossing rejects the maintained restrictions). Then the bounds on the $r$-th exiter's cost are $(g_r,\, u_r)$, and the delay bounds are (2)–(4) with the own floor $f_r$ replaced by the chain floor $g_r$:*

$$
\begin{aligned}
t_r^* - t_r^R &\in \Big[\, T_{(r)} - \Psi_{n_r}^{-1}(g_r), \quad T_{(r)} - \Psi_{n_r}^{-1}(u_r) \,\Big], \\
t_r^C - t_r^R &\in \Big[\, \min_{\theta \in [g_r, u_r]} w_{n_r}(\theta), \quad \max_{\theta \in [g_r, u_r]} w_{n_r}(\theta) \,\Big], \\
t_r^* - t_r^C &\in \Big[\, T_{(r)} - \pi_{n_r}^{-1}(g_r), \quad T_{(r)} - \pi_{n_r}^{-1}(u_r) \,\Big].
\end{aligned}
$$

*Two properties follow.*

1. *The total interval is strictly positive. For $n_r \geq 2$, the competition interval is strictly positive wherever the exiter's ceiling lies below the start-of-decline flow, $u_r < \pi_{n_r}(0)$.*

2. *The informational lower endpoint is strictly positive exactly when $g_r > f_r$, that is, when some later exit's floor exceeds the exiter's own. The last exit, with no later exit to chain from, has $g_R = f_R$ and retains the zero floor.*

*Proof.* Write $\vartheta_r \equiv \theta_{(n_r)}$ for the $r$-th exiter's cost and $\vartheta = (\vartheta_1, \dots, \vartheta_R)$ for the exiter cost vector. Under Assumption 1 the maintained restrictions say exactly this about the realized costs: $\vartheta_1 > \cdots > \vartheta_R$, with every survivor cost below $\vartheta_R$ (Lemma 4) and below $\pi_1(h)$ (Lemma 3), and $f_{r'} < \vartheta_{r'} < u_{r'}$ for every $r'$ (Lemma 2). Each delay of exit $r$ depends on the cost vector only through $\vartheta_r$, by (1) and Lemma 5.

*From costs to delays.* $T_{(r)} - \Psi_{n_r}^{-1}(\vartheta_r)$ and $T_{(r)} - \pi_{n_r}^{-1}(\vartheta_r)$ are continuous and nondecreasing in $\vartheta_r$, strictly increasing at interior crossings, so the total and informational sets are intervals whose closures have the displayed endpoints. The competition delay $w_{n_r}(\vartheta_r)$ is continuous but need not be monotone, so its set is the range of $w_{n_r}$ over $(g_r, u_r)$, an interval whose closure is the displayed min/max over $[g_r, u_r]$.

*Signs (properties 1 and 2).* The informational floor $T_{(r)} - \pi_{n_r}^{-1}(g_r)$ is nonnegative, and strictly positive iff $g_r > f_r$: $\pi_{n_r}^{-1}$ is strictly decreasing on $[f_r, \pi_{n_r}(0)]$ with $\pi_{n_r}^{-1}(f_r) = T_{(r)}$, and the floor equals $T_{(r)} > 0$ if $g_r \geq \pi_{n_r}(0)$; the last exit has $g_R = f_R$ by construction. This is property 2. For the wedge at $n_r \geq 2$, $\Psi_{n_r} < \pi_{n_r}$ puts $\Psi_{n_r}^{-1}$ strictly below $\pi_{n_r}^{-1}$ at every cost below $\pi_{n_r}(0)$, so $w_{n_r} > 0$ there, while costs at or above $\pi_{n_r}(0)$ truncate both dates at 0 and give $w_{n_r} = 0$; the competition interval is therefore strictly positive iff $u_r < \pi_{n_r}(0)$, with floor zero otherwise. The total floor is $T_{(r)} - \Psi_{n_r}^{-1}(g_r) = \big[T_{(r)} - \pi_{n_r}^{-1}(g_r)\big] + w_{n_r}(g_r)$: below $\Psi_{n_r}(0)$ the first term is nonnegative and $w_{n_r}(g_r) > 0$, and at $g_r \geq \Psi_{n_r}(0)$ the floor equals $T_{(r)} > 0$ directly, so the total interval is strictly positive in every case. This is property 1. □

The direct bounds are the relaxation of Proposition 2 obtained by replacing $g_r$ with $f_r$: this substitution reproduces (2)–(4).

## 4.3 Point identification of the total delay

When the profit function falls fast enough in the number of competitors, $n\,\pi_n(t) < \pi_1(t)$ at every date (e.g., in a Cournot quantity competition game with linear demand with profits of the form $\pi_n \propto (n+1)^{-2}$), the marginal social flow of a second firm is negative, $\Psi_2 = 2\pi_2 - \pi_1 < 0$. The profit-maximizing planner thus runs the industry as a monopoly from the outset, and the bounds degenerate: the total strategic delay is point identified and equals the observed survival time, while the two components remain interval-identified and trade off one for one along a segment of slope $-1$. Appendix B defines "$n\,\pi_n(t) < \pi_1(t)$ at every date" as Assumption 2 and proves the result above (Proposition 3).

## 4.4 Dropping the cost-ordered exits assumption

The delay bounds above use the ordering of Lemma 4: the observed $r$-th exiter is assigned cost rank $n_r$, so the benchmark maps $\pi_{n_r}^{-1}$ and $\Psi_{n_r}^{-1}$ are evaluated at the observed active count. The assignment matters because both counterfactual paths order exits by cost, so a firm's coordination and regulator dates depend on its rank in the cost order, not on its place in the observed exit sequence. Without cost-ordered exits, the two need not coincide: the observed $r$-th exiter may hold any cost rank consistent with the dominance bounds.

Dropping the cost-ordered exits assumption, therefore, makes the exiter's cost rank set-valued. Each firm still carries dominance bounds, $(f_{r'}, u_{r'})$ for the firm observed to exit $r'$-th and $(0, \pi_1(h))$ for each firm surviving to the horizon $h$, and a candidate rank for an exiter is feasible when the other firms' bounds can be arranged around its cost. For each feasible rank, the counterfactual dates are well defined, and the identified set for each delay is the union over feasible ranks and bounds-consistent costs.

Proposition 4 (see Appendix C for the statement and proof) characterizes the delay bounds when dropping the cost-ordered exits assumption. Crucially, the total and informational delays are unsigned (i.e., they can take both positive and negative values), whereas the competition delay remains weakly positive. That is, the data and the cost-ordered exits assumption contribute as summarized by the table below. From an empirical perspective, assuming symmetric strategies (and thus cost-ordered exits) provides identifying power in that the bounds become weakly positive always. In the absence of symmetry, the data do not provide an answer as to whether strategic delays existed.

| component | dominance only (Proposition 4) | dominance + symmetry (Proposition 2) |
|---|---|---|
| total $t^* - t^R$ | unsigned | $> 0$ |
| competition $t^C - t^R$ | $\geq 0$ | $> 0$ |
| informational $t^* - t^C$ | unsigned | $\geq 0$; $> 0$ where the chain binds |

We conclude the section with an example that illustrates why bounds become unsigned without cost-ordered exits.

**Example 1** (Duopoly: delays of either sign)**.** Consider a duopoly with monopoly flow $\pi_1(t) = e^{-t/10}$ and duopoly flow $\pi_2(t) = 0.75\,e^{-t/10}$, so that the marginal social flow is $\Psi_2(t) = 0.5\,e^{-t/10}$. The data record one exit at $t^* = 5$ and one survivor at the horizon $h = 6$, so the exiter's bracket is $(\pi_2(5), \pi_1(5)) = (0.455, 0.607)$ and the survivor's bracket is $(0, \pi_1(6)) = (0, 0.549)$. Fix the exiter's cost at $\theta = 0.48$ and consider two placements of the survivor's cost, each inside its bracket.

Placing the survivor at 0.30 gives the exiter the higher cost, so it holds rank two, as under the ordering of Lemma 4. Its benchmark dates are $t^C = \pi_2^{-1}(0.48) = 4.46$ and $t^R = \Psi_2^{-1}(0.48) = 0.41$, so the informational delay is $5 - 4.46 = 0.54$, the competition delay is 4.05, and the total delay is 4.59.

Placing the survivor at 0.52 instead, inside its bracket because $0.52 < \pi_1(6) = 0.549$, gives the exiter the lower cost, rank one. On both benchmark paths it outlasts its rival and leaves when the monopoly flow reaches its cost, at $t^C = t^R = \pi_1^{-1}(0.48) = 7.34$, so the informational and total delays equal $5 - 7.34 = -2.34$ and the competition delay is zero. That is, the firm conceded at date 5 a market that both benchmarks would have it hold until date 7.34.

The two cost vectors generate the same data and satisfy every dominance restriction, so the total and informational sets contain values of both signs: the data cannot separate a weak firm that overstayed from a strong firm that conceded.[2] Note that when assuming symmetric strategies (and thus that firms exit in cost order), we rule out all exit orders that deliver negative total and informational delays. Hence, the identifying power of the assumption.

# 5 The County Business Patterns application, 1986–2016

This section applies the bounds on the County Business Patterns (CBP) county panel, 1986–2016, for four declining local-oligopoly industries: motion picture exhibition, bowling centers, book stores and news dealers, and florists (U.S. Census Bureau, 2024). County establishment counts are a near-complete census through 2016, which allows us to measure exits after each industry's national peak (1986, 1987, 1993, and 1987 respectively). The panel covers 2,267 to 3,031 counties per industry, 3,096 distinct counties in all; 1,662 to 2,737 of them are active in the industry's peak year.

The four industries fit the description of a war of attrition, and Figure 1 presents motivating facts. Between their peaks and 2016, the industries lost 42 to 52 percent of their establishments, and the decline was steady over three decades (panel a). The capacity exited gradually: among counties that start with two to eight establishments, the share of peak capacity that has exited climbs roughly linearly for thirty years (panel b), so exits are staggered across decades rather than bunched into a collapse. Lastly, where a county's count falls, it falls by exactly one establishment 78 to 85 percent of the time, and 77 to 86 percent of the markets that record an exit still have an active establishment in 2016 (panel c), so the contest thins the market and leaves a survivor rather than emptying it.

[2] The negative arrangement requires the survivor's ceiling to exceed the exiter's floor, $\pi_1(h) > \pi_2(t^*)$, and it holds here because the horizon lies one year past the exit. When the horizon is distant relative to the clock, the inequality may fail, making rank one infeasible, and both sets positive.

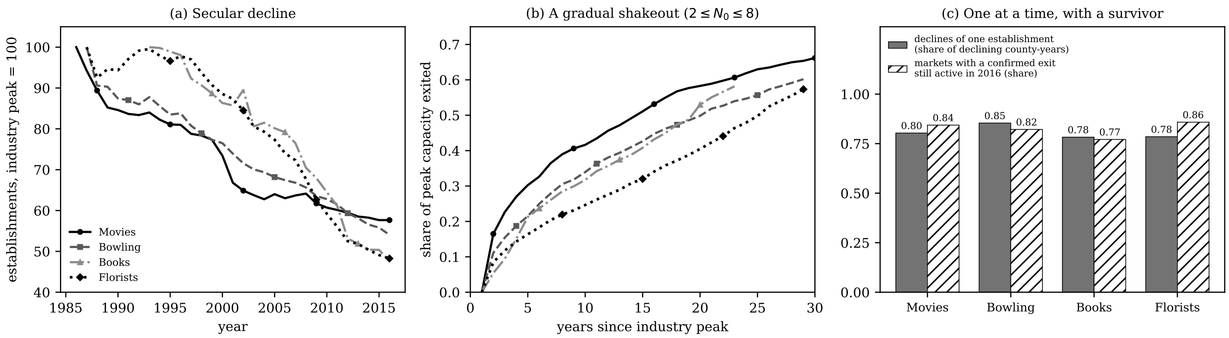


Figure 1: The shakeouts as wars of attrition. (a) National establishment counts in the county panel, indexed to each industry's peak. (b) Cumulative share of peak-year capacity confirmed exited (first passages, as in the text), counties with two to eight establishments at the peak; lines end at 2016. (c) Left bars: among county-years in which the count falls, the share falling by exactly one establishment. Right bars: among markets with at least one confirmed exit, the share still active in 2016.

## 5.1 Empirical model

The profit flow of a firm in market $m$ with $k$ active firms at time $\tau$ (in years since the industry's peak) is

$$\pi_k(m,\tau) \;=\; e^{a_m}\, k^{-\delta}\, e^{-\rho\tau}, \tag{5}$$

where $a_m$ is a county–industry level parameter capturing market size at the industry peak, $\delta$ is an industry-specific parameter capturing how profits decrease with the number of competitors in the market, and $\rho$ is a parameter governing the profit decay rate in a given industry over time. We estimate all the parameters of the profit function using two complementary approaches. First, we estimate $a_m$ for every county–industry and $\delta$ for every industry using an entry model along the lines of Bresnahan and Reiss (1991). Similar to Takahashi (2015a), we take the count of firms at the industry peak and assume it is the outcome of an entry game. A richer dataset that included quantities and prices would allow the researcher to estimate the profit function differently; no such data exists in the CBP. Second, we use the exit decision of monopolists—which lack the strategic considerations of the multi-firm game—to estimate $\rho$. We provide details in what follows.

**Entry model** Consider market $m$ at the industry's national peak, with $M$ potential entrants. Entrant $i$ pays a sunk entry cost $\kappa_i$, drawn independently from a parametric distribution $H$. We let the entry costs be a separate object from the exit opportunity costs $\theta$.

Under complete information, the cheapest entrants enter first, so the equilibrium count is

$$N_0 \;=\; \max\{\, n \leq M :\; \kappa_{(n:M)} \leq V_m(n)\,\}, \qquad \log V_m(n) \;=\; a_m - \delta \log n + \text{const}, \tag{6}$$

where $\kappa_{(n:M)}$ is the $n$-th order statistic of the $M$ entry-cost draws, and the probability distribution over the equilibrium count (conditional on the market size) is

$$\Pr\big(N_0 \geq n \mid \alpha_m\big) \;=\; \Pr\big(\mathrm{Bin}(M,\, H(u_{mn})) \geq n\big), \qquad u_{mn} = b_0 + b' x_m + \alpha_m - b_\delta \log n, \tag{7}$$

with $x_m = (\log \mathrm{pop}_m, \log \mathrm{inc}_m)$ and $\alpha_m \sim N(0, \sigma_a^2)$ a market effect.[3] In words, the market reaches $n$ firms if and only if at least $n$ of the $M$ entry-cost draws lie below the value of being one of $n$ firms, and that value falls across successive entrants by the rung $\delta \log n$ of (5). We estimate (7) by maximum likelihood on the count cells $0, \ldots, 7$ and $\geq 8$, over the 3,076 counties with peak-year population and income, including the 13 to 47 percent of counties with no establishment; the market effect is integrated by quadrature. We set $M = 11$ and assume that $H$ follows a Weibull distribution, with the Weibull scale absorbed into the intercept and the Weibull shape setting the units of the index coefficients.[4] We measure the profit function in the population units of Bresnahan and Reiss (1991), under which market size multiplies per-capita profit: log population then enters log profit with a coefficient of one, so dividing the estimated index by its coefficient $b_{\mathrm{pop}}$ on $\log \mathrm{pop}_m$ expresses the profit function in log-profit units. We use those units throughout; Appendix D reports the index coefficients and the conversion. Figure 2 shows that the fitted model tracks the count distribution closely in all four industries.

**Monopolist exit** To estimate the decay, we use the monopolists. An $N_0 = 1$ market has a single firm with an opportunity cost draw $\theta$, solving a single-agent stopping problem. Leaving at date $\tau$ is worth

$$V_m(\tau) \;=\; \int_0^{\tau} e^{-rs}\, \pi_1(m,s)\, ds \;+\; \frac{\theta}{r}\, e^{-r\tau}, \qquad V_m'(\tau) \;=\; e^{-r\tau}\big[\pi_1(m,\tau) - \theta\big],$$

the flow given up at the margin against the cost saved. The discount factor is a strictly positive common factor, so $V_m'$ carries the sign of $\pi_1(m,\tau) - \theta$, positive and then negative because $\pi_1(m,\cdot)$ is strictly decreasing. The monopolist therefore leaves at the myopic date at which its flow falls to its cost, $\pi_1(m,\tau) = \theta$, which under (5) is

$$\tau_m \;=\; \frac{a_m - \log \theta}{\rho}. \tag{8}$$

Because $\tau_m$ is decreasing in $\theta$, the $q$-th quantile of the exit date at a given market level is $\big(a_m - \log Q_\theta(1-q)\big)/\rho$, with $Q_\theta$ the quantile function of the cost distribution: for any fixed quantile the exit date is linear in the market level with slope $1/\rho$ and an intercept that depends on the quantile but not on the level, whatever the cost distribution is.

We group the monopoly markets into bins of the fitted market level, compute Kaplan–Meier exit-date quantiles within each bin, censoring at 2016, and estimate the common contour slope across bins. The market effect and the cost draw are both independent of $x_m$, so their contribution to (8) is distributed the same way in every bin; each quantile contour is then a line in the fitted market level with slope $1/\rho$ and a bin-invariant intercept, the estimated slope is consistent for the inverse of the decay rate, and a market bootstrap gives the confidence interval. Table 1 presents the estimates of $\delta$ and $\rho$.

[3] Population: Census Bureau intercensal county estimates. Per-capita income: BEA CAINC1.

[4] The Weibull distribution achieved the best fit to the data among the lognormal, log-logistic, Weibull, and Fréchet families.

Table 1: Entry-based estimates of the profit function

| Industry | $\hat{\delta}$ [95% CI] | $\hat{\rho}$ [95% CI] |
|---|---|---|
| Movie exhibition | 0.11 [0.00, 0.23] | 0.259 [0.198, 0.419] |
| Bowling centers | 0.54 [0.35, 0.84] | 0.285 [0.222, 0.441] |
| Book stores | 0.05 [0.00, 0.15] | 0.410 [0.268, 1.045] |
| Florists | 0.20 [0.17, 0.29] | 0.130 [0.094, 0.219] |

Notes: Entry model with Weibull costs and $M = 11$, Bresnahan–Reiss population units. Brackets are bootstrapped 95% confidence intervals: a county bootstrap of the count model for $\hat{\delta}$ (200 replications) and a market bootstrap of the survival-contour slope for $\hat{\rho}$ (500 replications), both on the Weibull baseline. The contour sample is the 482 to 769 monopoly markets per industry with peak-year population and income, of which 318 to 538 record an exit by 2016. See Table 4 in the Appendix for the full set of estimates.

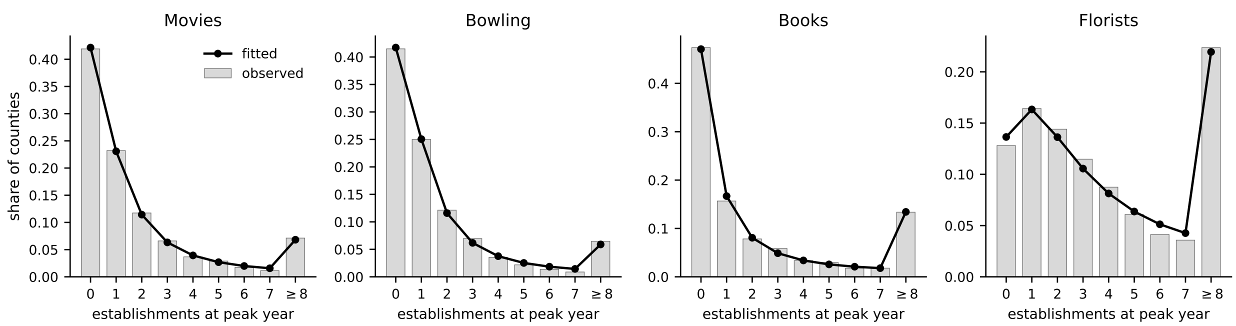


Figure 2: Observed (bars) and fitted (dots) distributions of establishment counts at the industry peak, count cells $0, \ldots, 7$ and $\geq 8$, baseline entry model (Weibull costs, $M = 11$).

## 5.2 Strategic exit delay decomposition

Table 2 reports the bounds of Proposition 2 evaluated market by market on the entry-based profit function (5); the bounds are within-market objects and use only $\hat{\delta}$, $\hat{\rho}$, and observed exit dates. The sample is the 837, 813, 694, and 1,409 markets per industry with two to six establishments at the peak.[5] An exit is recorded when the county count first falls below a threshold and remains there for a second year, which screens out one-year reporting dips, and it is timed to the first of the two years. Of these markets, 792, 741, 606, and 1,271 record at least one exit with two or more firms active, contributing the 1,135 to 2,551 exits the table averages over. The final monopolist's exit enters none of the bounds: it leaves at its myopic boundary, which is also the coordination date and the regulator date for a firm with no rivals, so its three delays are zero by construction.

[5]A market with one establishment at the peak has only a final monopolist, whose three delays are zero by construction, so monopoly markets contribute nothing to the bounds; they identify the decay instead (Section 5.1). The cap at six is the same rule in all four industries and reflects the unit of observation: the bounds read each county as one market and each establishment as one firm, and both readings weaken in larger counties, which span several local submarkets and multi-establishment firms; net counts, which date the exits, also mix closings with offsetting openings more often as counts grow. Markets above six are also scarce in movies and bowling, 27 to 37 per initial size, while every initial size from two to six has at least 42 markets in every industry.

Table 2: Strategic exit delays on the entry-based profit function (years)

| Industry | Exits | Survival | Total | Informational | Competition |
|---|---|---|---|---|---|
| Movie exhibition | 1436 | 7.6 | [0.3, 1.6] | [0.0, 1.3] | [0.2, 0.3] |
| Bowling centers | 1211 | 9.1 | [1.5, 3.9] | [0.1, 2.4] | [1.4, 2.3] |
| Book stores | 1135 | 9.0 | [0.1, 1.2] | [0.0, 1.1] | [0.1, 0.1] |
| Florists | 2551 | 11.9 | [1.1, 3.3] | [0.0, 2.3] | [1.0, 1.4] |

Notes: Means over bounded exits in markets with initial size $N_0 \leq 6$; survival is mean post-peak survival of exiting firms; totals are capped at observed survival. Each endpoint mean carries a two-sided market-bootstrap 95% interval with half-width at most 0.5 years, computed holding the profit function at the point estimates $(\hat{\delta}, \hat{\rho})$; first-stage uncertainty is taken up in the text. The delay bounds pool nothing across markets.

The maintained restrictions are refutable, and this sample does not refute them. Proposition 2 requires the chain floor to lie below the ceiling at every exit, $g_r < u_r$, and a crossing would reject dominance and cost-ordered exits jointly. Under (5) the market level cancels from the two sides, so the test turns on the curvature, the decay, and the observed dates alone. No exit in any of the four industries has $g_r > u_r$. Exact ties, $g_r = u_r$, occur at 61, 23, 21, and 24 exits, between 0.9 and 4.3 percent of each industry's total, and every one of them arises the same way: the last firm in the market is recorded as leaving in the same year as an earlier exit, and its own floor is the monopoly flow, so it coincides with that earlier exiter's ceiling at the shared date. The ties are a consequence of annual dating rather than a violation of the restrictions, and they do not survive separating the two exits within the year.

Three findings emerge. First, the war is present in every industry but explains a minority of observed survival: the total-delay floors attribute 0.1 to 1.5 years to the war, 1 to 16 percent of mean post-peak survival, and the ceilings cap it at 1.2 to 3.9 years, between a seventh (book stores) and about two fifths (bowling) of survival. Second, the competition component is strictly positive and tight, suggesting the incentive to free-ride on the exits of rivals played a role in explaining the strategic exit delay in all industries. Third, the informational component is signed but its floor stays at 0.0 to 0.1 years, because the estimated clocks are fast relative to the competition steps, so a later exit's floor rarely exceeds an earlier exiter's own, and the chain of Proposition 2 rarely binds.

All four industries sit in the interior regime of Lemma 5. Every competition curvature estimate lies below one, so the marginal social flow is positive and the planner holds each exit until decay brings $\Psi_n$ down to the firm's cost.

The remainder of post-peak survival is the decline clock, not the war: the regulator dates account for at least 5.3 (bowling) to 8.6 (florists) of the observed years. That is, most of the decades-long shakeout in Figure 1 is efficient survival under decline. The split between the components remains open: at each exit the identified set admits allocations in which competition supplies nearly all of the delay and allocations in which information supplies the majority, and selecting a point between them requires solving the equilibrium.

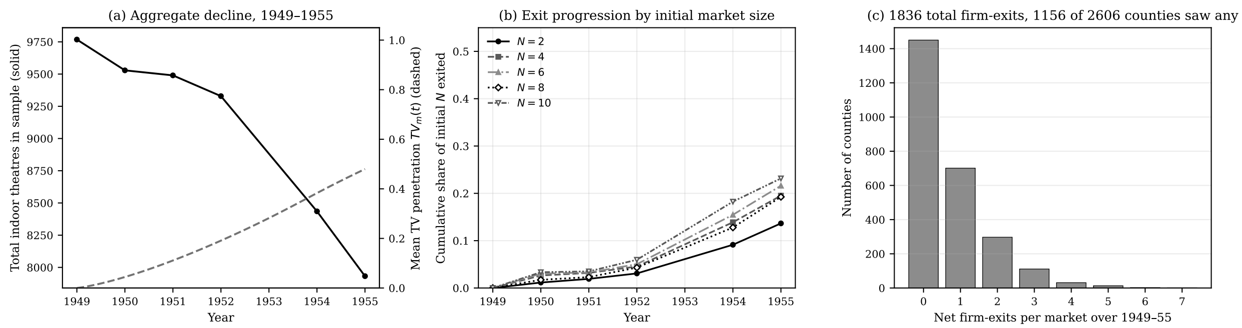


Figure 3: Descriptive evidence on the shakeout. (a) Aggregate indoor theater count across the 2,606 sample counties (blue, left axis) and the average market-level TV penetration (red dashed, right axis). (b) Cumulative share of the initial cohort that has exited, by initial market size $N$; larger markets shake out faster, consistent with stronger business stealing. (c) Distribution of firm exits per county: most counties saw no exit, and a tail of one to seven exits accounts for the 1,836 total.

# 6 The US movie-theater shakeout, 1949–1955

The second application returns to the shakeout Takahashi (2015a) studies, using the county-level theater panel from his replication package (Takahashi, 2015b). The panel covers 2,606 US counties with indoor theater counts in 1949–1952, 1954, and 1955. The aggregate count falls from 9,768 in 1949 to 7,932 in 1955 while average television penetration rises from zero to nearly one half; county counts never rise within the sample, so gross and net exits coincide at 1,836, spread over 1,156 counties, and Figure 3 summarizes the shakeout. The counts are observed at six dates, so an exit is known only up to the interval between the last date the theater was counted and the first date it was not, and the missing 1953 makes the 1952–1954 interval two years wide; the bounds below take unions over the feasible exit dates within each interval, which keeps all bounds valid under the coarse timing.

We evaluate the bounds under Takahashi's estimated profit function: base demand $(\alpha_m + \beta' X_m)\, n^{-0.242}$ and television decay $(1 - \lambda_1 TV_m(t))^{A_m}$, with the market unobservables drawn from their conditional posteriors and the bracket endpoints posterior-averaged. That is, we take the profit function his estimator delivers and drop everything else it imposes: the equilibrium selection, the parametric cost distribution, and the tie between entry and exit values.

Table 3 reports the bounds by initial market size, means over the 1,777 exits in 1,113 markets.[6] Two readings follow. First, the bounds contain the structural benchmark: the total delay is [2.4, 4.2] years overall, and Takahashi's equilibrium-selected estimate of 2.68 lies inside, obtained here without selecting an equilibrium, restricting the cost distribution, or solving the game. Second, the split between the components remains open on his data under his profit function: competition is [0.1, 3.3] and information [0.2, 3.9] overall, so the exit record is consistent with a war sustained mostly by business stealing and with one

[6]The 59 exits held out of the table are the ones that empty their market: a firm leaving with no rival active has $\pi_1$ on both arms of Lemma 2, so its cost is point identified and its three delays are zero, and the 43 counties whose only exit is of this kind leave the sample with them.

Table 3: Strategic exit delays under Takahashi's estimated profit function (years)

| $N$ | Exits | Total | Competition | Informational |
|---|---|---|---|---|
| 2 | 130 | [2.51, 4.13] | [0.29, 3.49] | [0.17, 3.50] |
| 3 | 198 | [2.51, 4.18] | [0.19, 3.48] | [0.17, 3.77] |
| 4 | 287 | [2.45, 4.22] | [0.16, 3.36] | [0.19, 3.87] |
| 5 | 245 | [2.44, 4.33] | [0.14, 3.40] | [0.18, 4.00] |
| 6 | 271 | [2.36, 4.21] | [0.14, 3.17] | [0.19, 3.90] |
| 7 | 221 | [2.24, 4.07] | [0.11, 3.12] | [0.13, 3.81] |
| 8 | 134 | [2.49, 4.36] | [0.13, 3.18] | [0.22, 4.08] |
| 9 | 159 | [2.31, 4.17] | [0.08, 3.08] | [0.20, 3.98] |
| 10 | 132 | [2.21, 4.13] | [0.08, 3.05] | [0.13, 3.92] |
| All | 1777 | [2.39, 4.20] | [0.14, 3.27] | [0.18, 3.88] |

Notes: Bounds under symmetry with interval-censored timing, means over the 1,777 exits in 1,113 markets. The profit function is Takahashi's estimated vector, held fixed: the intervals do not propagate his reported sampling uncertainty. Market-bootstrap 95% confidence intervals on the overall endpoints: competition floor [0.13, 0.15], informational floor [0.15, 0.20].

sustained mostly by private information, and closing that gap requires the equilibrium his estimator selects.

# 7 Concluding remarks

We studied what exit data from declining industries reveal about strategic exit delay, and which assumptions of the war-of-attrition exit model have identifying power. To do so, we replaced the equilibrium inversion of the structural approach with two dominance restrictions that hold in every equilibrium of the exit game: a firm does not exit while staying pays, and a firm does not stay when even monopoly profit would not cover its cost. The restrictions bound each exiting firm's cost market by market, with no assumption on how rivals play, and the bounds map into bounds on the strategic delay and its decomposition into a competition and an information component.

The bounds allocate the identifying weight as follows. Dominance alone leaves the total and information delays unsigned and puts only a nonnegative floor under the competition delay. Adding symmetric strategies, used only through the exit ordering it implies, bounds the total delay and its competition component away from zero and lifts the informational floor off zero where a later exit's floor binds. No layer of assumptions short of the equilibrium indifference condition delivers the split between the two components: the size of the delay is a property of the data and the ordering, while the division of the delay between competition and information is a property of the equilibrium. The selection question has the same structure: under a common strategy the survivors are the lowest-cost firms by construction, so whether wars of attrition select efficient survivors cannot be answered by a model that assumes symmetry.

The applications quantify the answer. In four declining industries in the County Business Patterns county panel, exiting firms survive 7.6 to 11.9 years past the industry's peak on average. The bounds attribute between a seventh and about two fifths of that survival to the war at the ceiling, and they cap the information component at 1.1 to 2.4 years; the war's share is largest in bowling, whose curvature estimate is the largest of the four. In the movie-theater shakeout of 1949–55, the bounds contain the 2.68-year structural estimate of Takahashi (2015a) under his own estimated profit function, and they do so without a stand on which equilibrium each of the 2,606 county markets was playing.

Two limitations qualify the exercise. First, the bounds take the profit function as an input, so their credibility is that of its measurement; we measured every component outside the exit game, but note that the entry and survival designs carry assumptions of their own. Second, the regulator benchmark maximizes industry profit rather than welfare because the model has no consumer side, so the bounds are silent on the consumer value of the capacity that exit withdraws. Within these limits, the implication for structural practice is direct: the size of the strategic delay may be reported without solving the model, and the split of the delay into competition and information may be presented as what it is, a property of the selected equilibrium.

# References


Jose Miguel Abito and Cuicui Chen. A partial identification framework for dynamic games. *International Journal of Industrial Organization*, 87:102915, 2023.

Andrés Aradillas-López, Amit Gandhi, and Daniel Quint. Identification and inference in ascending auctions with correlated private values. *Econometrica*, 81(2):489–534, 2013.

Timothy B. Armstrong. Bounds in auctions with unobserved heterogeneity. *Quantitative Economics*, 4(3):377–415, 2013.

Timothy F. Bresnahan and Peter C. Reiss. Entry and competition in concentrated markets. *Journal of Political Economy*, 99(5):977–1009, 1991.

Jeremy Bulow and Paul Klemperer. The generalized war of attrition. *American Economic Review*, 89(1):175–189, 1999.

Federico Ciliberto and Elie Tamer. Market structure and multiple equilibria in airline markets. *Econometrica*, 77(6):1791–1828, 2009.

Jean-Paul Décamps, Fabien Gensbittel, and Thomas Mariotti. The war of attrition under uncertainty: Theory and robust testable implications. TSE Working Paper 22-1374, Toulouse School of Economics, 2024. Revised June 2024.

Drew Fudenberg and Jean Tirole. A theory of exit in duopoly. *Econometrica*, 54(4):943–960, 1986.

Matthew Gentry and Tong Li. Identification in auctions with selective entry. *Econometrica*, 82(1):315–344, 2014.

Pankaj Ghemawat and Barry Nalebuff. Exit. *RAND Journal of Economics*, 16(2):184–194, 1985.

Philip A. Haile and Elie Tamer. Inference with an incomplete model of English auctions. *Journal of Political Economy*, 111(1):1–51, 2003.

Ken Hendricks, Andrew Weiss, and Charles Wilson. The war of attrition in continuous time with complete information. *International Economic Review*, 29(4):663–680, 1988.

Bo E. Honoré and Áureo de Paula. Interdependent durations. *Review of Economic Studies*, 77(3):1138–1163, 2010.

John G. Riley. Strong evolutionary equilibrium and the war of attrition. *Journal of Theoretical Biology*, 82(3):383–400, 1980.

Yuya Takahashi. Estimating a war of attrition: The case of the US movie theater industry. *American Economic Review*, 105(7):2204–2241, 2015a.

Yuya Takahashi. Replication data for: Estimating a war of attrition: The case of the US movie theater industry. American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor], 2015b. `https://doi.org/10.3886/E116132V1`.

U.S. Census Bureau. County business patterns: Complete county files, 1986–2023. U.S. Department of Commerce, 2024.

# A The equilibrium indifference condition

The standard estimator recovers a cost from each observed exit by inverting the equilibrium indifference condition of the war of attrition. We derive it here in the duopoly case of Fudenberg and Tirole (1986), by optimizing one firm's exit time against a fixed rival strategy, so that the object our bounds replace is explicit. No result in the body uses it.

**Firm $i$'s problem against a fixed rival.** Take $N = 2$ and fix the rival's strategy $T_j(\cdot)$, decreasing in cost. From firm $i$'s vantage the rival's exit time is random through $\theta_j \sim G$: the rival is still active at $t$ if and only if $\theta_j < T_j^{-1}(t)$, so its probability of surviving to $t$ and the density of its concession are

$$S_j(t) = G\big(T_j^{-1}(t)\big), \qquad f_j(t) = -S_j'(t) = -\,g\big(T_j^{-1}(t)\big)\,\big(T_j^{-1}\big)'(t),$$

the density nonnegative because $T_j$ is decreasing, and its concession hazard given survival is $h_j(t) = f_j(t)/S_j(t)$. Firm $i$, of cost $\theta_i$, earns the duopoly flow $\pi_2$ while both are active, becomes a monopolist earning $\pi_1$ if the rival concedes first, and collects the outside option worth $\theta_i/r$ once it exits. A monopolist of cost $\theta_i$ stays while $\pi_1$ covers $\theta_i$ and quits at its myopic date $T_1(\theta_i)$, defined by $\pi_1(T_1(\theta_i)) = \theta_i$, so its value at $s$ is

$$M(s;\theta_i) = \int_s^{T_1(\theta_i)} e^{-r(u-s)}\pi_1(u)\,du + \frac{\theta_i}{r}\,e^{-r(T_1(\theta_i)-s)}.$$

Planning to exit at $t$ unless the rival concedes first, firm $i$'s expected payoff, in present value at 0, is

$$\Pi_i(t;\theta_i) = \int_0^t f_j(s)\Big[A(s) + e^{-rs}M(s;\theta_i)\Big]ds \;+\; S_j(t)\Big[A(t) + e^{-rt}\tfrac{\theta_i}{r}\Big], \tag{9}$$

with $A(t) = \int_0^t e^{-ru}\pi_2(u)\,du$ the discounted duopoly flow earned while both are active. The first term is the rival conceding at some $s < t$, after which firm $i$ continues as a monopolist; the second is the rival still active at $t$, when firm $i$ exits and takes its outside option.

**The first-order condition.** Differentiate (9) in the exit time $t$. Using $S_j' = -f_j$ and $A'(t) = e^{-rt}\pi_2(t)$, the terms in $A$ cancel and

$$\Pi_i'(t;\theta_i) = e^{-rt}\Big\{\, f_j(t)\big[M(t;\theta_i) - \theta_i/r\big] + S_j(t)\big[\pi_2(t) - \theta_i\big]\Big\}.$$

Setting $\Pi_i' = 0$ and dividing by $S_j(t)$ gives the indifference condition at firm $i$'s optimal exit time $t = T_i(\theta_i)$,

$$\underbrace{\theta_i - \pi_2(t)}_{\text{flow loss from staying}} \quad = \quad h_j(t) \quad \underbrace{\big[\,M(t;\theta_i) - \theta_i/r\,\big]}_{\text{gain from becoming a monopolist}} \quad , \tag{10}$$

a stationary point that is firm $i$'s optimum under the regularity conditions of Fudenberg and Tirole (1986).

**What it says.** At its exit time the firm is exactly willing to wait an instant longer. The left-hand side is the flow loss it absorbs by staying: it is already past its myopic point, $\theta_i > \pi_2(t)$, so each instant of delay costs its cost net of the duopoly flow. The right-hand side is the reward for waiting: with hazard $h_j(t)$ the rival concedes, and the firm inherits the monopoly, its value rising from the exit value $\theta_i/r$ to $M(t;\theta_i) > \theta_i/r$. Equation (10) equates the marginal cost and marginal benefit of waiting. A higher-cost firm faces a larger flow loss and a smaller monopoly gain, so it exits earlier: $T_i$ is decreasing, and the exit date reveals the cost. The delay past the myopic date $\pi_2(t) = \theta_i$ reflects the informational friction: were the rival's cost common knowledge the higher-cost firm would concede at once, but against an unknown rival it waits because a concession may come.

**The equilibrium boundary.** In the symmetric equilibrium (Assumption 1) both firms use the same $T(\cdot)$, so $T_j = T_i = T$ and the type exiting at $t$ is $\theta_i = \Phi(t) \equiv T^{-1}(t)$; substituting $T_j^{-1} = \Phi$ turns (10) into a differential equation for the boundary $\Phi$, whose solution is the map the standard estimator inverts to read a cost off each exit time. Our bounds keep only the two inequalities (10) implies without being solved, the dominance bounds $\pi_2(t) < \theta_i < \pi_1(t)$ of Lemma 2: the positive flow loss on the left, and the requirement that even the monopoly continuation be worth staying for, $\theta_i < \pi_1(t)$.

# B Point identification of the total delay

Section 4.3 states that the total-delay bounds degenerate when a monopolist out-earns any larger active set. This appendix records the condition and the proof.

**Assumption 2** (Monopoly dominance)**.** At every date a monopolist earns more than the combined flow of any larger active set: $n\,\pi_n(t) < \pi_1(t)$ for all $n \geq 2$ and all $t$.

Total industry profit is then maximized by a single firm at every instant. The condition holds for the Cournot form $\pi_n(t) = A(t)/(n+1)^2$, where $n\pi_n < \pi_1$ reduces to $4n < (n+1)^2$, that is $(n-1)^2 > 0$, for every $n \geq 2$; more generally it holds whenever the per-firm flow decays in the active count faster than $1/n$. It is the complement of the regime in Lemma 5: at $n = 2$ it reads $\Psi_2 = 2\pi_2 - \pi_1 < 0$, the opposite sign to that lemma's premise, so the interior first-order condition $\Psi_n(t) = \theta_{(n)}$ has no solution and the planner chooses a corner solution.

**Proposition 3** (Point-identified total delay)**.** *Under Assumptions 1 and 2 the industry-profit-maximising schedule keeps only the lowest-cost firm and drops every other firm at the start, so $t_r^R = 0$ for every observed exit. The total strategic delay is then point identified and equals the observed survival time,*

$$t_r^* - t_r^R = T_{(r)}.$$

*Equivalently,* (2) *degenerates because $\Psi_{n_r}^{-1}$ evaluates to* 0 *at both endpoints of the cost bounds. The competition and informational components remain interval-identified as in Proposition 2; since they sum to the point-identified total $T_{(r)}$, they trade off one for one, and the identified set for the pair is a segment of slope* $-1$.

*Proof.* With an active set $S$ of size $m$ the net industry flow at $t$ is $m\,\pi_m(t) - \sum_{k\in S}\theta_k$. By Assumption 2, $m\pi_m(t) < \pi_1(t)$ for $m \geq 2$, and $\sum_{k\in S}\theta_k \geq \theta_{\min(S)}$, so replacing $S$ by its single lowest-cost member raises the net flow at every $t$. The pointwise optimum keeps at most one firm active, the globally lowest-cost firm, until $\pi_1(t) = \theta_{(1)}$; this schedule is nonincreasing in the active set over $t$ and hence feasible under irreversible exit. The planner therefore drops $\theta_{(2)}, \ldots, \theta_{(N)}$ at the start. By Lemma 4 the observed exiters are exactly these non-survivors, so each has $t_r^R = 0$ and $t_r^* - t_r^R = T_{(r)}$. The components are unchanged because $t_r^C = \pi_{n_r}^{-1}(\theta_{(n_r)})$ still depends on the cost bounds; their sum is $t_r^* - t_r^R = T_{(r)}$. □

# C Ordering-free delay bounds: statement and proof

Section 4.4 drops both Assumption 1 and the ordering of Lemma 4 and works from the dominance restrictions alone. This appendix states and proves the characterization of the resulting identified sets, and it helps to fix the logic in words first. Without the ordering, an observed exit no longer has a known cost rank: the $r$-th firm to leave may be the market's strongest firm or its weakest, so the rank the ordering would have assigned, $n_r = N - r + 1$, is only one candidate among $N$. What the data still deliver is one bracket per firm, in the sense of Proposition 1: the firm observed to exit $r'$-th carries the bracket $(f_{r'}, u_{r'})$, with bracket floor $f_{r'}$ and bracket ceiling $u_{r'}$, and each firm surviving to the horizon $h$ carries $(0, \pi_1(h))$. A candidate rank for an exiter is feasible when the other firms' brackets can be arranged around its cost, some above and some below. For each feasible rank and cost the counterfactual dates are the point objects $t^C = \pi_\nu^{-1}(\theta)$ and $t^R = \Psi_\nu^{-1}(\theta)$, with the generalized inverses of Section 4, and the identified set for each delay is the closure of the union over feasible pairs.

**Proposition 4** (Ordering-free delay bounds)**.** *Maintain the dominance restrictions and impose neither Assumption 1 nor the ordering of Lemma 4. Give each firm its bracket: exiter $r'$ carries $(f_{r'}, u_{r'})$, and each firm surviving to the horizon $h$ carries $(0, \pi_1(h))$. For exit $r$, cost rank $\nu \in \{1, \dots, N\}$ is feasible at $\theta \in (f_r, u_r)$ if and only if at most $N - \nu$ of the other firms have bracket floors at or above $\theta$ and at most $\nu - 1$ have bracket ceilings at or below $\theta$. The identified set for each delay of exit $r$ is the closure of the union over feasible $(\nu, \theta)$ of the point delays computed from $t^C = \pi_\nu^{-1}(\theta)$ and $t^R = \Psi_\nu^{-1}(\theta)$. We establish the following results.*

1. *The total and informational delays are unsigned.*

2. *The competition delay is weakly positive.*

*Proof. Part 1.* Rank-one feasibility at $\bar{\theta}$ requires no other firm's bracket ceiling at or below $\bar{\theta}$, so feasibility at $\bar{\theta}$ implies feasibility at every $\theta \in (f_r, \bar{\theta}]$, which is the downward closure. At any feasible $(1, \theta)$ both benchmark flows are $\pi_1$ and $\theta < u_r = \pi_1(T_{(r)})$, so $t^C = t^R = \pi_1^{-1}(\theta) > T_{(r)}$ and both delays are negative. The general statements follow from $\pi_\nu^{-1}(\theta) > T_{(r)}$ iff $\theta < \pi_\nu(T_{(r)})$, and its $\Psi_\nu$ analogue. If $\theta$ lies at or below the level the benchmark flow declines toward, the generalized inverse is $+\infty$, the corresponding delay is $-\infty$, and the lower endpoint of the closure is $-\infty$.

That the total and informational delay bounds include positive values follows from the fact that the cost-ordered exit sequence is always feasible and from the analysis in Section 4.

*Part 2.* $\Psi_\nu \le \pi_\nu$ pointwise implies $\{t : \pi_\nu(t) \le \theta\} \subseteq \{t : \Psi_\nu(t) \le \theta\}$, so $\Psi_\nu^{-1} \le \pi_\nu^{-1}$ and the wedge $t^C - t^R$ is nonnegative at every pair at which it is defined. At $\nu = 1$ the wedge is identically zero, because $\Psi_1 = \pi_1$. At $\theta \ge \pi_\nu(0)$ both dates are zero and the wedge is zero. For the divergence, a pair with $\lim_{t\to\infty} \Psi_\nu(t) < \theta \le \lim_{t\to\infty} \pi_\nu(t)$ has $t^R$ finite and $t^C = +\infty$; pairs with $\theta$ at or below both levels leave both dates infinite and contribute no competition value, by the convention in the statement. □

# D Entry-model estimates

Table 4 reports the full maximum-likelihood estimates behind Table 1. Writing out the components of $b$, the index of (7) is $u_{mn} = b_0 + b_{\text{pop}} \log \text{pop}_m + b_{\text{inc}} \log \text{inc}_m + \alpha_m - b_\delta \log n$, with the covariates demeaned at their industry peak-year means, so the intercept is the index of a market with mean log population and mean log income. The coefficients are measured in the units of the standardized log-cost shape (Section 5.1), and the population anchor converts them into log-profit units: the curvature is $\hat{\delta} = b_\delta / b_{\text{pop}}$, the income elasticity of profit is $b_{\text{inc}} / b_{\text{pop}}$, the market-effect dispersion is $\sigma_a / b_{\text{pop}}$, and the decay is $\hat{\rho} = \rho_\kappa / b_{\text{pop}}$, with $\rho_\kappa$ the reciprocal of the monopolist contour slope estimated on the fitted index, the same regression the text runs on the fitted market level, rescaled by $b_{\text{pop}}$.

Table 4: Entry-model estimates

| | Movie exhibition | Bowling centers | Book stores | Florists |
|---|---|---|---|---|
| *Index coefficients, standardized log-cost units* | | | | |
| Intercept $b_0$ | $-2.40$ | $-2.28$ | $-2.52$ | $-0.66$ |
| Log population $b_{\text{pop}}$ | 1.12 | 1.15 | 1.54 | 1.26 |
| Log income $b_{\text{inc}}$ | 0.79 | 2.64 | 2.17 | 0.47 |
| Competition rung $b_\delta$ | 0.12 | 0.62 | 0.07 | 0.25 |
| Market effect $\sigma_a$ | 0.51 | 1.02 | 0.78 | 0.50 |
| *Population units, decay, and fit* | | | | |
| Curvature $\hat{\delta} = b_\delta / b_{\text{pop}}$ | 0.11 [0.00, 0.23] | 0.54 [0.35, 0.84] | 0.05 [0.00, 0.15] | 0.20 [0.17, 0.29] |
| Income elasticity $b_{\text{inc}}/b_{\text{pop}}$ | 0.71 | 2.30 | 1.40 | 0.37 |
| Market effect $\sigma_a / b_{\text{pop}}$ | 0.45 | 0.89 | 0.51 | 0.40 |
| Contour reciprocal slope $\rho_\kappa$ | 0.289 [0.222, 0.469] | 0.327 [0.255, 0.507] | 0.633 [0.414, 1.613] | 0.164 [0.119, 0.277] |
| Decay $\hat{\rho} = \rho_\kappa / b_{\text{pop}}$ | 0.259 [0.198, 0.419] | 0.285 [0.222, 0.441] | 0.410 [0.268, 1.045] | 0.130 [0.094, 0.219] |
| Negative log-likelihood | 3512.4 | 3663.5 | 3158.7 | 4138.2 |
| Monopoly markets (exits) | 714 (527) | 769 (538) | 482 (388) | 505 (318) |

Notes: Count-cell maximum likelihood of (7) on cells $0, \ldots, 7$ and $\geq 8$, Weibull entry costs, $M = 11$, 3,076 counties per industry, market effect integrated by quadrature; covariates demeaned at industry peak-year means. Brackets are 95% confidence intervals: a county bootstrap of the count model for $\hat{\delta}$ (200 replications) and a market bootstrap of the contour slope for $\rho_\kappa$ (500 replications); the $\hat{\rho}$ interval divides the $\rho_\kappa$ interval by the point estimate of $b_{\text{pop}}$. The monopoly-markets row gives the contour sample and, in parentheses, the markets recording an exit by 2016.