Finnish Competition and Consumer Authority

DETECTING COLLUSION IN ENTRY

Riku Buri
Adriano De Leverano
Teemu Pekkarinen
Carlo Perroni
Vitezslav Titl
Janne Tukiainen

Working Papers 4/2026

Authors: Riku Buri, Adriano De Leverano, Teemu Pekkarinen, Carlo Perroni, Vitezslav Titl, Janne Tukiainen
Publication: Working Papers 4/2026: Detecting Collusion in Entry
Publisher: Finnish Competition and Consumer Authority
Postal address: Finnish Competition and Consumer Authority, POB 5, 00531 Helsinki, Finland
Visiting address: Lintulahdenkuja 2, 00530 Helsinki, Finland
kkv.fi

ISSN 2954-1859

Detecting Collusion in Entry*

Riku Buri†
Finnish Competition and Consumer Authority

Adriano De Leverano‡
ANAC

Teemu Pekkarinen¶
University of Vaasa

Carlo Perroni||
University of Warwick

Vitezslav Titl**
Utrecht University

Janne Tukiainen††
University of Turku

October 2026

Abstract

We show how collusion in market entry can be detected from entry choices alone, without price data. In an infinitely repeated entry game, competition and collusion generate opposite responses to persistent cost shocks: an adverse shock reduces joint entry under competition, but under collusion, it can reduce the value of punishment by more than it increases the gain from a one-shot deviation, so that the cartel must allow more entry. Monte Carlo evidence shows adequate power with a few hundred markets. In the Ohio school milk cartel, participation rises with persistent cost increases where cartel activity concentrates.

KEY WORDS: Collusion, entry, cartel screen, repeated games
JEL CLASSIFICATION: L41, C73, L13

---

*We thank Leonardo Giuffrida and Otto Toivanen for useful discussions. The views expressed in this paper are those of the authors and do not reflect those of ANAC.

†Finnish Competition and Consumer Authority, Finland; e-mail: riku.buri@kkv.fi

‡Italian National Anticorruption Authority (ANAC), Italy; e-mail: adriano.deleve@gmail.com

¶University of Vaasa, Finland; e-mail: teemu.pekkarinen@uwasa.fi

||University of Warwick, UK; e-mail: c.perroni@warwick.ac.uk

**Utrecht University, Utrecht University School of Economics, Netherlands; Charles University, Faculty of Law, Department of Economics and Empirical Legal Studies, Czechia; e-mail: V.Titl@uu.nl

††University of Turku, Finland; e-mail: janne.tukiainen@utu.fi

1 Introduction

Detecting collusive behavior among firms has long been a central challenge in industrial organization and antitrust economics. Competition authorities use statistical methods to detect collusion. Besides being preventive due to increasing the expected likelihood of detection, these methods can help in detecting and prosecuting existing cartels. Most existing cartel screens have concentrated on identifying collusion in pricing conduct or bidding behavior, using patterns that deviate from the predictions of competitive models. Phlips [1996] and Porter [2005] discuss theoretical and empirical approaches to detecting collusion, particularly in settings where firms' prices, quantities, or other market outcomes are observable. A parallel literature focuses on auctions, where collusion can be inferred from noncompetitive bidding patterns, as shown in empirical studies such as Baldwin et al. [1997], Pesendorfer [2000] and Chassang et al. [2022].

Our contribution is to develop a formal theoretical framework and practical empirical tests to detect collusion in market entry, distinguishing competitive from collusive conduct through observable entry patterns, without any price or bid data. While existing cartel detection approaches have advanced our understanding of how collusion manifests in market behavior, these methods often rely on detailed data on prices, quantities, or bids—data that may be unavailable or noisy in many real-world settings. Moreover, in some markets there is very little variation in the price data. Franchising chains, for example, often employ nationwide pricing policies. In contrast, market entry decisions are often directly observable. Entry choices are also inherently strategic and forward-looking, which makes them a potentially rich source of information about underlying conduct. Despite this, entry behavior has received relatively little attention in the literature on cartel detection. Exceptions include structural estimation studies of entry conduct under competition [e.g., Jia, 2008], but frameworks that link observed entry behavior to collusive conduct remain scarce. To our knowledge, no existing framework explicitly uses observed entry choices as a basis for testing for collusion grounded in a repeated-game interpretation of firm behavior.

We begin by characterizing competitive and collusive equilibria in a complete-information setting. We then embed the collusive allocation in an infinitely repeated entry game in which firms sustain cooperation using subgame-perfect punishment strategies. Our main result is a sign reversal in the response of joint entry to a persistent adverse profit shock. Under competition, a decrease in operating profits reduces the set of entry-cost realizations for which both firms enter and therefore lowers the probability of joint entry. Under collusion, the same shock also reduces the continuation-value loss generated by punishment. If this enforcement effect is sufficiently strong, exclusion be-

comes harder to sustain and the cartel must tolerate joint entry for a larger set of cost realizations. Thus, a positive response of joint entry to a persistent adverse profit shock is inconsistent with the maintained competitive model but can arise under collusion.

We derive this result in an infinitely repeated entry game between two symmetric firms with heterogeneous entry costs. Under competition, both firms enter whenever their entry costs are below duopoly profits. On the collusive path, the firms instead seek to allocate the market to the lower-cost firm. A firm that was supposed to stay out can deviate by entering and earning duopoly profits in the current period, but doing so triggers a temporary reversion to competitive play. Therefore, the collusive entry rule is determined by the comparison between the one-period gain from the deviation and the future loss from punishment. We characterize necessary and sufficient local conditions under which a marginal persistent decline in profits raises the collusive accommodation cutoff. Moreover, the main theoretical results and testing framework are robust to an incomplete information setting, despite private information complicating collusive enforcement, because observed entry choices do not necessarily reveal deviations.

The theoretical result leads directly to a reduced-form empirical test. Under the maintained competitive model, the response of joint entry to an adverse cost shifter is nonpositive. A statistically significant positive response therefore rejects the competitive benchmark and is consistent with the collusive mechanism identified in the theory. The test requires entry data and a persistent profit shifter, but not data on prices, quantities, or submitted bids. Our testing approach can be robust to unobserved heterogeneity as, depending on the available data, also quasi-experimental approaches are feasible.

We also provide Monte Carlo simulation evidence to demonstrate the feasibility and robustness of these tests and that they achieve reasonable power under samples typical in real-world data. Finally, we show that the test works in the well-known Ohio school milk cartel, where the cartel operated through a market allocation scheme based on respecting incumbencies, making entry restraint an integral part of the collusive agreement [Porter and Zona, 1999].

The existing literature on the detection of collusion uses almost solely price or bid data. Harrington [2008] provides a survey on the earlier approaches to detecting cartels. Prominent examples include Bajari and Ye [2003], Baldwin et al. [1997], Banerji and Meenakshi [2008], Porter and Zona [1993] and Price [2008]. Our test is related to Porter [1983] and Ellison [1994], who test for collusion by examining how prices and quantities respond to demand and cost shifters, because also we use cost shifters to identify collusion. Among more recent collusion screens Chassang et al. [2022] and Kawai et al. [2023] develop bid-pattern tests based on quasi-experimental reduced form estimation strategies as do we.

Few existing approaches use entry choices to test for collusion. First, among the several collusion tests proposed by Porter and Zona [1999], one requires no bid data: they estimate single-equation participation models for suspected members and examine the correlation of the residuals across bidders. This test requires the assumption of no unobserved heterogeneity. Second, the entry test of Conley and Decarolis [2016] is tailored to a setting of average price auctions and compares the frequency of joint entry for members in a suspect group with a set of comparable bidders in terms of various determinants of entry. Moreover, their test requires a rich set of observed firm covariates and the presence of a complementary dataset of first-price auctions. Relatedly, Kawai et al. [2023] propose a regression discontinuity based test using past allocation of public procurement contracts as the outcome variable to test for bid rotation, but that test is not directly based on entry choices and constructing the forcing variable requires bid information. Also they verify their test using the same Porter and Zona [1999] data that we use. To our knowledge, ours is the first entry-based collusion screen derived from the comparative statics of a repeated entry game. The distinguishing implication is not merely that collusion reduces entry, but that collusive entry discipline can respond to persistent profit shocks with the opposite sign from competitive entry. Moreover, our test is not limited to auction settings, but can test for market allocation type collusion in market entry games more generally.

Our approach is also related to the large literature on empirical estimation of entry games, where seminal contributions include Bresnahan and Reiss [1990, 1991], Berry [1992], Tamer [2003], Seim [2006]. That literature uses entry outcomes to recover market structure, competitive effects, and firms' profit functions. The key difference is that our objective is not to structurally estimate the primitives of the entry game, but rather use the comparative statics of an entry-game model to construct a reduced form cartel detection test.

The theoretical mechanism behind our empirical test is related to the repeated-game theory of self-enforcing collusion. Pioneered by Stigler [1964] and Friedman [1971], a central insight in the oligopoly theory is that collusion is sustained not by static optimality but by dynamic behavior and incentives; firms refrain from profitable deviations because such deviations reduce future continuation payoffs. This logic underlies the price-war models of Green and Porter [1984] and the general theory of repeated games with imperfect monitoring developed by Abreu [1988], Abreu et al. [1990], and Fudenberg et al. [1994]. In our setting, the relevant deviation is not a secret price cut but over-entry: a firm that is supposed to stay out can enter and earn duopoly profits today, but doing so weakens future collusive profits via punishments.

Our main comparative static is closely related to the literature on collusion over the

business cycle. Rotemberg and Saloner [1986] show that collusive prices may become more competitive in booms because high current demand increases the one-shot gain from deviation. Haltiwanger and Harrington [1991] emphasize that anticipated movements in profitability also affect the value of future punishment, so that the sustainability of collusion depends on where the industry is in the cycle. Kandori [1991] and Bagwell and Staiger [1997] extend this logic to serially correlated demand shocks, while Staiger and Wolak [1992] and Fabra [2006] show that capacity constraints can alter, and in some cases reverse, the cyclical pattern of collusive conduct. Our contribution is to apply the same incentive-constraint logic to the extensive margin of market participation. A persistent adverse profit shock reduces the static profitability of joint entry under competition, but under collusion it may reduce the continuation-value loss from punishment even more, forcing the cartel to relax exclusion and tolerate more joint entry.

The model is also related to theoretical work on collusion with private information and market allocation. In Athey and Bagwell [2001], productive efficiency under private costs requires high-cost firms to relinquish market share today in exchange for future market-share favors. Athey et al. [2004] and Athey and Bagwell [2008] study related tradeoffs when private cost shocks affect collusive pricing and when those shocks are persistent. This connection is natural in our environment because collusion in entry similarly allocates market participation across firms with heterogeneous short-run costs. Relatedly, Harrington and Skrzypacz [2011] show how communication and transfers can sustain cartel agreements when prices and quantities are privately observed. Our incomplete-information extension follows this tradition by showing that the entry-based comparative static can survive when participation costs are private but deviations become verifiable, or sufficiently well monitored, over time.

Finally, the paper is connected to the literature on collusion through market allocation and entry restraint. Bernheim and Whinston [1990] show how multimarket contact can facilitate collusion across markets, while Byford and Gans [2014] study collusion at the extensive margin, where firms sustain collusion by avoiding entry into each other's markets. Vasconcelos [2004] analyzes how entry affects cartel stability. Our focus differs from these papers: we use the comparative statics of collusive entry discipline to derive a detection test based on observed entry responses to persistent profit shocks.

2 Theoretical Framework

We study an infinitely repeated market-entry game between two symmetric firms, A and B. In each period, the firms simultaneously choose whether to enter the market. Prior to the entry decision, each firm $j \in \{A, B\}$ draws an entry cost k_j. In each period, a

new entry-cost vector (k_A, k_B) is drawn. The vectors are independently and identically distributed across periods, and their two components are independently and identically distributed according to a continuous distribution F with support on $\mathbb{R}_+$. The realized entry-cost vector (k_A, k_B) is observed by both firms before they choose entry. In Section 2.3, we extend the analysis to an environment in which each firm privately observes its own entry cost.

Let $y_j \in \{0, 1\}$ denote firm j's entry decision, where $y_j = 1$ indicates entry. If exactly one firm enters, it earns gross monopoly profit $\pi_M > 0$; if both firms enter, each earns gross duopoly profit $\pi_D \in (0, \pi_M/2)$. If a firm does not enter, it earns zero. Net payoffs equal gross profits minus the realized entry cost when $y_j = 1$. Firms discount future payoffs with a common discount factor $\delta \in (0, 1)$.

2.1 Competitive Equilibrium

Under complete information and competitive behavior, firms choose entry strategies to maximize individual expected profits. A firm j's payoff from entering, conditional on the opponent's strategy, y_{-j}, is

$$\phi_j(y_{-j}, k_j) = (1 - y_{-j})\pi_M + y_{-j}\pi_D - k_j.$$

Allowing for mixed strategies, define

$$\Phi_j(q_{-j}, k_j) = q_{-j}\phi_j(1, k_j) + (1 - q_{-j})\phi_j(0, k_j),$$

where q_{-j} is the probability that the opposing firm enters. In equilibrium, firm j chooses $q_j \in [0, 1]$ to satisfy

$$q_j\Phi_j(q_{-j}, k_j) \geq 0, \qquad (1 - q_j)\Phi_j(q_{-j}, k_j) \leq 0.$$

When $k_A, k_B \in (\pi_D, \pi_M)$, the entry subgame has two strict pure-strategy Nash equilibria and one mixed-strategy equilibrium. We select, among the Nash equilibria, the equilibrium that maximizes total industry profit. Hence, in this region, the lower-cost firm enters while the higher-cost firm stays out.[1] With this refinement, the selected pure equilibrium strategies are characterized by

$$y_j\phi_j(y_{-j}, k_j) \geq 0, \qquad (1 - y_j)\phi_j(y_{-j}, k_j) \leq 0.$$

[1]This type of equilibrium-selection assumption is standard in the empirical industrial-organization literature; see, for example, Seim [2006]. Alternatively, one could avoid point selection using the partial-identification approach of Tamer [2003]. For our purposes, the selection is innocuous because the joint-entry region is not affected by how the unilateral-entry equilibria are selected.

2.2 Collusive Equilibrium

We now turn to collusion sustained by dynamic incentives in an infinitely repeated setting. Firms can support cooperative behavior by conditioning future entry decisions on past play.

Candidate collusive entry rules. For any accommodation cutoff $k \in [0, \pi_D]$, consider the following collusive entry rule. If both firms' entry costs are below k, both firms enter. Otherwise, the lower-cost firm enters unilaterally whenever monopoly entry is profitable. A lower cutoff therefore represents stricter collusive exclusion and less joint entry. When $k < \pi_D$, the rule features less joint entry and more unilateral entry than the competitive benchmark.

To be precise, let

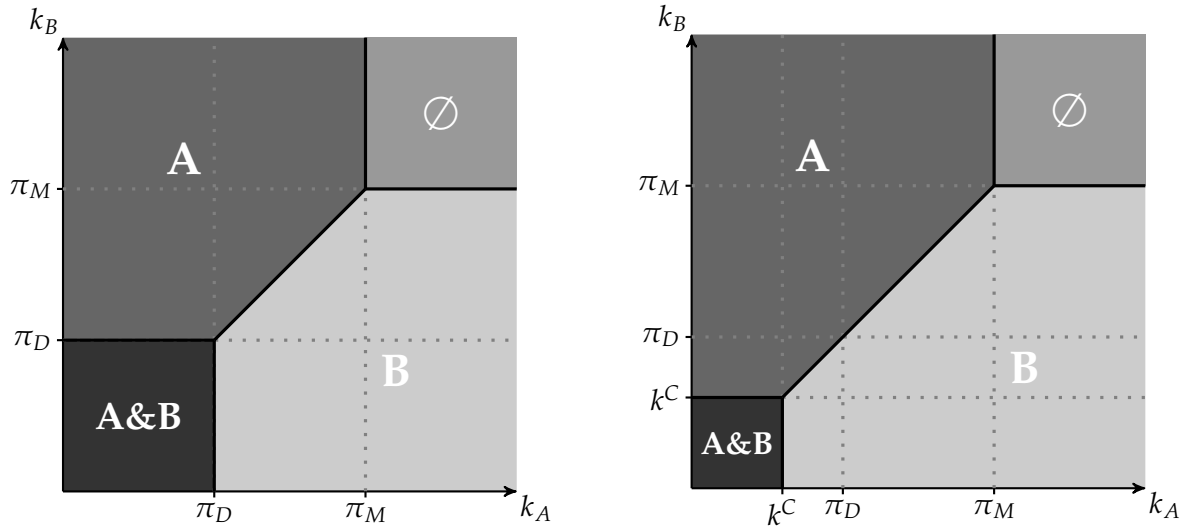
$$k_{(1)} = \min\{k_A, k_B\}, \qquad k_{(2)} = \max\{k_A, k_B\}.$$

If $k_{(1)} \leq \pi_M$ and $k_{(2)} \geq k$, the lower-cost firm enters while the higher-cost firm stays out. If instead $k_{(2)} < k$, both firms enter. Under competition, joint entry occurs whenever $k_A \leq \pi_D$ and $k_B \leq \pi_D$, so the competitive joint-entry cutoff is π_D. Thus, any collusive rule with $k < \pi_D$ replaces some joint-entry outcomes with unilateral monopoly entry and raises expected industry profit before taking incentive constraints into account.

Figure 1 plots the entry-cost realizations (k_A, k_B) and the corresponding entry outcomes. In each panel, the dark bottom-left square indicates joint entry: under competition the cutoff is π_D, while under the selected collusive rule it is k^C. Outside the joint-entry region, at most one firm enters; the shading distinguishes which firm enters, with the lower-cost firm selected in the unilateral-entry region, and the upper-right square $(k_A > \pi_M,\ k_B > \pi_M)$ corresponds to no entry.

Punishments. Consider the following public strategy profile. Following any public history without a deviation, the firms play the candidate collusive entry rule. Following a detected deviation, they play the selected competitive equilibrium for Z periods. Any deviation during the punishment phase restarts the Z-period punishment. After Z consecutive punishment periods without a further deviation, the firms return to collusive play.

The competitive continuation is sequentially rational because it consists of stage-game Nash equilibria. On the collusive path, a low-cost firm prescribed to enter unilaterally earns $\pi_M - k_{(1)} \geq 0$ and therefore does not gain by staying out. A firm prescribed to enter jointly has $k_j \leq k \leq \pi_D$ and hence does not gain by staying out. The potentially binding deviation is therefore over-entry by a firm that is prescribed to stay out.



(a) Competitive equilibrium.

(b) Collusive equilibrium.

Figure 1: Entry decisions under competition and collusion.

Let $\Pi^C(k;\pi_D)$ and $\Pi^N(\pi_D)$ denote a representative firm's ex-ante expected per-period profits under the candidate collusive rule and the selected competitive equilibrium, respectively. The discounted loss from triggering Z periods of competitive reversion is

$$\Gamma(k,\pi_D) = H_0\big(\Pi^C(k;\pi_D) - \Pi^N(\pi_D)\big), \qquad H_0 := \sum_{t=1}^{Z} \delta^t = \frac{\delta(1-\delta^Z)}{1-\delta}.$$

By symmetry, this punishment loss is the same for both firms. As shown in Appendix A, $\Pi^C(k;\pi_D) \geq \Pi^N(\pi_D)$ for every $k \in [0,\pi_D]$. Hence restarting the punishment phase weakly lowers a firm's continuation payoff. Because the prescribed actions during punishment constitute a stage-game Nash equilibrium, and a deviation would additionally restart the punishment, no one-shot deviation during the punishment phase is profitable.

Sustainability and selection. Consider a history in which firm j is prescribed to stay out while the other firm enters. If firm j complies, it earns zero in the current period. If it deviates by entering, it earns $\pi_D - k_j$ and triggers the punishment phase. Its incremental gain from deviating is therefore

$$\Delta(k_j;k,\pi_D) = \pi_D - k_j - \Gamma(k,\pi_D).$$

A firm prescribed to stay out has $k_j \geq k$. Because $\Delta(k_j;k,\pi_D)$ is decreasing in k_j, the marginal excluded firm has the strongest incentive to deviate. Define

$$\Psi(k,\pi_D) := \pi_D - k - \Gamma(k,\pi_D).$$

We call a cutoff k *sustainable* if the associated collusive entry rule can be supported by the specified punishment strategy. By the one-shot deviation principle and the preceding observations, the cutoff k is sustainable if and only if

$$\Psi(k, \pi_D) \leq 0.$$

Accordingly, let

$$\mathcal{S}(\pi_D) := \{k \in [0, \pi_D] : \Psi(k, \pi_D) \leq 0\}$$

denote the set of sustainable cutoffs. This set is nonempty because $k = \pi_D$ is always sustainable: at that cutoff the collusive entry rule coincides with the selected competitive allocation, so $\Gamma(\pi_D, \pi_D) = 0$ and $\Psi(\pi_D, \pi_D) = 0$.

We define the cartel-selected accommodation cutoff as

$$k^C(\pi_D) := \min \mathcal{S}(\pi_D). \tag{C}$$

Under the continuity conditions imposed in Appendix A, $\mathcal{S}(\pi_D)$ is closed, so the minimum exists. If $k^C(\pi_D) > 0$, the marginal exclusion constraint binds:

$$\Psi\big(k^C(\pi_D), \pi_D\big) = 0.$$

Indeed, if the inequality were strict, continuity would imply that a slightly smaller cutoff is also sustainable, contradicting the definition of $k^C(\pi_D)$. At the boundary $k^C(\pi_D) = 0$, the constraint may hold strictly.

Appendix A also shows that expected collusive profit is decreasing in the accommodation cutoff. Hence, among the sustainable cutoff rules considered here, both firms prefer the smallest cutoff, and (C) selects the cartel-optimal rule.

Comparative statics. To represent a persistent profit shock, take π_D as the public profit index and allow monopoly profit to co-move with it according to a differentiable increasing function $\pi_M = \pi_M(\pi_D)$. Thus, an adverse shock is a decrease in π_D, accompanied by a decrease in $\pi_M(\pi_D)$. The derivative Γ_{π_D} below includes both the direct effect through duopoly profit and the indirect effect through monopoly profit.

Under competition, the probability of joint entry is

$$P^N(\pi_D) = F(\pi_D)^2,$$

which is increasing in π_D. Under collusion, the probability of joint entry is

$$P^C(\pi_D) = F\big(k^C(\pi_D)\big)^2.$$

Hence, an adverse shock can increase joint entry under collusion precisely when the selected accommodation cutoff rises as π_D falls.

Proposition 1. *Suppose that Γ is continuously differentiable, that F is strictly increasing at π_D and $k^C(\pi_D)$, and that the cartel-selected cutoff is interior and regular:*

$$k^C(\pi_D) \in (0, \pi_D) \qquad \textit{and} \qquad 1 + \Gamma_k\big(k^C(\pi_D), \pi_D\big) > 0.$$

A marginal permanent decrease in duopoly profits lowers the probability of joint entry under competition. Locally, under collusion, the same shock raises the probability of joint entry if and only if

$$\Gamma_{\pi_D}\big(k^C(\pi_D), \pi_D\big) > 1.$$

The intuition is as follows. Under competition, joint entry requires both firms to find duopoly entry profitable. Lower duopoly profits therefore shrink the set of entry-cost realizations for which both firms enter.

Under collusion, joint entry occurs when exclusion is not supportable. Exclusion is sustained by comparing the one-shot gain from over-entry with the discounted loss from triggering punishment. A decrease in π_D reduces the one-shot gain from over-entry one-for-one. If it reduces the punishment loss by more than this, so that $\Gamma_{\pi_D} > 1$, exclusion becomes harder to sustain and the cartel must accommodate joint entry for a larger set of cost realizations. Appendix A proves the proposition, establishes that the local comparative-static branch remains the selected smallest sustainable cutoff, and gives conditions under which such a cutoff exists.

This mechanism is the entry-margin analogue of the cyclical-collusion literature. In those models, the effect of demand or cost movements on collusive conduct depends on how the shock changes the incentive constraint: the current gain from deviation must be compared with the future loss from punishment. Rotemberg and Saloner [1986] show that high current demand can make collusion harder to sustain because the deviation gain is large. Haltiwanger and Harrington [1991] emphasize that anticipated future profitability also matters, since it affects the value of the punishment phase. Subsequent work extends this logic to serially correlated shocks [Kandori, 1991, Bagwell and Staiger, 1997] and to capacity-constrained environments in which the sign of the cyclical pattern may change [Staiger and Wolak, 1992, Fabra, 2006]. Our comparative static works through the same dynamic incentives, but from the entry perspective: an adverse persistent profit shock may reduce the punishment loss by more than it reduces the one-shot gain from over-entry. Consequently, exclusion becomes harder to sustain, and joint entry rises.

2.3 Incomplete-information robustness

Our baseline analysis assumes that entry costs are commonly known. In many settings, however, each bidder or firm privately knows its own short-run participation cost. Pri-

vate entry costs complicate collusive enforcement because observed entry choices do not necessarily reveal deviations. Appendix B shows that the identifying comparative static can nevertheless survive under particular verification technologies.

Competition. With privately observed entry costs, the competitive entry game admits a symmetric cutoff equilibrium: each firm enters if and only if its cost is below a threshold that depends on the public profit shifters. An increase in marginal cost c lowers operating profits, so the competitive cutoff falls and joint entry weakly decreases. Hence, even under incomplete information, an adverse cost shock cannot generate a positive relationship between c and joint entry within the maintained competitive model.

Collusion. Incomplete information matters for enforcement because the public entry profile does not perfectly reveal whether a firm deviated, for example by entering with a high private cost. Nevertheless, collusion can be sustained in monitoring environments in which costs, or evidence about costs, become verifiable with some delay, for example through occasional audits, procurement scrutiny, or ex-post documentation. In such environments, the cartel can implement a cutoff-style entry rule and punish detected over-entry. The collusive cutoff may then move in the opposite direction from the competitive cutoff: when c rises, the continuation-value wedge supporting discipline can shrink because punishments become less severe when future profits are lower, inducing the cartel to relax entry discipline and tolerate more joint entry.

Implication. The competitive sign restriction extends to the private-cost benchmark analyzed in Appendix B. On the collusive side, Appendix B provides particular verification environments in which joint entry increases after an adverse cost shock. Thus, the opposite comparative static can survive private information, although the result is an existence statement for those verification environments rather than a general characterization of private-information collusion.

3 Empirical Identification and Test Design

Having formalized competitive and collusive models of entry, we now turn to the question of how to empirically detect collusive conduct using observed data. The data consist of a sample of entry decisions across markets or time periods, accompanied by covariates that plausibly affect profits or entry costs. Our goal is to test whether observed behavior is consistent with competitive conduct or instead provides evidence against the competitive model.

Formally, suppose we observe N market outcomes:

$$\{(y^t, w^t),\ t = 1, \ldots, N\},$$

where $y^t = (y_A^t, y_B^t)$ are the entry choices in market t, and w^t is a vector of covariates. Let

$$\rho(c, w) = \Pr\big(y_A^t = 1, y_B^t = 1 \mid c, w\big)$$

denote the probability of joint entry conditional on a persistent marginal-cost shifter c and other observable market characteristics w.

Under competition, an increase in marginal cost c reduces both π_M and π_D, and therefore weakly reduces the probability of joint entry:

$$\frac{\partial \rho(c, w)}{\partial c} \leq 0.$$

Under collusion, however, an increase in marginal cost also reduces the continuation-value loss from punishment. This can increase the accommodation cutoff k^C and thereby raise the probability of joint entry; see Proposition 1.

Our empirical test is therefore based on the one-sided population restriction

$$H_0 : \frac{\partial \rho(c, w)}{\partial c} \leq 0 \qquad \text{against} \qquad H_1 : \frac{\partial \rho(c, w)}{\partial c} > 0.$$

A statistically significant positive relationship between joint entry and a persistent adverse cost shifter rejects the maintained competitive model and is consistent with the collusive mechanism identified in the theory. In our empirical analysis, we exploit this result by examining the relationship between costs and entry.

Although our focus is on entry, the same logic applies to other dynamic collusion environments such as pricing or investment. For instance, observing price reductions in response to permanent cost increases, but not to temporary ones, is a possible signal of collusive behavior. This suggests that tests that are based on asymmetric responses to shocks with different persistence could be applied in possibly many different empirical contexts.

3.1 Monte Carlo simulation

We proceed with a Monte Carlo simulation. In what follows, we set $p = D(q) = a - q$ with $a = 11/5$. As a reference point, when marginal cost equals $c = 1/5$, monopoly and duopoly profits are $\pi_M = 1$ and $\pi_D = 4/9$, respectively. In the Monte Carlo simulations,

Table 1: Thresholds and joint entry probabilities

c	π_D	k_C	Pr(joint entry | com.)	Pr(joint entry | col.)
0.20	0.444	0.195	0.198	0.038
0.22	0.436	0.208	0.190	0.043
0.24	0.427	0.221	0.182	0.049
0.26	0.418	0.233	0.175	0.054
0.28	0.410	0.245	0.168	0.060
0.30	0.401	0.257	0.161	0.066

Notes: π_D is duopoly profit under linear demand $p = a - q$ with $a = 11/5$. k_C is the collusive cutoff. Entry costs are i.i.d. Uniform$[0, 1]$.

marginal costs vary across markets and are drawn independently from the uniform distribution on $[0, 2/5]$. We set $H_0 = \delta \frac{1-\delta^Z}{1-\delta} \approx 7$.[2] Entry costs are uniformly distributed on $[0, 1]$.

The assumed data structure is as follows. We observe a total of N entry decisions. Each observation corresponds to a market with a permanent marginal-cost level c. Marginal cost varies across markets but remains constant over the repeated interaction relevant for sustaining collusion.

Because each market's cost level is permanent, firms expect the current monopoly and duopoly profits to continue throughout the punishment phase. The stationary punishment value used in the theoretical calculation therefore applies directly.

In this setting our model predicts the following: in the competitive market, the probability of joint entry decreases after an adverse cost shock, while in the collusive market joint entry can either decrease or increase.

Table 1 illustrates this mechanism by varying marginal costs from 0.2 to 0.3 and reporting the implied probabilities of joint entry under competition and collusion. A 10% increase in marginal cost reduces duopoly profits by approximately 2% while increasing the collusive entry threshold by about 7%. At the other extreme, a 50% increase in marginal cost reduces duopoly profits by about 10% and increases the entry threshold by roughly 32%.

To estimate the relationship between marginal cost and joint entry, we estimate separate probit models for competitive and collusive markets:

$$\Pr(y_{ir} = 1 \mid c_{ir}) = \Phi(\beta_{0r} + \beta_{1r} c_{ir}), \qquad r \in \{\text{com}, \text{col}\},$$

[2]This is, for example, supported by a discount factor of $\delta = 0.95$ and the number of punishment periods equal to 9 (see Appendix A).

where i indexes markets, y_{ir} equals one if both firms enter market i, c_{ir} is the marginal cost in that market, and Φ denotes the standard normal cumulative distribution function. The models are estimated separately by maximum likelihood for the two regimes.

The sign of β_{1r} determines the direction of the relationship between marginal cost and the probability of joint entry. Under competition, the model predicts $\beta_{1,\text{com}} < 0$. Under the collusive calibration used in the simulation, it predicts $\beta_{1,\text{col}} > 0$. The latter prediction is specific to the chosen calibration, since under other collusive parameterizations the response may be nonpositive, resulting in a false negative.

In each Monte Carlo replication, we apply the one-sided test introduced in the previous section. At the 5 percent significance level, the null hypothesis is rejected when the estimated coefficient-to-standard-error ratio exceeds the one-sided critical value of 1.645. The rejection frequency under the competitive data-generating process is the empirical false-positive rate. The rejection frequency under the collusive data-generating process is empirical power, and its complement is the false-negative rate. We set the number of Monte Carlo replications to 1000. We let the number of observations (i.e., entry events) vary from 200 to 1,000.

To evaluate the finite-sample size of the test, we additionally simulate the least-favorable boundary of the null hypothesis, where the population relationship between marginal cost and joint entry is zero. This distinction is important because, under the baseline competitive data-generating process, higher marginal costs reduce the probability of joint entry, so the population coefficient on costs is strictly negative. A false positive would therefore require sampling variation not only to offset this negative relationship but also to produce a sufficiently large positive estimate, which becomes increasingly unlikely as the sample size grows. The conventional five-percent rejection probability instead applies at the boundary of the null, where joint entry is independent of marginal cost. In the boundary simulation, marginal costs are drawn from the same uniform distribution as in the baseline simulations, but joint entry is generated with a constant probability equal to the competitive joint-entry probability at $c = 1/5$. For each simulated sample, we estimate the same probit model and apply the one-sided five-percent rejection rule. Empirical size is calculated as the proportion of the 1,000 boundary-null replications in which the estimated coefficient-to-standard-error ratio exceeds 1.645.

The rejection frequencies are presented in Table 2. Across all sample sizes, the empirical false-positive rate under competition is zero, meaning that none of the 1,000 competitive Monte Carlo replications incorrectly rejects the null. At the least-favorable boundary of the null, the empirical size ranges from 0.053 to 0.060 across the reported sample sizes, remaining close to the nominal five-percent level. Under collusion, the test has a power

Table 2: Monte Carlo performance of the one-sided collusion screen

Observations	Empirical size	False-positive rate	Power	False-negative rate
200	0.060	0.000	0.731	0.269
400	0.054	0.000	0.938	0.062
600	0.053	0.000	0.987	0.013
800	0.054	0.000	0.998	0.002
1000	0.053	0.000	0.998	0.002

Notes: Results use 1,000 Monte Carlo replications. The first column reports the number of observations in each separately simulated regime. Marginal cost is drawn independently across markets from the continuous distribution Uniform$[0, 0.40]$ and is permanent within each market's repeated interaction. The one-sided test rejects at the 5 percent level when the estimated coefficient-to-standard-error ratio exceeds 1.645. Standard errors use HC1. Empirical size is the rejection frequency under the least-favorable boundary null, where the population cost coefficient equals zero. The rejection frequency under the competitive data-generating process is the empirical false-positive rate. The rejection frequency under the collusive data-generating process is empirical power, and its complement is the false-negative rate.

of 0.731 with 200 observations, which rises to 0.938 with 400 observations and to 0.987 with 600 observations. Correspondingly, the false-negative rate falls from 0.269 at 200 observations to 0.062 at 400 observations and 0.013 at 600 observations. With 800 or more observations, the test detects the collusive alternative in almost all replications. Overall, the Monte Carlo results indicate that the proposed screen has a very low false-positive rate and achieves high power at empirically relevant sample sizes.

3.2 Remarks and Extensions

As discussed in Section 2, the frequency of duopoly can either decrease or increase in the collusive equilibria after marginal costs increase. It is possible to parametrize the model in a way that a cost shock has negative effect on the collusive entry threshold. This is the main limitation of our proposed test: it can fail to detect a collusive market i.e. a false negative. However, false positives are not possible since with any parametrization in a competitive market the relationship between costs and entry cannot be positive. We think that this feature makes our test particularly compelling for competition authorities to apply in practice. Ortner et al. [2026] emphasize the benefits of using safe tests for cartel detection, which are designed to minimize the risk of falsely flagging competitive firms. Our test fits this description in that rejection is based on a pattern that is incompatible with competitive behaviour.

A key determinant of the statistical power of the test is the amount of variation in

marginal costs across markets. In the baseline simulation, marginal costs are drawn independently from a uniform distribution between zero and two fifths. Greater variation in marginal costs makes the estimated relationship between costs and joint entry more precise and therefore increases the power of the test. Conversely, a narrower cost distribution reduces power.

The analysis can be extended in several directions without changing the identification logic. First, instead of relying on cross-market variation in marginal costs, the identification strategy can be implemented using a differences-in-differences design across markets exposed to a common cost shock. In this setting, competitive and collusive markets exhibit opposite responses of joint entry probabilities to an increase in marginal costs. This approach is particularly useful when markets can be credibly classified into treatment and control groups, but cross-market cost variation is limited.

Second, the cost process can be generalized to allow for both permanent and transitory shocks. When marginal cost innovations contain a transitory component, firms base entry decisions on expected future profits rather than contemporaneous profits alone. This dampens the magnitude of the response of joint entry to cost changes but does not alter its sign, provided that a nontrivial permanent component remains. As a result, the proposed test remains informative, though statistical power may be reduced relative to the baseline case of fully permanent shocks.

Third, the framework allows for correlation between observed covariates and unobserved marginal costs. Measurement error or imperfect observability of marginal costs attenuates estimated effects and reduces precision, but does not generate sign reversals in the relationship between costs and joint entry. Consequently, while power may be weakened, the test continues to distinguish competitive from collusive behavior under mild regularity conditions.

Fourth, the analysis can accommodate correlation in entry costs across firms or over time. Such correlation affects the level of joint entry probabilities and the degree of simultaneity in entry decisions, but does not overturn the comparative statics with respect to marginal costs. As long as the distribution of entry costs is stable across regimes, the sign of the estimated cost–entry relationship remains informative about the underlying conduct.

Finally, an additional extension concerns markets with more than two potential entrants. In such settings, the outcome of interest need not be joint entry, but rather the total number of firms entering the market. The identification logic extends naturally: under competition, an increase in marginal costs weakly reduces the expected number of entrants, whereas under collusion higher marginal costs may relax incentive constraints and lead to a higher tolerated level of entry. Accordingly, the sign of the relationship

between marginal costs and the number of entrants remains informative about the underlying conduct. With more than two firms, the outcome of interest is no longer binary. A natural empirical extension is therefore to model the number of entrants using either a linear specification or an ordered logit model.

4 Validation on the Ohio School Milk Cartel Data

4.1 Mapping the Model to the Data

To validate our collusion detection test, we apply it to data from the Ohio school milk cartel. This well-known cartel case in public procurement was analyzed by Porter and Zona [1999].[3]

Before turning to the data, we detail how the entry game of Section 2 maps into the Ohio school milk cartel setting. First, in procurement auctions, the entry decision is the decision to submit a bid. Bidding requires information acquisition and preparation costs that vary across contracts and firms, corresponding to the entry-cost draws in the model. The market outcome we observe in each auction—how many firms participate—is the empirical counterpart of the entry profile in the model.

Second, the Ohio case is well suited to an entry-based test because entry restraint was an integral part of the documented scheme. According to the trial testimony summarized in Porter and Zona [1999, p. 269], the Cincinnati defendants allocated school districts by respecting incumbencies: if a member had served a district in the previous year, the others were to either submit high complementary bids that would not undercut the incumbent, or else refrain from bidding. Consistent with this, Porter and Zona [1999] find that the defendants' bid-submission decisions deviate systematically from those of a competitive control group, and that the defendants competed aggressively in distant districts while showing restraint near their plants, a pattern they interpret as territorial allocation of the districts close to the dairies' plants. This incumbency-based market allocation is precisely the object our model formalizes: a rule prescribing which firm participates in which market, sustained by the threat of losing future allocations. The prescribed abstention of non-incumbent members corresponds to the exclusion of the high-cost firm on the collusive path, and entry into a fellow member's district corresponds to the over-entry deviation whose incentive constraint drives Proposition 1.

Third, the cost shifter in the model is a persistent shock. Raw milk is the dominant marginal cost of serving a school milk contract, and the federal marketing order price

[3]We obtained the data from Wachs and Kertész [2019].

exhibits substantial serial correlation [Nicholson and Stephenson, 2015]. Therefore, price movements shift firms' expectations of future profitability and hence the continuation values that sustain collusive discipline.

Fourth, the price of raw milk is plausibly exogenous to any individual auction. Through the mid-1980s, the Commodity Credit Corporation's price-support program, administered by the U.S. Department of Agriculture, kept manufacturing-grade milk prices close to an administratively set floor, and after 1987 market prices rose above that floor. Throughout the studied period, the federal marketing-order price equals the Minnesota–Wisconsin manufacturing-milk price plus a fixed location-specific differential. Its variation therefore reflects national dairy-market conditions and federal policy rather than procurement decisions by individual Ohio school districts. Reverse causality from local entry decisions or collusive conduct to the input price appears unlikely as a small district's auction or a cartel's conduct cannot move the price.

Two features of the setting fall outside the model. First, the cartel also employed complementary bidding, and thus, a submitted bid need not represent genuine competitive entry. We discuss the implications for our test among the caveats below. Second, the model has two symmetric firms, whereas the auctions here attract up to six bidders. As discussed in Section 3.2, the identification logic extends to the number of entrants, which is the outcome we use. With these qualifications, the Ohio data provide a demanding but appropriate testing ground: a documented market-allocation cartel, an observable and persistent cost shifter, and entry decisions recorded for both cartelized and non-cartelized auctions.

4.2 Data Description

We use two data sources. The main dataset from Porter and Zona [1999] contains information on school milk procurement contracts in Ohio awarded between 1980 and 1990. It reports the contract ID, the year of publication, the submitted bid, an indicator for whether the bid is a winning bid, and bids adjusted for inflation. The dataset covers 506 school districts. Porter and Zona [1999] identify three firms from Cincinnati that were involved in the collusive scheme. Bids submitted by these firms are flagged using a binary indicator identifying cartel participation.

Second, our test is based on cost shocks; therefore, we augment the procurement and cartel data with the average federal marketing order price for raw milk, as reported in Porter and Zona [1999]. The price is quoted for July of each year. The dataset contains 3,754 contracts. The median number of bidders per contract is 2, the mean is 1.9, the minimum is 1, and the maximum is 6.

4.3 Validation Findings

To validate the theory, we examine whether firm entry decisions in response to price shocks differ conditional on cartel activity. Specifically, we regress the number of bidders on the raw milk price, and on an indicator for the school district having high cartel activity (i.e., share of cartel bidders is above median), and the interaction between the two variables. Our theory requires firms to perceive cost shocks as persistent for these shocks to affect entry decisions. A transitory shock raises current costs, but does not materially affect the long-run profitability of collusion. By contrast, a persistent increase in costs reduces the future gains from collusion and therefore weakens incentives to sustain the cartel.

The U.S. farm milk price exhibits this type of persistence. Nicholson and Stephenson [2015] document large-amplitude cycles in the 1996–2014 milk price series, reflecting substantial serial correlation rather than white noise. This allows us to test whether cost increases are associated with greater entry in districts with a larger cartel presence, as predicted by the model. In our baseline test, we estimate the following regression:

$$\begin{aligned} \mathit{NumBidders}_a = \; & \beta_0 + \beta_1 \, \mathit{PriceRawMilk}_{t(a)} + \beta_2 \, \mathit{AboveMedianCartelShare}_{d(a)} \\ & + \beta_3 \, \mathit{PriceRawMilk}_{t(a)} \times \mathit{AboveMedianCartelShare}_{d(a)} \\ & + \gamma \, t(a) + \epsilon_a, \end{aligned} \tag{1}$$

where $\mathit{NumBidders}_a$ is the number of bidders in auction a, $t(a)$ and $d(a)$ denote the year and the school district of auction a, and $\mathit{PriceRawMilk}_t$ is the July federal marketing order price of raw milk (FMO #33, Class 1 Fluid Milk, in dollars per half-pint) in year t. To measure cartel exposure, we first compute the cartel share of district d, defined as the share of cartel firms among all firms that ever bid in the district during the studied period. This share captures how much of a district's potential competition the cartel controls. $\mathit{AboveMedianCartelShare}_d$ is an indicator equal to one if the cartel share exceeds its median among districts with any cartel presence (0.333), and zero otherwise. The design compares how the number of bidders responds to the common milk price across districts with high and low cartel exposure. Under the competitive benchmark $\beta_1 \leq 0$, and collusion implies that $\beta_3 > 0$. A stricter version of the test requires $\beta_1 + \beta_3 > 0$.

In Table 3, we report the estimation results. All columns are estimated by OLS. Column (1) includes only the variables of interest. Column (2) adds a linear time trend, which the remaining columns without year fixed effects also include. Column (3) adds school district fixed effects, column (4) instead adds fixed effects for the winning firm, and column (5) includes both. Columns (6)–(7) replace the linear trend with year fixed effects, which absorb the raw milk price, so that only the interaction term is identified.

Table 3: Bidder participation and the raw milk price, by cartel exposure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Raw milk price β_1	-3.46	-7.70**	-6.19	-6.06	-7.39*		
	(3.88)	(3.76)	(3.83)	(3.82)	(3.90)		
Price × Above-median share β_3	27.69*	26.94*	31.85*	24.14	31.97*	31.48*	32.02*
	(15.28)	(15.10)	(16.63)	(15.34)	(16.68)	(16.43)	(16.52)
Cartel-regime slope $\beta_1 + \beta_3$	24.23	19.25	25.66	18.08	24.58		
	(14.77)	(14.52)	(16.18)	(14.83)	(16.26)		
Mean of dependent variable	1.91	1.91	1.91	1.91	1.91	1.91	1.91
Linear trend	No	Yes	Yes	Yes	Yes		
School district FE	No	No	Yes	No	Yes	Yes	Yes
Winning bidder FE	No	No	No	Yes	Yes	No	Yes
Year FE	No	No	No	No	No	Yes	Yes
R^2	0.031	0.035	0.525	0.147	0.554	0.531	0.559
Observations	3,754	3,754	3,754	3,754	3,754	3,754	3,754

Dependent variable: number of bidders per contract. All columns are estimated by OLS. Standard errors, clustered at school district level, are reported in parentheses. Significance at 10% (*), 5% (**), and 1% (***), based on district-clustered p-values. The raw milk price is the July federal marketing order price (FMO #33, Class 1 Fluid Milk) in dollars per half-pint, and it takes one value per year. Above-median share equals one if the cartel share of a district's potential bidders exceeds the median among districts with any cartel presence (0.333), and zero otherwise. The cartel-regime slope is the price slope in above-median districts: $\beta_1 \leq 0$ is the competitive prediction and $\beta_1 + \beta_3 > 0$ is the collusive reversal. Column (1) includes no time controls. Columns (2)–(5) include a linear time trend. Columns (6)–(7) replace the linear trend in columns (3) and (5) with year fixed effects, which absorb the trend and the raw milk price, so only β_3 is identified in these columns.

The coefficient on the raw milk price is negative in all five columns and statistically significant in two of them. The interaction coefficient is positive in all seven columns and statistically significant in six of them. This means that conditional on cartel activity, higher milk prices are associated with more bidders in auctions with stronger cartel presence. This pattern is consistent with the model's prediction that persistent cost increases weaken incentives to sustain collusion and induce cartel members to participate more aggressively.

The historical record provides further independent support for this mechanism. The two cartel breakdowns reported in the trial testimony, the 1983–1984 and 1989–1990 school years [Porter and Zona, 1999, p. 269], coincide with periods of elevated or sharply rising raw milk prices in the federal marketing order series, and the year-by-year esti-

mates of Porter and Zona [1999] show the collusive markup collapsing in precisely those years. These are only two episodes, and other factors may have played a role. Nonetheless, the timing is consistent with our proposed mechanism.

Caveats. Three caveats apply. First, the interaction term is not identified from exogenous variation in both of its components. The raw milk price is a single aggregate series. Although it is plausibly exogenous to any individual auction, which rules out reverse causality from local entry or collusion, it may move with other year-level factors, such as regional weather shocks or changes in federal dairy policy, that independently affect participation. This type of variation should be captured by year fixed effects. The cartel share, in turn, is not randomly assigned: districts served by the defendants differ systematically from others, for instance in their distance to the cartel's plants, and the share is constructed from realized bidding as the fraction of cartel firms among all firms that ever bid in the district, so it is related to entry by construction. School district fixed effects absorb differences in levels, and year fixed effects absorb shocks common to all districts, but neither rules out that districts with a high cartel share respond differently to the milk price or to other year-specific shocks for reasons unrelated to collusion. We therefore interpret the interaction as descriptive validation of the predicted sign pattern rather than as a causal effect of cartel presence.

Second, our sample spans only eleven years (1980–1990), so the milk price takes only eleven distinct annual values. Although the number of contracts is large (3,754), the time-series variation used for identification is limited. As a result, statistical power is constrained and the estimates may be sensitive to individual years.

Third, the cartel also employed complementary bidding, so bid submission is an imperfect measure of genuine competitive entry. Indeed, Porter and Zona [1999] find that the defendants bid more often than a competitive benchmark predicts within thirty miles of their plants, and that their bid-submission residuals are positively correlated pairwise—patterns suggesting that complementary bids inflated observed participation in cartel districts. Our interaction term is identified from how participation co-moves with persistent cost shocks rather than from its level, and the level contamination is absorbed by the cartel-intensity main effect and the school district fixed effects. Moreover, there is no evident mechanism by which the number of complementary bids would rise with the milk price specifically in cartel auctions; if anything, weakening discipline may convert complementary bids into serious ones without changing bid counts, biasing the test toward a false negative rather than a false positive, consistent with the asymmetric error structure discussed in Section 3.2.

5 Conclusion

This paper develops the first collusion screen based on the comparative statics of a repeated entry game. Our starting point is that entry decisions are often observable in settings where prices, quantities, or bids are unavailable, uninformative, or centrally set—franchising chains with national pricing being one example. We show that competitive and collusive conduct generate opposite predictions about how joint entry responds to persistent profit shocks. Under competition, an adverse cost shock unambiguously reduces joint entry. Under collusion sustained by dynamic incentives, the same shock erodes the value of future punishment. When this enforcement effect dominates the reduction in the one-shot gain from over-entry, the cartel is forced to relax exclusion, and joint entry *rises*. This mechanism is the entry-margin analogue of the cyclical-collusion logic of Rotemberg and Saloner [1986] and Haltiwanger and Harrington [1991]. Appendix B shows that it can also survive privately observed entry costs under particular verification technologies.

The resulting test is deliberately asymmetric regarding type I and type II errors. Because no parametrization of the competitive model can generate a positive relationship between persistent cost increases and joint entry, a positive estimated relationship constitutes evidence against competition. False positives are ruled out by the theory, while false negatives remain possible when the enforcement channel is weak or the application has limited statistical power. We view this asymmetry as attractive for screening purposes, since the cost of falsely flagging a competitive market is typically higher for a competition authority than the cost of failing to flag a cartel that other tools might catch. False positives are nonetheless in practice possible, for example, due to sampling bias or misspecified empirical model.

Monte Carlo evidence shows that a simple probit of joint entry on marginal cost recovers correctly signed and statistically significant coefficients at sample sizes of a few hundred entry events—well within the range of typical procurement or market entry datasets.

Applying the test to the Ohio school milk cartel of Porter and Zona [1999], in which market allocation based on respecting incumbencies made entry restraint an integral part of the collusive scheme, we find that bidder participation responds positively to persistent increases in raw milk prices precisely where cartel activity is concentrated, while the relationship is negative elsewhere. The validation is subject to caveats—a short time window, and a cartel that also employed complementary bidding—but the sign pattern matches the theory's distinguishing prediction.

Several extensions merit future work. The framework extends naturally to more than

two potential entrants, where the outcome of interest becomes the number of entrants; to differences-in-differences designs exploiting common cost shocks across markets; and to other forward-looking strategic margins such as investment or capacity, where asymmetric responses to shocks of different persistence could similarly separate conduct. More broadly, our results suggest that the extensive margin, long treated in the entry literature as a window onto market structure, can also serve as a window onto conduct.

References

Dilip Abreu. On the theory of infinitely repeated games with discounting. *Econometrica*, 56(2):383–396, 1988. doi: 10.2307/1911077.

Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063, September 1990. doi: 10.2307/2938299.

Susan Athey and Kyle Bagwell. Optimal collusion with private information. *The RAND Journal of Economics*, 32(3):428–465, 2001. doi: 10.2307/2696363.

Susan Athey and Kyle Bagwell. Collusion with persistent cost shocks. *Econometrica*, 76 (3):493–540, 2008. doi: 10.1111/j.1468-0262.2008.00845.x.

Susan Athey, Kyle Bagwell, and Chris Sanchirico. Collusion and price rigidity. *Review of Economic Studies*, 71(2):317–349, 2004. doi: 10.1111/0034-6527.00286.

Kyle Bagwell and Robert W. Staiger. Collusion over the business cycle. *RAND Journal of Economics*, 28(1):82–106, 1997. doi: 10.2307/2555941.

Patrick Bajari and Lixin Ye. Deciding between competition and collusion. *The Review of Economics and Statistics*, 85(4):971–989, 2003. doi: 10.1162/003465303772815871.

Laura H. Baldwin, Robert C. Marshall, and Jean-François Richard. Bidder collusion at Forest Service timber sales. *Journal of Political Economy*, 105(4):657–699, 1997. doi: 10.1086/262089.

Abhijit Banerji and J. V. Meenakshi. Millers, commission agents and collusion in grain markets: Evidence from basmati auctions in North India. *The B.E. Journal of Economic Analysis & Policy*, 8(1), 2008. doi: 10.2202/1935-1682.1786. (Contributions), Article 4.

B. Douglas Bernheim and Michael D. Whinston. Multimarket contact and collusive behavior. *RAND Journal of Economics*, 21(1):1–26, 1990. URL `https://www.jstor.org/stable/2555490`.

Steven T. Berry. Estimation of a model of entry in the airline industry. *Econometrica*, 60 (4):889–917, 1992. doi: 10.2307/2951571.

Timothy F. Bresnahan and Peter C. Reiss. Entry in monopoly markets. *Review of Economic*

Studies, 57(4):531–553, 1990. doi: 10.2307/2298085.

Timothy F. Bresnahan and Peter C. Reiss. Entry and competition in concentrated markets. *Journal of Political Economy*, 99(5):977–1009, 1991. doi: 10.1086/261786.

Martin C. Byford and Joshua S. Gans. Collusion at the extensive margin. *International Journal of Industrial Organization*, 37:75–83, 2014. doi: 10.1016/j.ijindorg.2014.08.001.

Sylvain Chassang, Kei Kawai, Jun Nakabayashi, and Juan Ortner. Robust screens for noncompetitive bidding in procurement auctions. *Econometrica*, 90(1):315–346, 2022. doi: 10.3982/ECTA17155.

Olivier Compte. Communication in repeated games with imperfect private monitoring. *Econometrica*, 66(3):597–626, May 1998. doi: 10.2307/2998576.

Timothy G. Conley and Francesco Decarolis. Detecting bidders groups in collusive auctions. *American Economic Journal: Microeconomics*, 8(2):1–38, 2016. doi: 10.1257/mic.20130254.

Glenn Ellison. Theories of cartel stability and the Joint Executive Committee. *The RAND Journal of Economics*, 25(1):37–57, 1994.

Natalia Fabra. Collusion with capacity constraints over the business cycle. *International Journal of Industrial Organization*, 24(1):69–81, 2006. doi: 10.1016/j.ijindorg.2005.01.014.

James W. Friedman. A non-cooperative equilibrium for supergames. *Review of Economic Studies*, 38(1):1–12, 1971. doi: 10.2307/2296617.

Drew Fudenberg, David K. Levine, and Eric Maskin. The folk theorem with imperfect public information. *Econometrica*, 62(5):997–1039, September 1994. doi: 10.2307/2951505.

Edward J. Green and Robert H. Porter. Noncooperative collusion under imperfect price information. *Econometrica*, 52(1):87–100, January 1984. doi: 10.2307/1911462.

John Haltiwanger and Joseph E. Harrington. The impact of cyclical demand movements on collusive behavior. *RAND Journal of Economics*, 22(1):89–106, 1991. doi: 10.2307/2601009.

Joseph E. Harrington. Detecting cartels. In Paolo Buccirossi, editor, *Handbook of Antitrust Economics*, chapter 6, pages 213–258. MIT Press, Cambridge, MA, 2008.

Joseph E. Harrington and Andrzej Skrzypacz. Private monitoring and communication in cartels: Explaining recent collusive practices. *American Economic Review*, 101(6):2425–2449, 2011. doi: 10.1257/aer.101.6.2425.

Panle Jia. What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica*, 76(6):1263–1316, 2008. doi: 10.3982/ECTA6649.

Michihiro Kandori. Correlated demand shocks and price wars during booms. *Review of Economic Studies*, 58(1):171–180, 1991. doi: 10.2307/2298053.

Michihiro Kandori and Hitoshi Matsushima. Private observation, communication and collusion. *Econometrica*, 66(3):627–652, May 1998. doi: 10.2307/2998577.

Kei Kawai, Jun Nakabayashi, Juan Ortner, and Sylvain Chassang. Using bid rotation and incumbency to detect collusion: A regression discontinuity approach. *Review of Economic Studies*, 90(1):376–403, 2023. doi: 10.1093/restud/rdac013.

Charles F Nicholson and Mark W Stephenson. Milk price cycles in the U.S. dairy supply chain and their management implications. *Agribusiness*, 31(4):507–520, 2015.

Juan Ortner, Sylvain Chassang, and Kei Kawai. Screening adaptive cartels. *American Economic Review: Insights*, 2026. doi: 10.1257/aeri.20250233. Forthcoming.

Martin Pesendorfer. A study of collusion in first-price auctions. *Review of Economic Studies*, 67(3):381–411, 2000.

Louis Phlips. On the detection of collusion and predation. *European Economic Review*, 40 (3–5):495–510, 1996. doi: 10.1016/0014-2921(95)00064-X.

Robert H. Porter. A study of cartel stability: The Joint Executive Committee, 1880–1886. *The Bell Journal of Economics*, 14(2):301–314, 1983.

Robert H. Porter. Detecting collusion. *Review of Industrial Organization*, 26(2):147–167, 2005. doi: 10.1007/s11151-004-7294-7.

Robert H. Porter and J. Douglas Zona. Detection of bid rigging in procurement auctions. *Journal of Political Economy*, 101(3):518–538, 1993. doi: 10.1086/261885.

Robert H Porter and J Douglas Zona. Ohio School Milk Markets: An Analysis of Bidding. *RAND Journal of Economics*, 30(2):263–288, 1999.

Michael K. Price. Using the spatial distribution of bidders to detect collusion in the marketplace: Evidence from timber auctions. *Journal of Regional Science*, 48(2):399–417, 2008. doi: 10.1111/j.1467-9787.2008.00557.x.

Julio J. Rotemberg and Garth Saloner. A supergame-theoretic model of price wars during booms. *American Economic Review*, 76(3):390–407, 1986. URL `https://www.jstor.org/stable/1813358`.

Katja Seim. An empirical model of firm entry with endogenous product-type choices. *RAND Journal of Economics*, 37(3):619–640, 2006. doi: 10.1111/j.1756-2171.2006.tb00034.x.

Robert W. Staiger and Frank A. Wolak. Collusive pricing with capacity constraints in the presence of demand uncertainty. *RAND Journal of Economics*, 23(2):203–220, 1992. doi: 10.2307/2555984.

George J. Stigler. A theory of oligopoly. *Journal of Political Economy*, 72(1):44–61, 1964. doi: 10.1086/258853.

Elie Tamer. Incomplete simultaneous discrete response model with multiple equilibria. *Review of Economic Studies*, 70(1):147–165, 2003. doi: 10.1111/1467-937X.00240.

Hélder Vasconcelos. Entry effects on cartel stability and the Joint Executive Committee. *Review of Industrial Organization*, 24(3):219–241, 2004. doi: 10.1023/B:REIO.0000038247.28065.33.

Johannes Wachs and János Kertész. A network approach to cartel detection in public auction markets. *Scientific Reports*, 9(1):10818, 2019.

Yuichi Yamamoto. Characterizing belief-free review-strategy equilibrium payoffs under conditional independence. *Journal of Economic Theory*, 147(5):1998–2027, September 2012. doi: 10.1016/j.jet.2012.05.016.

A Appendix: Theory – Complete Information

This appendix formalizes the theoretical arguments used in the main text. We allow monopoly profits to co-move with duopoly profits by writing $\pi_M = \pi_M(\pi_D)$, where $\pi_M(\cdot)$ is differentiable and increasing. This is natural, for example, when the underlying shock affects marginal costs and therefore shifts both monopoly and duopoly profits. Throughout this appendix, whenever derivatives with respect to the cutoff are used, F is assumed to be continuously differentiable with density f.

Let $\Pi^N(\pi_D)$ and $\Pi^C(k;\pi_D)$ denote a representative firm's ex-ante expected per-period profit net of entry costs under competition and under the collusive cutoff rule with accommodation cutoff k, respectively. With the tie-breaking refinement that the lower-cost firm enters whenever unilateral entry is feasible, these payoffs can be written as

$$\begin{aligned}\Pi^N(\pi_D) = &\int_0^{\pi_D} \Big[F(\pi_D)(\pi_D - x) + \big(1 - F(\pi_D)\big)\big(\pi_M(\pi_D) - x\big)\Big]\, dF(x) \\ &+ \int_{\pi_D}^{\pi_M(\pi_D)} \big(\pi_M(\pi_D) - x\big)\big(1 - F(x)\big)\, dF(x),\end{aligned} \tag{2}$$

$$\begin{aligned}\Pi^C(k;\pi_D) = &\int_0^{k} \Big[F(k)(\pi_D - x) + \big(1 - F(k)\big)\big(\pi_M(\pi_D) - x\big)\Big]\, dF(x) \\ &+ \int_{k}^{\pi_M(\pi_D)} \big(\pi_M(\pi_D) - x\big)\big(1 - F(x)\big)\, dF(x).\end{aligned} \tag{3}$$

Define the per-period collusive advantage by

$$D(k,\pi_D) := \Pi^C(k;\pi_D) - \Pi^N(\pi_D),$$

and let the discounted loss from triggering Z periods of competitive reversion be

$$\Gamma(k,\pi_D) := H_0 D(k,\pi_D), \qquad H_0 := \sum_{t=1}^{Z} \delta^t = \delta\frac{1-\delta^Z}{1-\delta}.$$

Sustainability and selection. For a candidate cutoff $k \in [0,\pi_D]$, define

$$\Psi(k,\pi_D) := \pi_D - k - \Gamma(k,\pi_D).$$

We call a cutoff *sustainable* if the associated collusive entry rule satisfies the exclusion incentive constraint under the specified punishment strategy. Because the gain from over-entry is decreasing in the excluded firm's entry cost, the cutoff k is sustainable if and only if

$$\Psi(k,\pi_D) \le 0.$$

Accordingly, define the set of sustainable cutoffs by

$$\mathcal{S}(\pi_D) := \big\{k \in [0,\pi_D] : \Psi(k,\pi_D) \le 0\big\}. \tag{4}$$

The set $\mathcal{S}(\pi_D)$ is nonempty because $k = \pi_D$ is always sustainable: at this cutoff the collusive allocation coincides with the selected competitive allocation, so $D(\pi_D,\pi_D) = 0$ and hence

$\Psi(\pi_D, \pi_D) = 0$. Continuity of Ψ implies that $\mathcal{S}(\pi_D)$ is closed, and therefore its minimum exists. As in (C), the cartel-selected accommodation cutoff is

$$k^C(\pi_D) = \min \mathcal{S}(\pi_D).$$

If $k^C(\pi_D) > 0$, then the incentive constraint binds:

$$\Psi\big(k^C(\pi_D), \pi_D\big) = \pi_D - k^C(\pi_D) - \Gamma\big(k^C(\pi_D), \pi_D\big) = 0. \tag{5}$$

Indeed, if the inequality were strict, continuity would imply that a slightly smaller cutoff is also sustainable, contradicting the definition of $k^C(\pi_D)$. At the boundary $k^C(\pi_D) = 0$, the incentive constraint may hold strictly.

Lemma 1. *For a fixed profit vector (π_M, π_D), expected collusive profit is weakly decreasing in the accommodation cutoff k. It is strictly decreasing whenever increasing the cutoff changes the prescribed entry outcome with positive probability.*

Proof. Consider an increase in the accommodation cutoff. The only states in which the prescribed allocation changes are states that previously featured unilateral entry by the lower-cost firm but, under the higher cutoff, feature joint entry. Let $k_{(1)}$ and $k_{(2)}$ denote the lower and higher entry costs in such a state.

Under unilateral entry, total industry profit is

$$\pi_M - k_{(1)}.$$

Under joint entry, total industry profit is

$$2\pi_D - k_{(1)} - k_{(2)}.$$

The change in total industry profit is therefore

$$2\pi_D - \pi_M - k_{(2)} < 0,$$

because $2\pi_D < \pi_M$. Thus, increasing the cutoff weakly reduces expected industry profit and strictly reduces it whenever the affected set has positive probability. By symmetry, each firm's ex-ante expected profit is one half of expected industry profit.

Equivalently, when F has density f, differentiation of (3) gives

$$\frac{\partial \Pi^C(k; \pi_D)}{\partial k} = F(k)f(k)\big(2\pi_D - \pi_M(\pi_D) - k\big) \leq 0, \tag{6}$$

with strict inequality whenever $F(k)f(k) > 0$. □

Lemma 1 implies that, among sustainable cutoff rules, both firms prefer the smallest cutoff. Hence the selection in (C) is cartel-optimal within the class of cutoff rules considered here. Since

$$\Pi^C(\pi_D; \pi_D) = \Pi^N(\pi_D),$$

Lemma 1 implies

$$D(k, \pi_D) \geq 0 \qquad \text{for every } k \in [0, \pi_D].$$

Proposition 2. *Fix a duopoly-profit level π_D^0, and let*

$$k_0 := k^C(\pi_D^0) \in (0, \pi_D^0)$$

be the cartel-selected cutoff. Suppose that Γ is continuously differentiable in a neighborhood of (k_0, π_D^0), that F is strictly increasing at π_D^0 and k_0, and that

$$1 + \Gamma_k(k_0, \pi_D^0) > 0. \tag{7}$$

Then, in a neighborhood of π_D^0, the cartel-selected cutoff is continuously differentiable and satisfies

$$\frac{dk^C}{d\pi_D} = \frac{1 - \Gamma_{\pi_D}\big(k^C(\pi_D), \pi_D\big)}{1 + \Gamma_k\big(k^C(\pi_D), \pi_D\big)}. \tag{8}$$

Consequently, a permanent decrease in duopoly profits raises joint entry under collusion if and only if

$$\Gamma_{\pi_D}(k_0, \pi_D^0) > 1. \tag{9}$$

Under competition, the same decrease lowers joint entry.

Proof. Because $k_0 > 0$ is the selected cutoff, (5) implies

$$\Psi(k_0, \pi_D^0) = 0.$$

Moreover, (7) gives

$$\Psi_k(k_0, \pi_D^0) = -1 - \Gamma_k(k_0, \pi_D^0) < 0.$$

The implicit-function theorem therefore yields a unique continuously differentiable local root $\widehat{k}(\pi_D)$ satisfying

$$\Psi\big(\widehat{k}(\pi_D), \pi_D\big) = 0$$

and $\widehat{k}(\pi_D^0) = k_0$.

It remains to verify that this local root is still the smallest sustainable cutoff. By continuity of Ψ_k, choose $\varepsilon > 0$ and a neighborhood of π_D^0 such that

$$\Psi_k(k, \pi_D) < 0$$

whenever $k \in [k_0 - \varepsilon, k_0 + \varepsilon]$ and π_D lies in that neighborhood. Shrinking the neighborhood if necessary, the local root satisfies

$$\widehat{k}(\pi_D) \in (k_0 - \varepsilon/2, k_0 + \varepsilon/2).$$

Since k_0 is the smallest sustainable cutoff at π_D^0,

$$\Psi(k, \pi_D^0) > 0 \qquad \text{for every } k < k_0.$$

In particular, compactness of $[0, k_0 - \varepsilon]$ implies

$$\min_{k \in [0, k_0 - \varepsilon]} \Psi(k, \pi_D^0) > 0.$$

By continuity, this inequality remains strict for all π_D sufficiently close to π_D^0. For $k \in [k_0 - \varepsilon, \widehat{k}(\pi_D))$, strict monotonicity of $\Psi(\cdot, \pi_D)$ on the local neighborhood implies

$$\Psi(k, \pi_D) > \Psi\big(\widehat{k}(\pi_D), \pi_D\big) = 0.$$

Thus no cutoff below $\widehat{k}(\pi_D)$ is sustainable, while $\widehat{k}(\pi_D)$ itself is sustainable. Hence

$$k^C(\pi_D) = \widehat{k}(\pi_D)$$

locally.

Implicit differentiation of

$$\pi_D - k^C(\pi_D) - \Gamma\big(k^C(\pi_D), \pi_D\big) = 0$$

yields (8). Because its denominator is positive,

$$\frac{dk^C}{d\pi_D} < 0 \quad \Longleftrightarrow \quad \Gamma_{\pi_D}(k_0, \pi_D^0) > 1.$$

The collusive joint-entry probability is

$$P^C(\pi_D) = F\big(k^C(\pi_D)\big)^2.$$

Since F is strictly increasing at k_0, a decrease in π_D raises joint entry under collusion if and only if (9) holds.

Under competition, the joint-entry probability is

$$P^N(\pi_D) = F(\pi_D)^2,$$

which is strictly increasing in π_D because F is strictly increasing at π_D^0. Hence a decrease in π_D lowers joint entry under competition. □

Next, we provide a sufficient condition for the existence of a punishment length Z and discount factor δ such that a prescribed interior cutoff is the cartel-selected sustainable cutoff and has the opposite comparative static from the competitive cutoff.

Proposition 3. *Fix a duopoly-profit level π_D and suppose that D is continuously differentiable. Suppose there exists $k^C \in (0, \pi_D)$ such that*

$$D(k^C, \pi_D) > 0$$

and

$$D_{\pi_D}(k^C, \pi_D) > \frac{D(k^C, \pi_D)}{\pi_D - k^C} > \sup_{x \in [0, k^C]} \big\{-D_k(x, \pi_D)\big\}. \tag{E}$$

Here D_{π_D} is the derivative with respect to π_D holding the cutoff fixed, including the effect operating through $\pi_M(\pi_D)$. Then there exist $Z \in \mathbb{N}_+$ and $\delta \in (0, 1)$ such that k^C is the cartel-selected smallest sustainable cutoff, lies on a regular branch, and the locally selected cutoff is strictly decreasing in π_D.

Proof. Choose

$$H_0 = \frac{\pi_D - k^C}{D(k^C, \pi_D)}. \tag{10}$$

Because $k^C < \pi_D$ and $D(k^C, \pi_D) > 0$, we have $H_0 > 0$. This choice gives

$$\Psi(k^C, \pi_D) = \pi_D - k^C - H_0 D(k^C, \pi_D) = 0,$$

so the candidate cutoff satisfies the binding exclusion constraint.

The second inequality in (E) implies that, for every $x \in [0, k^C]$,

$$-H_0 D_k(x, \pi_D) < 1.$$

Therefore,

$$1 + H_0 D_k(x, \pi_D) > 0$$

and hence

$$\Psi_k(x, \pi_D) = -1 - H_0 D_k(x, \pi_D) < 0 \qquad \text{for every } x \in [0, k^C].$$

Since $\Psi(k^C, \pi_D) = 0$, it follows that

$$\Psi(x, \pi_D) > 0 \qquad \text{for every } x < k^C.$$

Thus no smaller cutoff is sustainable, and k^C is the selected cutoff defined in (C). In particular,

$$1 + \Gamma_k(k^C, \pi_D) = 1 + H_0 D_k(k^C, \pi_D) > 0,$$

so the selected cutoff lies on a regular branch.

The first inequality in (E), together with (10), gives

$$\Gamma_{\pi_D}(k^C, \pi_D) = H_0 D_{\pi_D}(k^C, \pi_D) > 1.$$

Because

$$1 + \Gamma_k(k^C, \pi_D) = 1 + H_0 D_k(k^C, \pi_D) > 0,$$

the implicit-function theorem applies to the selected root. Moreover,

$$\frac{dk^C}{d\pi_D} = \frac{1 - H_0 D_{\pi_D}(k^C, \pi_D)}{1 + H_0 D_k(k^C, \pi_D)}.$$

The first inequality in (E) implies that the numerator is strictly negative, while the second inequality implies that the denominator is strictly positive. Therefore,

$$\frac{dk^C}{d\pi_D} < 0.$$

It remains to show that the chosen H_0 can be generated by some (Z, δ). Choose an integer $Z > H_0$ and define

$$h_Z(\delta) := \sum_{t=1}^{Z} \delta^t.$$

The function h_Z is continuous and strictly increasing on $(0,1)$, with

$$\lim_{\delta \downarrow 0} h_Z(\delta) = 0 \qquad \text{and} \qquad \lim_{\delta \uparrow 1} h_Z(\delta) = Z > H_0.$$

The intermediate-value theorem therefore implies that there exists $\delta \in (0,1)$ such that

$$\sum_{t=1}^{Z} \delta^t = H_0.$$

Thus the required punishment factor can be implemented by a finite punishment length and a discount factor in $(0,1)$. □

Example. We illustrate the framework under linear demand. Suppose inverse demand is $p = a - q$ and both firms have constant marginal cost c. Monopoly and duopoly profits are

$$\pi_M = \left(\frac{a-c}{2}\right)^2, \qquad \pi_D = \left(\frac{a-c}{3}\right)^2 = \frac{4}{9}\pi_M,$$

respectively.

In the simulations, entry costs are independently and identically distributed according to the uniform distribution on $[0,1]$. Take $a = 11/5$ and $c = 1/5$, so that $\pi_M = 1$ and $\pi_D = 4/9$. Thus monopoly profit is normalized to one. At this parameterization,

$$\Pi^N\left(\frac{4}{9}\right) = \frac{641}{2187}, \qquad \Pi^C\left(k; \frac{4}{9}\right) = \frac{1}{3} - \frac{1}{18}k^2 - \frac{1}{3}k^3.$$

Hence the binding exclusion constraint is

$$\frac{4}{9} - k - H_0\left[-\frac{1}{18}k^2 - \frac{1}{3}k^3 + \frac{88}{2187}\right] = 0. \tag{11}$$

The equation also admits the noncollusive solution $k = \pi_D = 4/9$. Consistent with (C), the cartel uses the smallest sustainable solution.

At $\pi_D = 4/9$, the relevant derivatives are

$$D_{\pi_D}\left(k, \frac{4}{9}\right) = \frac{1}{8}\left(\frac{176}{81} - k^2\right), \qquad -D_k\left(k, \frac{4}{9}\right) = k\left(\frac{1}{9} + k\right).$$

Therefore, the sign-reversal condition and regularity requirement become

$$\Gamma_{\pi_D} = \frac{H_0}{8}\left(\frac{176}{81} - k^2\right) > 1, \qquad -\Gamma_k = H_0 k\left(\frac{1}{9} + k\right) < 1. \tag{12}$$

Moreover,

$$-D_k\left(x, \frac{4}{9}\right) = x\left(\frac{1}{9} + x\right)$$

is increasing in x. Hence the second inequality in (12), verified at the candidate cutoff k, implies

$$H_0 \sup_{x\in[0,k]} \left\{-D_k(x, 4/9)\right\} < 1.$$

Together with the first inequality in (12) and the binding equation (11), this verifies the stronger selection condition (E). Thus the cutoff used below is the smallest sustainable cutoff, not merely one local solution of the binding equation.

For example, choosing $\delta = 0.99$ and $Z = 10$ gives $k^C \approx 0.0671$. Both inequalities in (12) hold, since $\Gamma_{\pi_D} \approx 2.57$ and $-\Gamma_k \approx 0.11$. Alternatively, choosing $k^C = 0.195$, the binding constraint determines $H_0 \approx 7$. One can then choose (δ, Z) such that

$$H_0 = \delta \frac{1-\delta^Z}{1-\delta} \approx 7.$$

For instance, $Z = 9$ and $\delta \approx 0.95$ deliver this value. With these parameters, $\Gamma_{\pi_D} \approx 1.87$ and $-\Gamma_k \approx 0.42$. Hence the candidate is the cartel-selected sustainable cutoff and satisfies

$$\frac{dk^C}{d\pi_D} < 0.$$

This is the baseline parameterization used in the simulations.

B Appendix: Theory – Incomplete Information

In this section, we discuss the model with incomplete information about entry costs.

Consider the same model as before, except suppose that k_j is private information of firm $j \in \{A, B\}$. Assume that k_A and k_B are independently drawn from the common distribution F. We write duopoly profit π_D as the primitive public profit index and allow monopoly profit to co-move with it, denoting monopoly profit by $\pi_M(\pi_D)$. Thus, an adverse profit shock is a decrease in π_D, which also reduces $\pi_M(\pi_D)$.

Assume: (i) $0 < \pi_D < \pi_M(\pi_D)/2$ for all π_D under consideration, (ii) $\pi_M'(\pi_D) > 0$, (iii) F is continuously differentiable with density f and full support on $[k_L, k_H] \subseteq \mathbb{R}_+$ such that $k_L < \pi_D$ and $k_H > \pi_M(\pi_D)$ for all π_D under consideration.[4]

Competitive Equilibrium. Let us analyze the competitive equilibrium. This entry game has a unique symmetric cutoff Bayesian Nash equilibrium. Fix a candidate cutoff $\bar{k}$ and suppose the opponent plays the strategy "enter iff $k \leq \bar{k}$", otherwise stay out.

If a firm with cost k enters, its expected net payoff is

$$\mathbb{E}[\phi]_{\text{ent}}(k; \bar{k}, \pi_D) = (1 - F(\bar{k}))\, \pi_M(\pi_D) + F(\bar{k})\, \pi_D - k. \tag{13}$$

[4] If $k_H = \infty$, then $f > 0$ for all $k \in [k_L, \infty)$.

Staying out yields 0, so the best response is to enter if and only if $\mathbb{E}[\phi]_{\text{ent}}(k; \bar{k}, \pi_D) \geq 0$, which is equivalent to $k \leq k^*(\bar{k}, \pi_D)$ where

$$k^*(\bar{k}, \pi_D) \equiv (1 - F(\bar{k}))\, \pi_M(\pi_D) + F(\bar{k})\, \pi_D.$$

We focus on symmetric cutoff equilibria. A symmetric cutoff equilibrium is a fixed point $\bar{k}$ such that $k^*(\bar{k}, \pi_D) = \bar{k}$, i.e.

$$\bar{k} = (1 - F(\bar{k}))\pi_M(\pi_D) + F(\bar{k})\pi_D,$$

and each firm enters iff $k_j \leq \bar{k}$.

Define

$$G(k, \pi_D) \equiv k - (1 - F(k))\pi_M(\pi_D) - F(k)\pi_D.$$

The equilibrium cutoff $\bar{k}(\pi_D)$ satisfies $G(\bar{k}(\pi_D), \pi_D) = 0$.

Proposition 4. *For each duopoly profit level π_D, there is a unique symmetric competitive entry cutoff $\bar{k}(\pi_D)$. This cutoff is increasing in duopoly profit: higher duopoly profits make firms willing to enter at higher entry costs. Equivalently, an adverse profit shock lowers the competitive entry cutoff.*

Proof. By the assumption that F is continuously differentiable, $G(k, \pi_D)$ is continuous in k. Moreover,

$$G_k(k, \pi_D) = 1 + f(k)\big(\pi_M(\pi_D) - \pi_D\big) > 0,$$

since $\pi_M(\pi_D) > \pi_D$. Hence G is strictly increasing in k.

Furthermore, $G(k_L, \pi_D) < 0$ and $G(k_H, \pi_D) > 0$ by the support assumptions. Therefore, there exists a unique $\bar{k}(\pi_D)$ such that $G(\bar{k}(\pi_D), \pi_D) = 0$.

Implicit differentiation yields

$$G_k(\bar{k}, \pi_D)\, \bar{k}'(\pi_D) + G_{\pi_D}(\bar{k}, \pi_D) = 0,$$

and hence

$$\bar{k}'(\pi_D) = -\frac{G_{\pi_D}(\bar{k}, \pi_D)}{G_k(\bar{k}, \pi_D)}.$$

Now,

$$G_{\pi_D}(k, \pi_D) = -(1 - F(k))\pi_M'(\pi_D) - F(k) < 0,$$

because $\pi_M'(\pi_D) > 0$. Therefore,

$$\bar{k}'(\pi_D) = \frac{(1 - F(\bar{k}))\pi_M'(\pi_D) + F(\bar{k})}{1 + f(\bar{k})\big(\pi_M(\pi_D) - \pi_D\big)} > 0.$$

□

This result implies that, in the incomplete-information competitive benchmark, an adverse profit shock cannot raise joint entry. Indeed, the probability of joint entry is

$$P^N(\pi_D) = F(\bar{k}(\pi_D))^2,$$

which is increasing in π_D. Hence, when π_D decreases, the competitive cutoff falls and joint entry weakly decreases.

Collusive Equilibrium. Let us now discuss collusive equilibrium. We have shown in the complete-information model that repeated interaction can generate the opposite comparative static: a decrease in duopoly profits can weaken punishments sufficiently strongly that the cartel is forced to tolerate more joint entry. Therefore, if empirically adverse profit shocks and entry move up together, one possible interpretation is that firms are playing a collusive equilibrium of this type.

With complete information, "collusion in entry" can be enforced by prescribing efficient participation and punishing observable deviations. With private costs, however, punishments can depend only on public histories of entry outcomes, so "unexpected" entry profiles may reflect either deviations or cost realizations. This enforcement problem is the entry analog of the imperfect public-monitoring logic in Green and Porter [1984], where a bad public outcome can arise either from cheating or from an adverse shock, and of the general imperfect public-monitoring frameworks developed by Abreu et al. [1990] and Fudenberg et al. [1994]. In particular, when a cartel follows Green and Porter [1984]'s construction and triggers punishment after an "undesirable" public signal, here an entry profile, then punishment necessarily occurs with positive probability even under compliant play. These on-path punishments reduce continuation values and can weaken incentives for collusion, as noted by Abreu et al. [1990].

To extend our model to incorporate private entry costs, there are several modeling directions. First, following the constructions of Yamamoto [2012], firms can cooperate for L periods, apply a statistical test to the realized public history of entry outcomes, and then switch to a Z-period punishment phase only if the test fails. Second, in the spirit of Athey and Bagwell [2001], collusion can be self-enforcing with private information when incentive compatibility is sustained via promised continuation payoffs, e.g., future market-share favors for firms that abstain from competing aggressively in high-cost states. Third, if firms can communicate their privately observed costs before choosing entry, then sustaining efficient entry allocations hinges on incentive compatibility constraints for that communication, as emphasized in the repeated-games-with-communication literature; see Compte [1998] and Kandori and Matsushima [1998].

For detection, the key question is whether the comparative statics in the main analysis continue to hold when entry costs are private information. Under competition, an adverse profit shock weakens entry incentives and therefore reduces everybody's entry. Under collusion, the same shock can instead weaken collusion-enforcing punishments, which may push the cartel to allow more joint entry. However, with private entry costs the punishment is constrained by imperfect inference and, if communication is allowed, by incentive compatibility for reports. Thus, robustness cannot be taken for granted and should be established.

To extend our results to incomplete information, it is enough to show that the same comparative static result can arise in a repeated-game equilibrium with private information. We do not aim to characterize collusion under incomplete information in general. Instead, we provide simple constructions under specific verification technologies. With private entry costs, one way to make deviations observable is to allow for verifiable disclosure.

Costly third-party verification. Suppose that each firm $j \in \{A, B\}$ privately observes k_j and, before choosing entry $y_j \in \{0, 1\}$, the firms simultaneously choose whether to verify, $v_j \in \{0, 1\}$. If $v_A = v_B = 1$, each firm pays a fee $m > 0$ and a third party reveals (k_A, k_B) to both firms. Otherwise, no disclosure takes place. If either firm refuses verification, the firms play the competitive Bayesian cutoff equilibrium in the current period and immediately enter the Z-period punishment phase.

Consider the following stationary strategy profile. On path, both firms verify in every period. After observing (k_A, k_B), they follow the complete-information collusive entry rule with accommodation cutoff $k_C(\pi_D)$: if $k_A, k_B \leq k_C(\pi_D)$, both enter; otherwise, the lower-cost firm enters whenever unilateral entry is profitable and the other stays out. Any publicly observed deviation—either a refusal to verify or an entry decision inconsistent with the disclosed costs—triggers Z periods of the competitive Bayesian cutoff equilibrium, after which play returns to the collusive path.

Besides the entry incentive constraint below, firms must also be willing to participate in verification. We therefore require that the largest current gain from refusing verification, including the saved fee and any change in the current entry payoff, is smaller than the resulting loss in continuation value. This condition is satisfied for sufficiently small m whenever collusion yields a strictly higher continuation value and the punishment is sufficiently severe.

Let $\Pi^C(k_C; \pi_D)$ denote the representative firm's ex-ante per-period payoff under the disclosed-cost collusive rule before verification costs, and let $\Pi^N(\pi_D)$ denote the representative firm's ex-ante per-period payoff in the competitive Bayesian cutoff equilibrium. The discounted punishment loss is

$$\Gamma^m(k_C, \pi_D) \equiv H_0\Big(\Pi^C(k_C; \pi_D) - m - \Pi^N(\pi_D)\Big), \qquad H_0 \equiv \sum_{t=1}^{Z} \delta^t = \delta\frac{1-\delta^Z}{1-\delta}.$$

As in the complete-information model, the tight incentive constraint is for the marginal excluded type and pins down $k_C(\pi_D)$ by

$$\pi_D - k_C(\pi_D) - \Gamma^m(k_C(\pi_D), \pi_D) = 0.$$

Since m enters the continuation-value gap additively and does not vary with either k_C or π_D, this verification model has the same local comparative-statics logic as the complete-information model. If the continuation-value loss falls sufficiently fast after a decrease in π_D, relative to the direct decrease in the one-shot deviation payoff, then the implicit function defined by the marginal excluded type's incentive constraint satisfies

$$\frac{dk_C(\pi_D)}{d\pi_D} < 0.$$

Hence, a decrease in π_D raises the collusive cutoff and increases joint entry.

Ex-post verification. Suppose firms cannot communicate costs prior to entry. In each period t, each firm $i \in \{A, B\}$ privately observes $k_{i,t}$ and then simultaneously chooses entry $y_{i,t} \in \{0, 1\}$. The entry profile is publicly observed.

After the entry decisions, an audit occurs exogenously with probability $\alpha \in (0, 1]$. The audit is compulsory, automatically charges each firm a fee $m > 0$, and publicly reveals the current period's entry costs $(k_{A,t}, k_{B,t})$. If the audit reveals that a firm violated the prescribed entry rule, a Z-period reversion to the competitive Bayesian cutoff equilibrium begins in period $t + 1$. During the punishment phase, no auditing is used.

Fix a cutoff $k_C(\pi_D)$ and consider the implementable symmetric collusive entry rule

$$y_{i,t} = \mathbf{1}\{k_{i,t} \leq k_C(\pi_D)\}, \qquad i \in \{A, B\}. \tag{14}$$

Because firms do not communicate before entry in this construction, the rule does not implement the complete-information lower-cost selection rule. Instead, it uses a symmetric cutoff rule that is directly implementable from each firm's own private cost. The purpose of the construction is only to show that the sign reversal can survive private information.

If the audit conducted after period-t entry reveals that (14) was violated in period t, that is, if $y_{i,t} = 1$ with $k_{i,t} > k_C(\pi_D)$, then Z periods of reversion to the competitive Bayesian cutoff equilibrium begin in period $t + 1$. After the punishment phase, play returns to the collusive path.[5] During reversion, no auditing is used.

The representative firm's ex-ante per-period payoff on the collusive path before audit fees is therefore

$$\Pi^C(k_C; \pi_D) = \int_{k_L}^{k_C} \Big[(1 - F(k_C))\pi_M(\pi_D) + F(k_C)\pi_D - k\Big]\, dF(k). \tag{15}$$

The corresponding competitive payoff is

$$\Pi^N(\pi_D) = \int_{k_L}^{\bar{k}(\pi_D)} \Big[(1 - F(\bar{k}(\pi_D)))\pi_M(\pi_D) + F(\bar{k}(\pi_D))\pi_D - k\Big]\, dF(k). \tag{16}$$

Since audits occur with probability α, the net on-path collusive payoff equals $\Pi^C(k_C; \pi_D) - \alpha m$.

Because an over-entry deviation is detected only if an audit occurs one period later, the marginal excluded type's incentive constraint becomes

$$\begin{aligned}(1 - F(k_C(\pi_D)))\pi_M(\pi_D) + F(k_C(\pi_D))\pi_D - k_C(\pi_D) \\ = \alpha H_0\Big[\Pi^C(k_C(\pi_D); \pi_D) - \alpha m - \Pi^N(\pi_D)\Big].\end{aligned} \tag{17}$$

If the continuation-value loss falls sufficiently fast after a decrease in π_D, relative to the direct decrease in the one-shot deviation payoff, then the implicit function defined by (17) satisfies

$$\frac{dk_C(\pi_D)}{d\pi_D} < 0.$$

Thus, even with private entry costs and delayed verification, a decrease in duopoly profits can raise the collusive entry cutoff and thereby increase joint entry. By contrast, under competition the cutoff $\bar{k}(\pi_D)$ is increasing in π_D, so the same adverse profit shock reduces joint entry.

[5] One may also declare a deviation when $y_{i,t} = 0$ with $k_{i,t} \leq k_C(\pi_D)$; this does not bind in the constructions below.

Table 4: Joint entry probabilities in the lagged-verification collusion model.

c	$\pi_D(c)$	$k_N(c)$	$k_N(c)^2$	$k_C(c)$	$k_C(c)^2$
0.20	0.4444	0.6429	0.4133	0.5221	0.2726
0.22	0.4356	0.6346	0.4027	0.5228	0.2733
0.24	0.4268	0.6263	0.3922	0.5236	0.2741
0.26	0.4182	0.6179	0.3818	0.5245	0.2751
0.28	0.4096	0.6095	0.3715	0.5256	0.2762
0.30	0.4011	0.6011	0.3613	0.5270	0.2778

Numerical illustration. Specialize to linear demand $p = a - q$ with constant marginal cost c, so that

$$\pi_M(c) = \left(\frac{a-c}{2}\right)^2, \qquad \pi_D(c) = \left(\frac{a-c}{3}\right)^2.$$

Assume $F(k) = k$ on $[0, 1]$ (i.i.d. uniform entry costs), take $a = 11/5$, $\delta = 0.95$, $Z = 9$, $\alpha = 1$, and $m = 0.001$.

For the lagged-verification, $k_C(c)$ solves (17) (numerically) and the probability of joint entry on the collusive path is $k_C(c)^2$. Table 4 shows that $k_C(c)$ is increasing in c while the probability of joint entry on the competitive path $k_N(c)^2$ decreases in c. For each value of c, equation (17) can have more than one root. The table reports the smaller regular root. At this root, the continuation-value gap is positive, the marginal excluded type's constraint binds, and all lower-cost types strictly prefer to enter.