

Information Sharing and Collusion: Theory and Evidence from Poultry Processing

Avner A. Kreps*

Job Market Paper

June 18, 2026

Abstract

Information exchanges among competitors have become more prevalent in recent years. While they may facilitate collusion by resolving monitoring problems, inviting antitrust scrutiny, they may also affect outcomes by informing firms about their economic environment. Empirically distinguishing between these two forces is crucial for conduct testing and policy evaluation. This paper develops tools to do so and applies them to the Agri Stats information exchange among turkey processors. In the year after the exchange ended, the retail price of participating processors fell by 5.1% relative to other processors, and descriptive evidence suggests both conduct and information effects played a role. I specify and estimate a structural model that identifies changes to conduct while controlling for information effects. I reject the null hypothesis of no conduct change, and estimate that Agri Stats clients acted like they internalized 47% of other subscribers' profits while sharing information. Information effects of Agri Stats alone reduce prices by 1.3%, but are outweighed by conduct effects. Bounds on collusive prices under counterfactual restrictions on information sharing indicate that firms can still attain near-monopoly profits even when sharing aggregated information.

JEL: D22, D43, K21, L13, L41, L66

Keywords: Information sharing, collusion, private monitoring, conduct testing

I thank my thesis committee—Rob Porter, Vivek Bhattacharya, and Gastón Illanes—for their invaluable guidance and support. Federico Bugni, Luis Cabral, Sophie Calder-Wang, Mert Demirer, Igal Hendel, Jean-François Houde, Mar Reguant, Bill Rogerson, Takaaki Sagawa, Mike Sinkinson, Alexander Wolitzky, and numerous seminar participants at Northwestern provided helpful comments and feedback. Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

*Northwestern University, Department of Economics. Email: avner@u.northwestern.edu.

1 Introduction

Information exchanges have become a topic of growing interest in antitrust policy. Motivated by advances in data processing and the increasing adoption of common price aggregators and revenue management systems, regulators have taken a more aggressive posture towards information exchanges in the last several years, leading to the recent filing of numerous antitrust cases across myriad industries (Harrington, 2026). In these cases, competing firms share “competitively sensitive” information, such as recent or future indicators of prices and quantities, with each other. Plaintiffs allege the exchange has a *conduct effect* (i.e., facilitates collusion): it makes markets less competitive even without an explicit price- or quantity-fixing agreement. Defendants usually maintain that there is a procompetitive *information effect*: the exchange teaches firms about economic conditions, making markets more efficient.

Evaluating these arguments requires accurately characterizing the information effects of information sharing. Conduct testing—empirically distinguishing between oligopoly solution concepts—has a rich tradition in industrial organization (Bresnahan, 1982). However, workhorse models of competition typically assume complete information, and may therefore misattribute information effects to conduct effects. Given that information effects are ambiguously signed in theory (Kühn and Vives, 1995; Raith, 1996) and may be large relative to conduct effects (Vives, 2002), tests built on complete-information models can be biased in *ex ante* unforeseeable directions and magnitudes. Further, while regulators have proposed restrictions on information sharing to prevent firms from colluding, the literature has not formally characterized the extent to which different restrictions affect firms’ ability to collude. Thus, there is little economic guidance on how to detect collusion through information sharing, nor how best to regulate information exchanges.

In this paper, I specify a structural incomplete-information oligopoly model that allows for conduct testing and policy evaluation in information sharing settings. The model endogenously captures information effects, allowing for identification of conduct effects, and can be identified and estimated using typically available market-level data. I also develop a computational framework to place bounds on outcomes under counterfactual information exchanges, leveraging recent results in the theory of private monitoring games (Sugaya and Wolitzky, 2023). These bounds show how different restrictions on the kind of information firms share affect the worst-case collusive outcome.

I implement my methods in the context of a unique natural experiment in the poultry processing industry. A third-party data provider named Agri Stats ran a prototypical information exchange among the largest poultry (chicken and turkey) processors in the United States. In May 2019, the information exchange among turkey processors broke down due to

the threat of an antitrust lawsuit. In the year after the exchange ended, Agri Stats clients' prices fell by 5% relative to other producers, with suggestive evidence of both conduct and information effects at play. With this as motivation, I estimate my oligopoly model, test whether the end of the information exchange caused a change in firm conduct, and simulate worst-case collusive outcomes under alternative information structures.

I ultimately reject the null hypothesis that firms maximized their own expected profits with and without the information exchange. I estimate that while sharing information, Agri Stats clients set prices as though they internalized 47% of their rivals' profits. I find that Agri Stats raised average prices by 2.5%—and those of its clients specifically by 3.5%—but this overall effect is the composition of a -1.3% information effect and a 3.9% conduct effect. Tests relying on complete information therefore underestimate the magnitude of the conduct effect. Finally, to inhibit collusion, preventing full industry participation in an exchange is more important than mandating aggregation: worst-case average prices are nearly at the monopoly benchmark if all firms observe only their own shares and the total share, and are 8.9 percentage points lower if only Agri Stats clients share the full vector of prices and quantities.

I begin in Section 2 by demonstrating how information exchanges affect markets through both conduct and information effects. It is relatively well understood how information exchanges facilitate collusion by alleviating monitoring problems, enabling firms to better punish deviation from a dynamic equilibrium. The theory of information effects is more subtle: firms sharing information about fundamentals shifts prices across states, which directly impacts welfare and can also impact average prices through nonlinearities in firms' marginal profit functions. In addition, since the shared objects are frequently both signals and equilibrium outcomes (e.g., prices and quantities), sharing signals changes the extent to which those signals themselves are predictive of the economic environment. The overall effect depends on primitives: the extent to which firms learn about their own demand versus rival demand and prices, the underlying correlation structure of the uncertainty, and the closeness of competition, to name several.

Section 3 details the empirical setting. In the exchange, Agri Stats clients shared recent prices and quantities, as well as indicators of future prices and quantities in the form of the number of chicks and poults (baby turkeys), which Agri Stats circulated among subscribers in weekly or monthly reports. In 2016, a class of chicken purchasers sued Agri Stats and its chicken subscribers. Concerned about potential litigation risk, turkey processors dropped out of the information exchange in May 2019, causing Agri Stats to suspend their turkey reports. This quasi-experimental variation in information sharing provides an ideal case study for investigating how and through what channels information exchanges affect market

outcomes.

Section 4 begins the empirical analysis by evaluating the aggregate effects of Agri Stats' exit from turkey. I use a difference-in-differences with turkey processors which did not subscribe to Agri Stats as a control group. I estimate that Agri Stats subscribers lowered their retail prices by 5.1% in the year after the end of the information exchange, but relative prices slowly rebounded to near pre-May 2019 levels. This effect is not driven by broader trends in poultry: I estimate a null effect of the turkey information exchange ending on Agri Stats clients' chicken prices as a placebo. Digging into these aggregate effects reveals further hallmarks suggestive of the presence of both conduct and information effects: the short-run effects were larger in markets which were structurally more susceptible to collusion, while Agri Stats clients' prices were more correlated with common demand shocks and with each other while the information exchange was ongoing.

I outline my structural model capturing information effects in Section 5. The model nests workhorse static supply-and-demand models such as that of [Berry et al. \(1995\)](#), but does not assume complete and perfect information. Instead, firms have incomplete information over potentially autocorrelated demand shocks, which are unobserved, and cost shocks, which are private information. They learn about these objects through *signals*: arbitrary functions of past prices and quantities and future prices that help them predict demand and their rivals' prices. When firms share information, they observe more informative signals, giving them more precise information about their residual demand curve. The key tractability assumption is that firms treat their rivals' prices as following a first-order Markov law of motion, conditional on observable supply and demand variables, meaning that firms need only incorporate period- t signals in order to update their beliefs about their residual demand curve from period t to period $t+1$. At a competitive solution, which I call a Bertrand-Markov equilibrium, firms set prices to maximize expected profits in each period given beliefs that update in a manner consistent with their rivals' long-run actions.

I then apply this model to the Agri Stats context. I capture information effects with different signal functions before and after the exchange: Agri Stats clients observed each others' prices and quantities while sharing information, while other firms just observed their own. Section 6 demonstrates that the entire system of demand and supply parameters, including beliefs, is identified from market-level data while firms are playing a Bertrand-Markov equilibrium. I assume that equilibrium is played after the exchange ends. I show that the model can be estimated using a combination of standard methods and Bayesian techniques, and that it provides sensible results.

Section 7 applies standard conduct testing techniques to test for the existence and magnitude of conduct effects while accounting for the information effects caused by Agri Stats'

exit. Using a test statistic based on the restriction that there was no discontinuity in the difference between Agri Stats subscribers’ average cost shocks and those of other producers in May 2019, I reject the null hypothesis that information effects alone can explain the Agri Stats effect. To quantify the magnitude of the conduct effects, I fit a parameterized conduct model and find that Agri Stats clients set prices as though they internalized 47% of their fellow subscribers’ expected profits while sharing information. The analogous conduct parameter with a complete-information supply side is not statistically distinguishable from zero, underscoring that not accounting for information effects can lead to biased estimates of conduct effects.

Finally, Section 8 simulates the effects of counterfactual information exchanges to unpack the economic mechanisms of Agri Stats and assess potential policy restrictions on data sharing. Agri Stats raises prices by 2.5% overall, which is the composition of a -1.3% information effect and a 3.9% conduct effect. The information effect alone is roughly 80% due to the direct effects of firms learning about the state, primarily about their rivals’ demand and prices, with most of the remainder attributable to equilibrium effects of firms best responding to changes to their rivals’ prices. However, while the estimated conduct parameter allows for computing the conduct effects of the Agri Stats exchange specifically, it may not apply to counterfactual monitoring structures, and finding exact equilibria in repeated games with private monitoring caused by incomplete information exchanges is known to be difficult. Thus, to characterize the potential effects of collusion under data sharing restrictions, I develop a computational framework to bound prices under alternative monitoring structures. The idea is to leverage necessary conditions that a repeated-game equilibrium must satisfy by solving a constrained maximization problem, where the objective is to maximize prices and the constraints are the necessary conditions. The main condition, from [Sugaya and Wolitzky \(2023\)](#), formalizes the idea that monitoring structures that allow firms to better detect deviation from their rivals enable more collusive outcomes.

Using this method, I assess the tradeoff between aggregation and participation: should regulators worry more about aggregate signals observed by all firms, or precise signals seen only by a subset? This tradeoff is policy-relevant given that the antitrust agencies have become more critical of aggregated exchanges in recent years, while the economics literature has disagreed over the years on the extent to which aggregation inhibits collusion.¹ I find that preventing full participation in an information exchange is more important than mandating

¹[Radner et al. \(1986\)](#) show that aggregation inhibits cooperation in a repeated partnership game in which each participant observes only an aggregate signal of actions. For collusion specifically, [Carlton et al. \(1997\)](#) state that “aggregating the data largely removes the value of information in facilitating collusion” (p. 434). Conversely, [Awaya and Krishna \(2020\)](#) present conditions under which firms can use aggregate sales information to enforce an equilibrium with near-monopoly profits.

aggregation of information. If all Agri Stats clients (which comprise roughly two-thirds of the market) were to continue sharing the full vector of price and quantity information, the upper bound collusive price attains 86% of the distance between monopoly and competitive prices on average, but this increases to 95% if all firms only observe total quantities and their own prices and shares, while under no information sharing 71% is attainable.

These results have several policy implications. Information effects can be economically meaningful, and thus it is important to account for them when attempting to detect conduct via information exchange. Further, information provision itself can make markets more efficient, if there were a way to prevent firms from using it to collude. However, even exchanges of aggregate information can facilitate near-monopoly collusion if the entire market participates. As such, the results support the notion, recently reflected in regulatory policy, that bright-line “safe harbors” on information exchanges satisfying aggregation requirements and timing requirements are a bad idea.

Related literature. The theoretical literature on information exchanges has historically been bifurcated. One literature, dating back to [Stigler \(1964\)](#), studies how information exchanges facilitate dynamic collusion through resolving monitoring problems. [Green and Porter \(1984\)](#) provide a formal treatment of how uncertainty impacts collusion, and subsequent contributions study collusion with different monitoring structures ([Athey and Bagwell, 2001; 2008](#); [Harrington and Skrzypacz, 2007](#); [Sugaya and Wolitzky, 2018b](#)) and the impact of self-reported communication ([Harrington and Skrzypacz, 2011](#); [Awaya and Krishna, 2016](#)). More recent papers derive bounds that repeated-game equilibria with private monitoring must satisfy ([Sugaya and Wolitzky, 2018a; 2023](#); [Awaya and Krishna, 2019; 2020](#)), providing results that I leverage in this paper. Another literature studying how information sharing affects outcomes in static equilibrium began with [Novshek and Sonnenschein \(1982\)](#); a series of related papers, reviewed by [Kühn and Vives \(1995\)](#) and [Raith \(1996\)](#), concludes that information effects are ambiguously signed and depend on modeling details such as the type of uncertainty and firm decision variable.² While these two literatures have largely evolved in parallel, I contribute an empirical framework weighing conduct and information effects against one another, enabling aggregate quantification and counterfactual predictions. My model also incorporates more general forms of uncertainty than are typically studied in this literature—while this generality precludes analytical results, it enables characterization of information effects in more realistic settings.

²For example, the information effects of sharing information about a linear demand intercept are procompetitive in Cournot competition and anticompetitive in Bertrand competition, while these are reversed for sharing information about costs. More recent contributions to this literature include [Asker et al. \(2020\)](#), who study information sharing in auctions, and [Bonatti et al. \(2017\)](#) and [Sweeting et al. \(2024\)](#), who incorporate information sharing in dynamic models with signaling incentives.

This paper also relates to an empirical literature investigating case studies of information exchanges. Findings have been mixed: [Albæk et al. \(1997\)](#) and [Byrne et al. \(2024\)](#) find anticompetitive effects of information sharing among Danish ready-mixed concrete producers and Australian gasoline retailers respectively,³ while [Doyle and Snyder \(1999\)](#) and [Borenstein \(2004\)](#) find procompetitive effects among automakers and airlines respectively.⁴ These papers typically study the aggregate impact of information sharing on market outcomes, without distinguishing whether and to what extent they are due to conduct or information effects. More recently, [Calder-Wang and Kim \(2025\)](#) and [Calder-Wang et al. \(2025\)](#) formally test whether adoption of a common pricing algorithm led to changes in competitive conduct among apartment owners. While our structural models are based on similar demand systems, our models of information effects differ due to the setting: it is difficult to model the signal transmitted by the pricing algorithm they study. [Calder-Wang and Kim \(2025\)](#) do not microfound information effects, instead assuming that non-subscribers have complete information but make markup-lowering mistakes, while [Calder-Wang et al. \(2025\)](#) test whether prices are rationalizable under any information structure consistent with Bayes-correlated equilibrium. Conversely, I study a classical information exchange in which firms directly share actions and outcomes, enabling me to model the information structure and account for information effects in a more direct manner.

Finally, this paper relates to the literature on conduct testing and estimation. One important and related early contribution is [Porter \(1983\)](#) which, like this paper, estimates conduct in a setting with incomplete information about demand shocks. This literature has received renewed attention recently, with contributions from [Berry and Haile \(2014\)](#) on identification of conduct; [Backus et al. \(2021\)](#), [Duarte et al. \(2024\)](#), and [Dearing et al. \(2025\)](#) on instrument relevance, model selection, and testing; and [Miller and Weinberg \(2017\)](#) providing an important empirical application. While these papers primarily restrict attention to complete-information supply-side models, this paper builds on these methods by incorporating incomplete and imperfect information, facilitating conduct testing in settings where firm conduct covaries with the information structure.

³While information sharing among Australian gasoline retailers raised prices overall, [Byrne et al. \(2025\)](#) find that removing one retailer from the exchange but leaving the rest intact raised prices by even more.

⁴Similar settings studied by prior papers include price transparency regulations, which combine information exchanges among firms with potential reductions in demand-side search costs ([Rossi and Chintagunta, 2016](#); [Luco, 2019](#); [Grennan and Swanson, 2020](#); [Ater and Rigbi, 2023](#); [Montag et al., 2026](#)), and common pricing algorithms, which recommend prices to competing firms using information that is pooled across competitors ([Assad et al., 2024](#); [Calder-Wang and Kim, 2025](#); [Calder-Wang et al., 2025](#)). [Clark and Houde \(2026\)](#) show indirect evidence of collusion via information exchange: firms that were shown to have colluded *ex post* were more likely to have joined an information sharing platform.

2 Conceptual framework

To provide intuition for the relevant economic forces, I briefly review how information sharing can affect market outcomes through conduct and (primarily) information effects, and what primitives affect these forces. The takeaway is that naive conduct tests can be biased in unexpected ways when information effects are also at play, underscoring the need to model information effects.

Suppose F single-product firms repeatedly compete by simultaneously setting prices. Firm f has differentiable profit function $\pi_f(p, \xi)$, where $p = (p_1, \dots, p_F)$ is the price vector and $\xi = (\xi_1, \dots, \xi_F)$ is a vector of demand shocks. Prices are strategic complements, while a higher realization of ξ_f incentivizes firm f to respond by raising its price and its rivals to respond by lowering their prices. The consumer surplus function, $CS(p, \xi)$, is decreasing in all p , increasing in all ξ , and has the standard feature that consumers care more about p_f when firm f 's residual demand is higher.⁵

Firms do not observe ξ directly. Instead, before setting prices, each firm observes a signal y_f of ξ . In practice, these signals are often past prices or quantities, which are informative of past demand and hence (if demand is autocorrelated) future demand. Bayes-Nash equilibrium prices for firm f satisfy the first-order condition

$$\mathbb{E}_{\xi, y_{-f}|y_f} [\pi'_f(p_f^{NS}(y_f), p_{-f}^{NS}(y_{-f}), \xi) | y_f] = 0 \quad \forall y_f, \quad (1)$$

where $\pi'_f(\cdot)$ is firm f 's marginal profit with respect to its own price, and the NS superscript denotes “no sharing.”

To build intuition, consider a simple example in which the only uncertainty is over $\xi_f \in \{\xi_f^L, \xi_f^H\}$, with a uniform prior. Figure 1(a) plots firm f 's demand curves at each value of ξ_f as p_f varies, with demand in the high state drawn in blue and in the low state drawn in orange.⁶ Rivals' prices are held fixed at the no-sharing equilibrium values. Each demand curve has an associated marginal profit function, drawn in dashed lines. Assuming for simplicity that y_f is uninformative about ξ_f , firm f sets its price p_f^{NS} such that its expected marginal profit—the vertical average of the marginal profit in the two states, drawn in the green dashed line—is zero.⁷ This yields a high quantity $q_f^{NS}(\xi_f^H)$ in the high-demand state and a low quantity $q_f^{NS}(\xi_f^L)$ in the low-demand state.

⁵Formally, for $f \neq f'$, $\partial^2 CS / \partial p_f \partial p_{f'} < 0$, $\partial^2 CS / \partial p_f \partial \xi_f < 0$, and $\partial^2 CS / \partial p_f \partial \xi_{f'} > 0$, while π_f is single-crossing in $(p_f, p_{f'})$, (p_f, ξ_f) , and $(p_f, -\xi_{f'})$.

⁶The plot is based on a duopoly facing logit demand, with indirect utilities $u_{if} = \xi_f - 3p_f + \varepsilon_{if}$, $\xi_1 \in \{2, 5\}$, $\xi_2 = 4$, and constant marginal costs equal to 0 for both firms, from the perspective of firm 1.

⁷One could also think about the uncertainty on ξ_f as the residual uncertainty after conditioning on a particular value of y_f .




Figure 1: Information effects of information sharing

Now suppose that firms start sharing information: all firms observe the full signal $y = (y_1, \dots, y_F)$ before setting prices. Antitrust enforcers might be concerned that this may facilitate collusion by resolving monitoring problems. The basic mechanism, originally described by [Stigler \(1964\)](#), is well-known. When colluding, firms set statically suboptimal prices, creating an incentive to deviate that is dissuaded by the threat of a price war. If a firm imperfectly observes its rivals' actions (say, only observing its own share), then it faces an inference problem: a low share is consistent with both a rival cutting its price and a low realization of its own demand shock (or a high realization of its rivals'), inhibiting its ability to punish deviation. Sharing information ameliorates this problem in two ways: directly, if signals directly inform firms of their rivals' actions, and indirectly, by giving firms more information about ξ , making them better at inferring their rivals' actions.

However, the information exchange affects prices and welfare even if firms keep maximiz-

ing expected profits. With the exchange, Bayes-Nash equilibrium prices solve

$$\mathbb{E}_{\xi|y} [\pi'_f(p_f^S(y), p_{-f}^S(y), \xi) | y] = 0 \quad \forall y, \quad (2)$$

where S stands for sharing. Comparing equations (1) and (2), there are three potential forces acting on equilibrium prices and welfare, which I group under the umbrella term “information effects”: the direct effect, equilibrium effect, and signal distribution effect. I walk through each force in the context of the example, with each corresponding to a panel in Figure 1.

Direct effect. Firm f conditions on the full signal vector y instead of just its own signal y_f , letting it better correlate prices with residual demand. This can change expected prices and welfare through nonlinearities in the marginal profit and CS functions. In the example, suppose again for simplicity that y fully informs firm f about ξ_f . Panel (b) shows how firm f updates its prices while only incorporating this new information, before any changes to its rivals’ strategies. Because firm f learns ξ_f , it sets its prices in each state ($p_f^S(\xi_f^L)$ and $p_f^S(\xi_f^H)$) such that the marginal profits are zero in both states, raising its price in the high-demand state and lowering it in the low-demand state. However, because firm f ’s marginal profit is steeper when demand is high than when it is low, firm f lowers its price in the low-demand state by more than it raises it in the high-demand state, shifting the expected price down. The reverse could occur: if marginal profit is steeper under ξ_f^L , then the direct effect shifts f ’s expected price up.

This reallocation also causes changes to CS that are asymmetric across states. CS—the area under the demand curve—decreases by the red shaded area in the high-demand state and increases by the green shaded area in the low-demand state. Because CS is more sensitive to p_f when f ’s demand is higher, the marginal increase in p_f in the high-demand state outweighs a marginal decrease in p_f in the low-demand state. Thus, even when expected prices go down, it is not guaranteed that expected CS increases. One can see this graphically in panel (b): the horizontal dimension of the CS change is the price change, which is greater in the low-demand than the high-demand state, but the vertical dimension is the demand curve, which is greater in the high-demand than the low-demand state.

Equilibrium effect. Now firm f responds to its rivals changing their strategies from $p_{-f}^{NS}(y_{-f})$ to $p_{-f}^S(y)$. In panel (c), I assume that rivals also learn ξ_f , causing them to lower their prices at ξ_f^H and raise their prices at ξ_f^L . This lowers firm f ’s residual demand curve in the high-demand state and raises it in the low-demand state, from the faded lines to the opaque lines, causing price movements in the opposite direction—relative to the direct effect, $p_f^S(\xi_f^H)$ decreases and $p_f^S(\xi_f^L)$ increases. This too need not be symmetric across states,

or with the direct effect: if firm f 's rivals change their prices by more in one state than another (or residual demand is more sensitive in one state than another), then equilibrium effects can outweigh direct effects. This occurs in the example: information sharing causes $p_f^S(\cdot)$ to decrease in *both* states relative to p_f^{NS} , and thus consumer surplus increases in both states. Intuitively, when the state is unknown, each firm sets its price to try to capitalize on the state in which it has higher demand than its rivals, but when the state is known, each firm can lower its price when its relative demand is low without worrying about reducing its margin when its relative demand is high, intensifying within-state competition.

Signal distribution effect. Finally, if the signals y are themselves equilibrium objects, then information sharing changes the distribution of $\xi \mid y$. Panel (d) highlights the movement in equilibrium prices and quantities when not sharing to when sharing, with arrows denoting the shift in equilibrium prices and quantities from the no-sharing to the with-sharing equilibrium, holding the signal structure fixed. When not sharing, firm f 's price does not contain any information about ξ_f —its price is exactly the same in both states, as firm f does not know ξ_f when setting p_f . Its quantities, however, vary substantially across states, as it soaks up the residual variation in demand. With more information about demand, firm f can correlate its price to the demand shock, making it a more informative signal of the demand state. However, firm f 's quantities move closer together, making them vary less with the demand state. Thus, a running assumption for this example—that firms' signals remain fixed with and without information sharing—may be violated in practice.

In summary, if a researcher had data generated with and without information sharing, and fit a model that assumes away information effects, they may make erroneous conclusions about whether the solution concept truly changed. While information effects tend to lower prices in this example, several economic primitives that are kept static in the example bear on the magnitude and direction of information effects, and may therefore reverse the direction of the effect. First, the fundamental underlying uncertainty in the example was about demand shocks, but there may be uncertainty about rivals' costs as well; such uncertainty is known to create different information effects (Raith, 1996). Second, in the example, firm f only learns about its own demand shock; learning about rivals' shocks causes movement in opposite directions. Third, the closeness of competition affects the extent to which equilibrium versus direct effects apply—in the limit, monopolists only have direct effects. Fourth, the example features independent-valued demand shocks, as ξ_{-f} was deterministic, but firms learning about a common-valued demand shock causes the equilibrium effect to reinforce the direct effect, flipping its sign.⁸ As such, it is important to directly model information effects, and

⁸Figure A.1 graphically shows (a) direct effects of a firm learning about its rival's demand shock, causing it to negatively correlate its price with the demand state, and (b) equilibrium effects under common-valued

estimate the primitives underlying them, to answer the question at hand.

3 Institutional details and data

3.1 Overview of poultry processing

This subsection briefly details background information about the poultry processing industry in the US that is important to know for the analysis to come. Section 5.3 will connect these facts to the implementation of the model.

Poultry processing is the preparation of live fowl (primarily chicken and turkey) for human consumption. It includes all steps in the production process starting from a live bird and finishing with a product that is sold to retailers, restaurants, or other food service firms. One firm, known as a processor or integrator, typically handles all steps of this process from slaughter to wholesale, as poultry must be quickly processed after slaughter before being chilled or frozen in order to prevent spoilage.⁹ The processors, or integrators, are present throughout the supply chain from wholesale to retail; while they contract with, rather than own, poultry farms that grow the birds themselves, the processors decide how many chicks and poults (baby chickens and turkeys) are placed (Zheng and Vukina, 2007). While there are differences between turkey and chicken processing, most notably relating to size and maturity time, the process is similar: feed components are the same, the USDA regulates the processing plants in the same way, and there are many suppliers in common (Fernández-López et al., 2010; USDA).

Wholesale poultry prices are relatively flexible.¹⁰ Retailers and other downstream consumers typically buy processed poultry through purchase orders, which authorize a retailer to purchase a prespecified quantity of poultry at a prespecified price, typically expressed as a proportion of a common price index that adjusts for cost shocks. While these purchase orders are occasionally components of longer-term master sales agreements, these agreements are flexibly renegotiated and generally do not constrain either party to purchase or sell any particular quantity of poultry—both processors and retailers generally do not list ongoing purchase orders as contractual obligations in their financial statements, as they are cancellable on short notice. Substantial variation in the terms of these agreements across retailers within a processor is relatively rare (Harjani, 2025).

demand shocks, in which residual demand curves across states diverge, rather than converge.

⁹The main exception is firms that purchase pre-processed raw poultry and process it further, often turning it into deli meat. Such firms, known as “value-added” processors, produce a very small fraction of the poultry consumed in the US, and I exclude them from my analysis; see Section 3.3.

¹⁰This paragraph draws on processor and retailer 10-Ks and financial statements available from the SEC EDGAR database.

(a)		(b)	(c)		(c.1)	(c.2)	(c.3)	(d)	(d.1)	(d.2)	(e)	(e.1)	(e.2)	(f)	
LINE	PRODUCT CODE	AGRI STATS/COMPANY PRODUCT DESCRIPTION	COMPANY POUNDS	COMPANY MIX %	NATIONAL POUNDS	NATIONAL % MIX	COMPANY NET PRICE	NAT'L NET PRICE	NAT'L TOP 25%	NAT'L RANK	NAT'L VAR	NAT'L TOP 25 VAR	ECONOMIC IMPACT DOLLARS	VAR NAT'L	VAR TOP 25
1	506/AD 00000A554	DRUMSTICK, V ADD, LW, P A/LR, CHILL, JUM	249,571	6.60	1,293,884	0.26	71.80	75.51	79.54	7-8	-3.71	-7.74	-9,255		-19,319

Figure 2: Product-level pricing information in Agri Stats reports

Source: Second Amended Complaint.

Poultry processors typically rely on private data providers to learn about market conditions. While the USDA publishes some aggregated price indices and summary statistics, the input data are voluntarily reported by processors. They are therefore viewed by market participants as less comprehensive than private sources and potentially subject to strategic nonreporting (Leonard, 2017; Boussios, 2023).

3.2 Agri Stats

Agri Stats is a data analytics and benchmarking firm that, at a high level, collects and exchanges information about prices, quantities, costs, and inventories among the 15-20 largest poultry processors in the US. These processors collectively accounted for 80% and 95% of US turkey and chicken production in the 2010s (Holdheim and Tameez, 2022).¹¹ Agri Stats has been sued several times recently—most notably, by the Department of Justice in late 2023—on the grounds that its information exchange facilitated collusion.¹²

Agri Stats exchanged information through printed reports, examples of which are shown in Figures 2 and 3. Figure 2 shows price and quantity information for one product. Columns (c) through (c.3) give the firm information on its market share, while columns (d) through (d.2) list the firm’s price for this product relative to the industry-wide average and 75th percentile price. Finally, column (e) lists that the firm’s price is seventh-highest out of eight firms making this product, with the remaining columns listing Agri Stats’ calculation of the impact on profits of not pricing at the national average or 75th percentile. Figure 3 contains plant-level information on the number of chicks placed at farms serving different processing plants, which directly feeds into future production. Other parts of the reports, examples of which are not available publicly, include plant-product-level prices, quantities, product mixes, and input usage (Leonard, 2017; Tunheim, 2026). The reports were disseminated on a weekly or monthly basis, depending on the data series, and contained up-to-date data from the past time period.

According to Agri Stats and its clients, processors primarily used the information ex-

¹¹It also operated among pork and egg processors, although it counted less of these markets as its clients.

¹²Much of the information presented below is taken from documents filed in these lawsuits. All footnote citations to legal documents are from *US et al. v. Agri Stats, Inc.*, case no. 0:23-cv-03009-JRD-JFD.

BREEDER SECTION PARTICIPANTS

(a) Participating plants

(a)										(b)		
FEMALE CHICKS-AVG WEEKLY PLACED										MALE CHICKS-AVG WEEKLY PLACED		
LIN	%	PLT.	SP	NUM.	% LAST	THIS	MONTH	YR AGO	MALE/	% LAST	THIS	MONTH
					YR MD	MONTH	YR AGO	100 FEM		YR MD	MONTH	YR AGO
1	100					29,952					2,774	
2	97					23,813					3,034	
3	94					21,740					2,596	
4	91				136.44	17,562	12,871	14.06	133.33		2,469	1,852
5	88				80.51	17,560	21,563	12.67	89.08		2,200	2,470
6	85				132.70	17,342	13,068	14.33	215.30		2,486	1,155
7	82				44.16	16,034	36,312	9.12	34.79		1,463	4,204
8	79					14,953					1,610	
9	76	GRNV			165.41	14,468	8,747	11.68	172.66		1,690	979
10	73				201.26	14,338	7,124	12.25	193.10		1,757	910
11	70				85.85	12,750	14,851	16.16	104.17		2,060	1,978
12	67				166.62	11,727	7,038	12.97	165.72		1,521	918
13	64				98.19	11,098	11,302	13.64	98.66		1,514	1,535
14	61					10,506					2,266	
15	58				68.33	10,250	15,000	12.20	61.73		1,250	2,025
16	55					9,690					1,239	
17	52				91.30	9,668	10,589	15.15	90.75		1,465	1,614
18	48					9,286					1,206	
19	45				128.21	9,180	7,160	11.11	120.28		1,020	848
20	42				91.82	7,894	8,597	14.47	91.30		1,142	1,251
21	39				91.82	7,894	8,597	14.18	91.30		1,120	1,226
22	36				81.89	5,813	7,099	18.06	108.45		1,050	968
23	33	COLA				5,800					825	
24	30	ARCADA			100.00	5,675	5,675	11.23	100.00		638	638
25	27				82.00	5,625	6,860	12.00	84.38		675	800
26	24				63.43	4,995	7,634	13.14	74.62		656	879
27	21				126.36	4,743	3,754	12.22	126.47		580	458
28	18					4,717					570	
29	15				71.58	4,716	6,589	13.65	72.66		644	886
30	12					4,488					632	
31	9				88.28	3,278	3,713	14.73	97.64		463	494
32	6					0	5,508	0.00			0	1,367
33	3					0	8,597	0.00			0	1,251
34												
35												
36												
Total										145.37	44,631	30,703

(b) Anonymized plant-level inventory data

Figure 3: Plant-level inventory information in Agri Stats reports

Source: Second Amended Complaint.

Note: Plant names, ownership, and locations are redacted *ex post*.

change to more accurately forecast market conditions. For example, in a 2009 earnings call, one executive at a poultry processor cited Agri Stats data on prices and the number of egg-laying hens placed by rivals as the basis for saying that “I see a lot of information from Agri Stats that tells me that nobody’s going to ramp up” production, and that therefore prices would remain elevated (Leonard, 2017). At the same earnings call, an Agri Stats executive stated that “the goal is obviously to provide the fact based forecasting and analysis that we talked about, that covers the supply, demand, price conditions and so forth,” while the reports themselves stated that one purpose was “to help forecast Broiler [chickens] & pounds produced for future months.”¹³ Agri Stats and its clients have also claimed that the reports were used to find production efficiencies: for example, Agri Stats states on its website that “protein producers depend on Agri Stats’ reports to help them identify opportunities to reduce production costs to keep prices low” (Agri Stats, 2023).¹⁴

While the reports themselves are largely outside the public domain, certain details of the reports have been made public, each of which has implications for the analysis to come. First, while aspects of the reports were either aggregated (as in Figure 2) or anonymized, the plant level data in the report could be deanonymized, and processors were allegedly

¹³Second Amended Complaint, p. 21.

¹⁴Some of these purported efficiencies may not be benign: separate lawsuits allege that processors used the Agri Stats reports to reduce costs by fixing wages for processing plant employees, as well as payments to poultry growers (Scarcella, 2026).

successful in doing so; as such, I will treat the Agri Stats reports as revealing information tied to specific firm identities and at the most detailed level. Second, the reports were based on automatic data transfers that Agri Stats regularly audited, not self-reports; therefore, Agri Stats exchanged true information without its clients being subject to any truth-telling or obedience constraints (Awaya and Krishna, 2016). Third, the reports were only available to subscribing processors that provided their own data, while Agri Stats refused to sell the reports to poultry purchasers and workers; thus, the information exchange only affected the supply side of the market, without any potential countervailing effects of price transparency on the demand side (Luco, 2019; Ater and Rigbi, 2023).

Agri Stats stopped publishing its turkey reports in May 2019, testifying that they did so because of a lack of data due to subscribers dropping out. This occurred after the first lawsuit against chicken processors began in late 2016, but before the first turkey lawsuit in late 2019. An Agri Stats executive testified that processors who dropped out “cited concerns regarding the broiler [chicken] litigation,” which was echoed in a 2020 internal Agri Stats presentation, suggesting that turkey processors hoped to avoid legal trouble (which chicken processors already were in) by exiting Agri Stats before facing litigation.¹⁵ Meanwhile, Agri Stats continued publishing its chicken reports without any changes until it entered into a behavioral remedy to settle with the Department of Justice in May 2026.¹⁶ The clean variation in information sharing in turkey provides an ideal setting for studying the effects of information sharing, while chicken provides a useful placebo given the overlap in the set of suppliers, inputs, and regulatory structures in the two industries.¹⁷

3.3 Data

My primary dataset is the NielsenIQ retail scanner data, provided by the [Kilts Center](#) for Marketing at Chicago Booth. This dataset contains weekly prices and quantities from a nationally representative set of retailers. I use data from two product modules (the most granular product category in the data): fresh meat and sliced deli meat. The predominant turkey product in the fresh meat module I use is ground turkey, which, unlike whole turkey, is not a seasonal product—the product module with whole turkey only has data beginning in 2020, foreclosing its inclusion in the analysis. The timeframe is June 2017 to May 2022, which

¹⁵Plaintiffs’ Opposition to Defendant’s Motion to Dismiss.

¹⁶The settlement prevents Agri Stats from sharing any nonpublic price information, requires that it aggregate any exchanged quantity or cost data, prevents sharing any data more recent than 45 days old, and requires Agri Stats to sell the reports to any potential purchaser.

¹⁷Agri Stats also shut down most aspects of their pork reports around this time. I focus on turkey for my analysis as there is cleaner variation in the information available to turkey processors than to pork processors, Agri Stats had a wider reach in turkey, and chicken provides a more natural comparison for turkey than for pork.

includes two years before to three years after the end of the turkey information exchange.¹⁸

I augment the scanner data with several auxiliary data sources. First, I incorporate the NielsenIQ consumer panel data, which contain purchase-level information from a panel of households, into my demand estimation procedure. Second, I use the [USDA’s Meat, Poultry, and Egg Product Inspection \(MPI\) directory](#), which contains ownership, addresses, licenses, and “doing business as” information of all USDA-regulated poultry processing plants. The current directory is publicly available online; I obtained monthly snapshots from throughout my time period via a Freedom of Information Act request. I geocode all processing plant locations and compute driving distances between processing plants and stores using the Google Maps API. Third, I obtain price indices of inputs, including poultry feed and diesel, as well as the CPI-U (which I use to deflate all prices to May 2019 levels), from the [Federal Reserve Bank of St. Louis’s FRED database](#). Fourth, I hand-collect ownership information of brands in the scanner data from the internet. Fifth, I identify which firms subscribed to the Agri Stats reports using information from legal filings.

I use the following procedure to process the data. I identify poultry products and define categorical product characteristics using string searches on product descriptions. I remove “value-added” processors—firms that purchase preprocessed turkey from other processors and resell it after further processing, most commonly turning it into deli meat—by removing firms that do not own a processing plant, as it is unclear whether these firms purchase turkey from Agri Stats subscribers. This step removes less than 1% of total turkey sales from the data. For private label products, I identify the original processor from publicly available sources, and exclude private label products whose origin I cannot identify in this way. I then aggregate these data to the barcode (UPC)-store-month level to arrive at the dataset used for the reduced-form evidence presented in Section 4. For estimation of the structural model and counterfactuals (Sections 6-8), I further refine the dataset by removing tiny-share products, filtering to regions with at least one panelist purchase in each month, and aggregating to the product-region-month level, where a product is defined as a unique value of brand- $\mathbb{1}$ (breast meat)- $\mathbb{1}$ (whole smoked turkey) for fresh turkey and brand- $\mathbb{1}$ (low fat) for deli turkey.¹⁹ I aggregate and report all prices and quantities on a per-pound basis throughout.

¹⁸I start my analysis in June 2017 to avoid contamination from a prior antitrust claim of poultry market manipulation in late 2016-early 2017 ([Welshans, 2017](#)), and end my analysis in May 2022 due to the outbreak of the Highly Pathogenic Avian Influenza (i.e., avian flu), which caused the death of millions of birds and substantially impacted poultry prices in mid-2022 ([Grossen, 2023](#)).

¹⁹For more details, see Appendix B.1. Throughout the paper, I use Nielsen Designated Market Areas as my primary geographical delineation, and refer to it as a region. A Nielsen DMA is a collection of counties that all receive the same television stations and advertising. They are commonly used to define geographic markets in studies of consumer packaged goods ([Nevo, 2001](#); [Miller and Weinberg, 2017](#); [Bhattacharya et al., 2026](#)).

	Fresh turkey		Deli turkey	
	Full sample	Est. sample	Full sample	Est. sample
Number of regions	202	110	205	118
Total monthly revenue	\$42.0m	\$40.4m	\$22.6m	\$21.5m
Number of products	19.8	8.7	43.3	15.5
	[12, 26]	[7, 10]	[30, 57]	[13, 18]
Number of firms	4.4	5.0	9.5	10.7
	[3, 6]	[4, 6]	[7, 12]	[9, 12]
Number of Agri Stats clients	2.7	3.2	4.7	5.0
	[2, 4]	[2, 4]	[4, 6]	[4, 6]
Agri Stats client share	0.67	0.61	0.77	0.71
	[0.38, 0.96]	[0.34, 0.89]	[0.63, 0.92]	[0.59, 0.83]
HHI	6,113	5,410	3,557	3,088
	[4,582, 7,440]	[4,312, 6,559]	[2,709, 4,023]	[2,522, 3,457]

Table 1: Summary statistics

Note: Est. sample refers to the sample used for the structural model. Number of products, firms, Agri Stats clients, Agri Stats client share, and HHI each feature the average followed by the interquartile range, both across region-months. Products are defined as a UPC for the full sample and, for the estimation sample, as a unique value of brand- $\mathbb{1}$ (breast meat)- $\mathbb{1}$ (whole smoked turkey) for fresh turkey and brand- $\mathbb{1}$ (low fat) for deli turkey. Shares are computed using sales volume in pounds.

Table 1 displays summary statistics of both samples. Both fresh and deli turkey are highly concentrated, but fresh turkey moreso: the average market (region-month) has an HHI of roughly 3,500 and around 10 competing firms for deli turkey, as compared to an average HHI of 6,000 and four to five competing firms for fresh turkey. Agri Stats clients dominate both markets, representing 60-70% of the fresh meat market and more than 70% of the deli market by volume.²⁰ Notably, however, there is substantial variation across markets in the Agri Stats share, particularly in fresh turkey; the top quartile of markets are nearly completely dominated by Agri Stats, with a share of more than 96%, while Agri Stats clients collectively have a share of less than 40% in the bottom quartile. Finally, while the filters used to construct the estimation sample result in nearly half of the regions being excluded, the remaining regions comprise more than 95% of the total turkey sold, and the concentration and penetration of Agri Stats subscribers is broadly similar in both samples.

²⁰This is slightly lower than the 80% figure cited in the complaint and several other sources. One may not expect this figure to line up exactly, as I analyze only a subset of turkey sold at retail, and grocery store markets only represent a subset of the turkey sold in the US; a substantial fraction of turkey is purchased by the food service industry, including restaurants, catering, and cafeterias at institutions like schools, hospitals, and prisons.




Figure 4: Average turkey price and sales quantity over time

Note: Average prices are weighted by total within-UPC-store quantity throughout the time period. The vertical dashed line represents May 2019.

4 Motivating evidence

I begin my analysis by establishing several facts about what happened after the Agri Stats information exchange ended. Before doing so, it is worth discussing what one might expect to happen when a long-running information exchange ends. If firms used the information exchange to collude, then ending the exchange will cause participating firms' prices to go down, while the information effect pushes average prices in an ambiguous direction. The conduct effect should occur relatively quickly; firms internalize that in the future their rivals will not be able to detect deviation, which by backwards induction makes them want to deviate as soon as possible. Conversely, the information effect may take more time to realize, as the information that firms have from just before Agri Stats stopped publishing its reports may still be predictive of the economic environment at first but stales over time.

4.1 Aggregate effects

Figure 4 shows the timeseries of average turkey prices (panel (a)) and total sales quantity (panel (b)) for Agri Stats clients in orange, and other producers in blue. The vertical dashed line is May 2019, the month in which Agri Stats stopped producing its turkey reports. While Agri Stats was still publishing, its subscribers' turkey was roughly 20 cents (5 percent) more expensive than non-Agri Stats clients' turkey on average. After the exchange ended, Agri Stats clients' turkey prices fell to around the level of their competitors within several months. Prices of both groups evolved relatively similarly for the next 18 months, including during




Figure 5: Event study results

Note: Regression is weighted by total within-UPC-region quantity throughout the time period. The shaded gray region is the 95% confidence interval, with standard errors clustered by UPC-region. The vertical dashed line represents May 2019.

a price spike at the beginning of the COVID-19 pandemic. Finally, towards the end of the period, Agri Stats' former clients' prices began rising relative to other producers' prices again. Meanwhile, the quantity sold by Agri Stats clients slightly rose after the end of the information exchange, both in absolute terms and relative to others, suggesting that this initial relative price decline was not solely due to relative demand trends.

To quantify this change in relative prices and verify that it holds within-product and region, I estimate an event study. For UPC j in store s , region $r(s)$, and month t , the estimating equation is

$$\log(p_{jst}) = \alpha_{jr} + \gamma_t + \beta_t \text{AS}_j + \varepsilon_{jst}, \quad (3)$$

where p_{jst} is UPC j 's price in store-month st , α_{jr} is a UPC-region fixed effect, γ_t is a month fixed effect, and AS_j is an indicator for whether UPC j is produced by an Agri Stats subscriber. β_t , the coefficient of interest, represents the average price difference between Agri Stats clients and other producers in month t ; I normalize its value in May 2019 to zero. Note that while prices of non-Agri Stats clients are used as a control to account for common demand and cost shocks, this exercise is fundamentally descriptive; one would expect an equilibrium effect on non-Agri Stats clients' prices from a price change for Agri Stats clients, inhibiting any causal interpretation.²¹

Figure 5 shows the results. The black line tracks the evolution of β_t over time, while the shaded gray region is the 95% confidence interval based on UPC-region clustered standard

²¹If prices are strategic complements, then this force in particular would cause this regression to underestimate the price effects of Agri Stats.

	(1)	(2)	(3)	(4)	(5)
$AS_j \times Post_t$	-0.051*** (0.006)	-0.034*** (0.006)	-0.052*** (0.007)	4.96×10^{-6} (0.021)	4.96×10^{-6} (0.021)
$AS_j \times Post_t \times t$			0.012*** (0.004)		
$AS_j \times Post_t \times Turkey_j$					-0.034 (0.021)
UPC-region FEs	✓	✓	✓	✓	✓
Month FEs	✓	✓	✓	✓	
Month-meat type FEs					✓
Observations	8,926,344	15,388,358	15,388,358	4,307,021	19,695,379
Meat	Turkey	Turkey	Turkey	Chicken	Both
Post-period	1 year	3 years	3 years	3 years	3 years

Table 2: Difference-in-differences results

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered by UPC-region are in parentheses. Regressions are weighted by total within-region quantity throughout the time period. AS_j is an indicator for whether UPC j is produced by an Agri Stats client. $Post_t$ is a post-May 2019 indicator. t is a linear time trend variable, denominated in years. $Turkey_j$ is an indicator for whether UPC j is a turkey product.

errors. Overall, the plot demonstrates similar patterns to the aggregate price timeseries. Agri Stats clients' prices followed roughly parallel trends to their competitors for the two years leading up to May 2019. Then, within several months of May 2019, Agri Stats' clients' prices fell by about 5% relative to their rivals, remaining depressed for about eighteen months after May 2019 before slowly climbing back up to slightly below its level from before May 2019. This result independently replicates a legal finding that Agri Stats' turkey clients' prices decreased upon the end of the information exchange.²²

I further quantify the effect of the information exchange by estimating two-by-two difference-in-differences specifications, in which I replace the time-varying β_t in (3) with one β multiplied by a post-May 2019 indicator. Results are in Table 2, for both short-run (one year) and long-run (three year) time horizons after the end of the information exchange. Columns (1) and (2) show that after the end of the exchange, Agri Stats clients' relative prices fell by 5.1% and 3.4% respectively. In column (3), I add an interaction term between the treatment variable $AS_j \times Post_t$ and a linear time trend to separately estimate an immediate price drop followed by growth over time, and find that Agri Stats prices immediately fell by 5.2% but then rose by 1.2% per year in the three years thereafter.


²²Plaintiffs' Opposition to Agri Stats' Motion to Exclude and Strike Expert Testimony. The magnitude of the decline is redacted from the public version of this filing.

I next check whether similar trends occur in chicken to see whether broader trends in poultry could explain these results. Chicken acts as a useful placebo because there is substantial overlap between the markets—90% of turkey producers and 85% of Agri Stats turkey producers by volume also produce chicken—and because the chicken information exchange continued throughout the analysis timeframe. I plot raw chicken prices and quantities in Figure A.2; the price of Agri Stats clients relative to others remains flat around the end of the turkey exchange. Column (4) of Table 2 runs the same specification as column (2) on chicken and finds a precisely estimated null, while (5) runs the full triple-differences specification, finding an estimated relative effect between turkey and chicken that is the same as the difference-in-differences specification for turkey, although introducing the third level of differences adds imprecision to the estimate. I also estimate the triple-differences counterpart to the event study depicted in (3), where the time-varying coefficients of interest represent the turkey-chicken difference in the Agri Stats-other producer difference over time (Figure A.3), and find similar overall patterns to Figure 5: Agri Stats clients’ relative prices drop in turkey relative to chicken after May 2019 but slowly increase thereafter.

These results paint a consistent picture of the aggregate effects of Agri Stats’ exit on turkey prices. When the exchange ended, relative prices of participants immediately fell, but eventually increased back to roughly the with-information sharing level. While the rebound could be due to a procompetitive information effect that takes a relatively long time to realize, as hypothesized at the beginning of the section, it could also be due to secular demand trends (as the contemporaneous relative increase in Agri Stats clients’ quantity late in the period in Figure 4(b) might suggest), foregone efficiencies causing Agri Stats clients to become relatively less efficient without the reports, or unrelated supply-side factors. The goal of the rest of this section is to demonstrate that additional patterns in the data are consistent with both conduct and information effects playing a role here.

4.2 Suggestive evidence of conduct effects

Tacit collusion is easier to sustain when there are fewer competitors, and when those competitors are more symmetric (Ivaldi et al., 2003; Garrod and Olczak, 2018). As such, if the information exchange had a conduct effect, one might expect it to be larger in markets where these “collusion plus factors” are more prevalent. To investigate, I first compute market-level short-run Agri Stats effects by regressing $\log(p_{jrt})$ on a UPC fixed effect, a month fixed effect, and a post-May 2019 indicator interacted with an Agri Stats indicator separately by market (unique combination of region and module), using only data from June 2017 to May 2020. Second, I aggregate the interaction coefficient across markets based on the number of Agri



(a) Effects by number of Agri Stats clients (b) Effects by symmetry of Agri Stats clients

Figure 6: Heterogeneity in local Agri Stats effects by collusion plus factors

Note: Panel (b) only includes markets with two or three Agri Stats clients as of May 2019. Market-by-market regressions are weighted by total within-UPC-region quantity, and standard errors are clustered by UPC. In panel (a), estimates are averaged, weighted by total within-market quantity. The crossbar represents the 95% confidence interval. In panel (b), each bubble is a market, and its size is proportional to total quantity in that market. The black line and shaded region are the predictions and 95% confidence intervals from a weighted regression of local short-run Agri Stats effect on asymmetry. Both panels feature second-step standard errors computed using the Murphy-Topel correction.

Stats clients in the market as of May 2019 and check whether markets with fewer Agri Stats clients had more negative short-run Agri Stats effects. Third, I define an asymmetry measure for each market as the largest within-market difference in two Agri Stats clients' shares, and assess its correlation with the short-run Agri Stats effect. If there were conduct effects that were predicted by collusion plus-factors, one would expect to see a positive correlation: more asymmetric markets will have larger (less negative) Agri Stats effects.

Figure 6 shows the results. Panel (a) finds that the short-run Agri Stats effect was largest in markets with two or three Agri Stats clients: it was 8.4% and 5.0% on average in markets with two or three subscribers, but 1.1% in all other markets. Panel (b) scatters market-level symmetry (which increases as points move to the left) against the short-run Agri Stats effect and finds a clear correlation: moving from the 25th to 75th percentile of symmetry results in a 2.6 percentage point more negative short-run Agri Stats effect ($SE = 1.1pp$). Both of these results point towards the end of the information exchange causing conduct effects, underscoring more direct evidence from business documents, presented in legal cases, of firms using the reports to raise prices.

4.3 Suggestive evidence of information effects

In this section, I test two predictions of the model of information effects presented in Section 2. First, Agri Stats clients should respond more to changes to demand when sharing information. In particular, if Agri Stats clients learn about a common-value demand shock,

then both direct and equilibrium effects should cause Agri Stats clients' prices to get more correlated with that shock (Figure A.1(b)). Second, sharing price information should cause correlation in Agri Stats clients' prices: prices are strategic complements, so learning a firm's price causes rivals' prices to move in the same direction.

I begin by estimating constant-elasticity total demand curves for each module:

$$\log(Q_{rmt}) = \gamma_{rm} + \delta_{\text{month}(t)} + \alpha_m \log(P_{rmt}) + \xi_{rmt} \quad (4)$$

where Q_{rmt} is total quantity in region-module-month rmt , γ_{rm} and $\delta_{\text{month}(t)}$ are region-module and month-of-year fixed effects, and P_{rmt} is a fixed-weight price index instrumented with poultry feed prices, an input cost.²³ I recover elasticities of -2.12 in fresh turkey and -7.03 in deli turkey, with first-stage F -statistics of 276 and 1,109 respectively. ξ_{rmt} represents a demand residual common to all producers in the market, some portion of which may be unknown without information sharing.

With the estimated residuals $\hat{\xi}_{rmt}$ in hand, I first test whether prices in markets with a greater Agri Stats presence tracked common demand shocks more while the information exchange was ongoing. Letting AS_{rm} be the share of production by Agri Stats clients in region-module rm as of May 2019, I estimate

$$\begin{aligned} \log(P_{rmt}) = & \alpha_{rm} + \gamma_t + \beta \hat{\xi}_{rmt} + \zeta AS_{rm} \text{Post}_t + \\ & \theta \hat{\xi}_{rmt} \text{Post}_t + \eta \hat{\xi}_{rmt} AS_{rm} + \mu \hat{\xi}_{rmt} AS_{rm} \text{Post}_t + \varepsilon_{rmt}. \end{aligned} \quad (5)$$

I standardize $\hat{\xi}_{rmt}$ by its within-module standard deviation, such that β can be interpreted as the increase in log prices associated with a 1 SD increase in the demand residual. If the exchange had information effects, we would expect $\eta > 0$ (before May 2019, prices in Agri Stats-heavy markets correlate more with the demand residual) and $\mu < 0$ (Agri Stats-heavy markets became less responsive to the demand residual after May 2019).

The results in Table 3 indicate that these hypotheses hold. The first three rows show that prices are correlated with the demand residual, post-May 2019 prices decreased by more in Agri Stats-heavy markets (although imprecisely so at this level), and non-Agri Stats clients' prices got slightly more correlated with demand shocks after the end of the exchange. $\hat{\eta}$ (fourth row) is positive and significant: in the full post-period time horizon, a 10 percentage point increase in the Agri Stats share is associated with a 0.004 log-point increase in the demand residual-price association, which is 10.3% of the baseline association of 0.039. Finally, $\hat{\mu}$ (fifth row) roughly equals $-\hat{\eta}$, suggesting that Agri Stats clients' and

²³Prices at the UPC-store-month level are aggregated to the region-module-month-level index by taking a weighted average, using total UPC-store-level quantity sold over the entire time period as weights.

	(1)	(2)
$\hat{\xi}_{rmt}$ (demand residual)	0.042*** (0.014)	0.039*** (0.015)
$AS_{rm} \times Post_t$	-0.024* (0.014)	-0.022* (0.013)
$\hat{\xi}_{rmt} \times Post_t$	0.023* (0.012)	0.021* (0.012)
$\hat{\xi}_{rmt} \times AS_{rm}$	0.031*** (0.009)	0.040*** (0.011)
$\hat{\xi}_{rmt} \times AS_{rm} \times Post_t$	-0.044*** (0.015)	-0.041*** (0.015)
Region-module FEs	✓	✓
Month FEs	✓	✓
Observations	14,486	24,152
Post-period	1 year	3 years

Table 3: Responsiveness of Agri Stats clients' prices to demand

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Observations are weighted by total region-module level quantity over the sample. Standard errors clustered by region-module, corrected for first-step uncertainty in $\hat{\xi}_{rmt}$ using the delta method, in parentheses.

rivals' prices were not differentially correlated with demand after May 2019. The pattern is similar in one-year and three-year time horizons, suggesting that this effect occurred quickly after the information exchange ended.

Next, I test whether Agri Stats clients' prices became less correlated with each other after the end of the information exchange. To do so, I estimate the following regression separately before and after May 2019, at the firm-store-month level:

$$\log(p_{fst}) = \alpha_{fs} + \gamma_t + \sum_{G \in \{AS, AS^C\}} \beta_G \mathbb{1}\{f \in G\} \log(\bar{p}_{AS \setminus f, st}) + \varepsilon_{fst} \quad (6)$$

where α_{fs} and γ_t are firm-store and month fixed effects, AS and AS^C are the set of Agri Stats subscribers and non-Agri Stats subscribers, and $\bar{p}_{AS \setminus f, st}$ is the average price of all Agri Stats products in the same store and module, excluding firm f (which removes the mechanical effect of firm f 's prices on the average). The coefficients of interest, β_{AS} and β_{AS^C} , represent the average change in log prices for Agri Stats and non-Agri Stats firms associated with a 1 log-point change in the average price of other Agri Stats firms at the same store. I also reestimate this specification while controlling for the common-value region-month demand residual $\hat{\xi}_{rmt}$ from (4), which isolates correlation due to information about this demand residual from observation of rivals' prices themselves.

The results, shown in Table 4, suggest that Agri Stats clients' prices become significantly less correlated after the exchange ended. Before May 2019, a 1% increase in the average Agri Stats price in the same store is associated with a 12pp higher increase in an Agri Stats

	(1)	(2)	(3)	(4)
β_{AS} (Agri Stats clients)	0.667*** (0.097)	0.613*** (0.100)	0.328*** (0.062)	0.329*** (0.063)
β_{ASC} (other producers)	0.548*** (0.053)	0.481*** (0.054)	0.531*** (0.045)	0.531*** (0.045)
$\hat{\xi}_{rmt}$ (demand residual)		0.081*** (0.015)		-0.002 (0.006)
Firm-store FEs	✓	✓	✓	✓
Month FEs	✓	✓	✓	✓
Period	Pre	Pre	Post	Post

Table 4: Correlation of prices with same-store average Agri Stats price

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Only includes store-module-months with at least two Agri Stats firms. Standard errors clustered by region-module, corrected for first-step uncertainty in $\hat{\xi}_{rmt}$ using the delta method, in parentheses. Average leave-out prices are weighted by total quantity throughout the time period.

client’s price than another producer’s price, but a roughly 20pp *lower* increase after May 2019. Some of the pre-period effect is due to correlated information about the demand state—conditioning on $\hat{\xi}_{rmt}$ reduces β_{AS}^{pre} by 5pp—but the gap between β_{AS}^{pre} and β_{ASC}^{pre} grows even wider, while not affecting the correlations post-May 2019.

Overall, firms seem to have used the Agri Stats reports to learn about their economic environment in a manner consistent with theoretical models of information effects. When combined with the evidence of conduct effects, there appear to be multiple forces affecting overall outcomes, which reduced-form methods are not equipped to separate. Further, given that information effects seemed to occur soon after May 2019, we do not know why Agri Stats prices rebounded slowly after May 2019, and whether this pattern can be explained by relative demand or supply trends. Finally, as discussed in Section 2, average prices are not sufficient statistics for the welfare effects of information sharing. These factors motivate the development of a structural model to estimate how firms’ information sets changed with the end of the exchange, test and quantify changes in conduct, assess welfare effects, and see how outcomes might change under counterfactual exchanges.

5 Model

Testing whether Agri Stats clients used the information exchange to collude requires a model of competition that accounts for information effects. However, leading conduct testing methods (Backus et al., 2021) typically employ complete-information “BLP-style” supply-side

models, hence will mistake information effects for conduct effects, resulting in bias. Thus, answering the question at hand requires incorporating incomplete information into the standard supply-side model in a manner that captures the effects of information sharing.

In this section, I lay out such a model. In it, firms have incomplete information about demand shocks and rivals' cost shocks, and potentially imperfect information about rivals' prices and quantities. The latter is captured in a general way by *signal functions* of the vector of prices and quantities that each firm observes after each period. When firms share information, they observe more precise signals. Firms use these signals to form beliefs about demand and their rivals' prices in each period, updating their beliefs from period to period based on what they observe. In a competitive Bertrand-Markov equilibrium, firms set prices that maximize their expected profits in each period given beliefs that are consistent with the signals they receive and the long-run autocorrelation of rivals' prices. I impose this solution concept in the period after firms stop sharing information, allowing me to test whether it holds while they were sharing information while accounting for information effects through changes to firms' signal functions.

Throughout, I use j to index products, f firms, r regions, and t months. I write a variable without certain subscripts to mean a vector stacking that variable over all omitted subscripts, while I add a minus sign to denote all entries except the one(s) corresponding to that subscript. For example, if x_{jrt} is a product-region-month-specific variable, then x_{rt} is the vector stacking values of those variables across products and within that region-month, while $x_{(-f)rt}$ is that vector excluding entries of products owned by firm f .

5.1 Supply

I begin with the supply side, which is novel; the demand side, which is described next, is a standard mixed logit (Berry et al., 1995). Firms $1, \dots, F$ compete in regions $1, \dots, R$ for an infinite horizon of time periods $1, 2, \dots$. I refer to a unique combination of a region and time period as a *market*; for the remainder of this subsection, I condition on a particular region and omit r subscripts for brevity. Each market is characterized by an exogenous tuple $(\mathcal{J}_t, \theta_t, x_t, w_t, \xi_t, \omega_t)$, where $\mathcal{J}_t \in 2^{\mathcal{J}}$ is the set of available products (and $\mathcal{J}$ is the set of all products ever sold), $\theta_t \in \Theta$ is a vector of consumer characteristics that includes the market size I_t , $x_t \in \mathcal{X}$ is a vector of observable product characteristics for each product, $w_t \in \mathcal{W}$ is a vector of observable cost shifters for each product, $\xi_t \in \Xi$ is a scalar demand shock for each product, and $\omega_t \in \prod_{j \in \mathcal{J}} \Omega_j = \Omega$ is a scalar mean-zero cost shock for each product. Product j in market t has market share $s_j(p_t; \mathcal{J}_t, \theta_t, x_t, \xi_t)$, where $p_{jt} \in \mathcal{P}_j$ is product j 's price in market t , and constant marginal cost $c_{jt} = h(w_{jt}) + \omega_{jt}$. For brevity, I write $s_{jt}(p_t; \xi_t)$ as

the demand for product j in market rt , fixing $(\mathcal{J}_{rt}, \theta_{rt}, x_{rt})$. Firm $f \in \mathcal{F}$ owns products $\mathcal{J}_f$, where $\bigcup_{f \in \mathcal{F}} \mathcal{J}_f = \mathcal{J}$; firm f 's product assortment in market t is $\mathcal{J}_{ft} = \mathcal{J}_f \cap \mathcal{J}_t$.²⁴ I assume that $\{\mathcal{J}_t, \theta_t, x_t, w_t, \xi_t, \omega_t\}_{t \geq 1}$ are generated by a stationary ergodic process.

Given a vector of prices p_t , demand shocks ξ_t , and cost shocks for firm f 's products $\omega_{ft} = \{\omega_{jt}\}_{j \in \mathcal{J}_{ft}}$, firm f 's profit in market t is

$$\pi_{ft}(p_t; \xi_t, \omega_{ft}) = I_t \sum_{j \in \mathcal{J}_{ft}} s_{jt}(p_t; \xi_t)(p_{jt} - h(w_{jt}) - \omega_{jt}). \quad (7)$$

Information and timing. The following steps take place in each period, which I discuss in turn below.

1. Primitives $(\mathcal{J}_t, \theta_t, x_t, w_t, \xi_t, \omega_t)$ are drawn. All primitives except (ξ_t, ω_t) are commonly known to all firms; ξ_{jt} is unobserved to all firms, while ω_{jt} is observed by firm f if $j \in \mathcal{J}_{ft}$.
2. Firm f observes cost and inventory information $(\omega_{f't}, \iota_{f't})$ for a subset of firms $f' \in \mathcal{F}_f^{\text{obs}} \subseteq \mathcal{F} \setminus \{f\}$, where the inventory information $\{\iota_{f't}\}_{f' \in \mathcal{F}_f^{\text{obs}}}$ is sufficient for firm f to predict the prices of firms in $\mathcal{F}_f^{\text{obs}}$.
3. All firms simultaneously set prices p_t ; shares s_t are realized.
4. Firms observe a signal of prices and shares, where firm f 's signal is $y_{ft} = y_f(p_t, s_t)$.

Step 1 states the fundamental uncertainty firms face: the vector of demand shocks ξ_t and rival cost shocks $\omega_{(-f)t}$. These, particularly ξ , traditionally correspond to characteristics that are unobserved to the economist but known to market participants, such as style or innate product quality (Berry, 1994). It may therefore seem odd to assume that firms do not observe them. However, as I explain below, firms learn about these objects through the information they observe about the market, and thus still have more information about them than the economist. Therefore, these variables' typical interpretation is still valid. Furthermore, the existence of Agri Stats implies there is economically significant information about the market that firms do not naturally observe.

Step 2 allows firms to share information about costs and the prices they are about to set with a subset of their rivals. Sharing advance notice of prices is a feature of many real-world information exchanges (Borenstein, 2004; Byrne et al., 2025), and the Agri Stats case record suggests firms interpreted inventory information in this way. I do not have inventory data, and directly modeling the inventory-setting process would introduce substantial complexity. I therefore take a reduced-form approach and assume that sharing inventories is equivalent

²⁴The indices and dimensions of product-level vectors such as ξ_t can change due to variation in product assortment. To avoid notational issues, I index these vectors by the full universe of products $\mathcal{J}$, with dummy values for $j \notin \mathcal{J}_t$ that are ignored by all operators.

to sharing advance notice of prices; I assess robustness to this modeling decision below.

Step 3 is a standard price-setting assumption, but there is nothing special about prices; the model extends to any setting with one flexible decision variable and where firms share information about both their decisions and their outcomes. Step 4 encodes sharing general information about past prices and shares. For example, a firm only observing its own prices and share has $y_f(p_t, s_t) = (p_{ft}, s_{ft})$, while observing the full vector of prices and shares corresponds to the identity function $y_f(p_t, s_t) = (p_t, s_t)$. Aggregates, such as average prices or total quantities, are also supported.²⁵ I consider deterministic signal functions here, but the model could accommodate stochastic signals as well.

Beliefs. The payoff-relevant uncertainty firm f faces in market rt is summarized by its beliefs, a probability distribution ρ_{ft} over $(\xi_t, p_{(-f)t})$. Notice that firms do not directly form beliefs over $\omega_{(-f)t}$, as they do not enter into firm f 's profits conditional on $p_{(-f)t}$. In each period, firms have *prior beliefs* (i.e., their beliefs as they are setting prices) and *posterior beliefs* (i.e., their beliefs after they set prices and observe the signal of prices and shares). Sharing information about past prices and shares affects posterior beliefs by giving firms a more precise signal of what occurred in the market, which affects their prior beliefs for the following period to the extent that ξ_t and $p_{(-f)t}$ are autocorrelated. Sharing cost and inventory information directly tightens firms' prior beliefs.

To be precise about this mechanism, suppose firm f has prior belief ρ_{frt} in market rt . If firm f observes inventory information from some of its rivals, then the prior belief over these rivals' prices will be degenerate. Suppose that firm f sets prices p_{ft} and receives signal $y_{ft} = y_f(p_t, s_t)$. It uses Bayes' rule to form its posterior as the regular conditional distribution $\rho_{ft}(\cdot | y_{ft})$, with a more precise signal tending to result in a more precise posterior. Firms use their posteriors to form their priors for the next period. For a more precise posterior to imply a more precise prior, demand and prices must be autocorrelated; otherwise, period t outcomes would be uninformative of period $t + 1$. I therefore place slightly more structure on ξ_t than is typical, by assuming that ξ_t follows a stationary first-order Markov process with a law of motion $F^\xi(\cdot | \xi_t)$. This nests as special cases common value and independent private value demand shocks, the prevailing assumptions in the information effects literature (Raith, 1996).

While ξ is a structural primitive, p is endogenous; its evolution over time is *ex ante* unclear and can be influenced by the structure of how firms' beliefs themselves evolve. For example, if firm 1 conjectures that firm 2's price depends on what happens more than one period ago, firm 1's price may then optimally depend on its beliefs about what happened more

²⁵Because quantities are just shares multiplied by market size and market size is commonly known, I use "quantities" and "shares" interchangeably here and below.

than one period ago, which then self-confirms the conjectured history-dependence. Without any restrictions, firms could in theory need to condition on the entire signal history as state variables, creating a curse of dimensionality on firms' state spaces.

I therefore assume that each firm f conjectures that $p_{(-f)t}$ follows a stationary conditional first-order Markov process, with kernel $F_f^p(\cdot \mid p_t, \xi_t, x_{t+1}, w_{t+1})$. This assumption essentially says that if firms knew all the prices and demand state last period, along with the demand and cost observables this period, then there would be no informational benefit of knowing prices more than one period ago in predicting the prices rivals set today.²⁶ Despite rival firms not knowing it, the transition kernel depends on ξ_t and p_t because they affect s_t , which enters $y_{(-f)t}$, which then affects $\rho_{(-f)t}(\cdot \mid y_{(-f)t})$, and hence $\rho_{(-f)(t+1)}(\cdot)$ and $p_{(-f)(t+1)}$; the distribution of period- t prices also depends on x_t and w_t because they are commonly known and shift demand and costs respectively.

Given these assumptions, firm f views the vector $(\xi_t, p_{(-f)t})$ as following a first-order Markov process, with law of motion

$$F_f(\cdot \mid \xi_t, p_t, x_{t+1}, w_{t+1}) = F^\xi(\cdot \mid \xi_t) F_f^p(\cdot \mid \xi_t, p_t, x_{t+1}, w_{t+1}). \quad (8)$$

Then, given a posterior belief $\rho_{ft}(\cdot \mid y_{ft})$, the prior belief for the following period is

$$\rho_{f(t+1)}(\cdot) = \int F_f(\cdot \mid \xi_t, p_t, x_{t+1}, w_{t+1}) d\rho_{ft}(\xi_t, p_{(-f)t} \mid y_{ft}). \quad (9)$$

The conjectured law of motion F_f^p and signal function y_f therefore pin down how firm f 's beliefs update from period to period.

Pricing rules and long-run consistency. To close the model, we must specify (a) firms' objectives and (b) an equilibrium condition pinning down firms' conjectured laws of motion. Let $\Lambda = 2^{\mathcal{J}} \times \Theta \times \mathcal{X} \times \mathcal{W}$ be the space of all commonly known exogenous variables, $\Omega_f^{\text{obs}} = \prod_{f' \in \{f\} \cup \mathcal{F}_f^{\text{obs}}} \prod_{j \in \mathcal{J}_{f'}} \Omega_j$ be the space of all cost shocks observed by firm f , $\mathcal{P}_f = \prod_{j \in \mathcal{J}_f} \mathcal{P}_j$ be the space of firm f 's prices, and $\mathcal{B}_f = \Delta \left(\Xi \times \prod_{f' \neq f} \mathcal{P}_{f'} \right)$ be the space of firm f 's beliefs. A *pricing rule* for firm f is a function $\sigma_f : \Lambda \times \Omega_f^{\text{obs}} \times \mathcal{B}_f \rightarrow \mathcal{P}_f$ that maps a set of observable exogenous variables, f -observed cost shocks, and beliefs to firm f 's prices.

The pricing rule summarizes the outcome of firm f 's period-by-period optimization prob-

²⁶This bears similarity to the central assumption in Markov perfect equilibrium (Ericson and Pakes, 1995; Maskin and Tirole, 2001), wherein agents conjecture that their rivals' actions combine to form a Markov process. There, the dependence over time in firms' actions is due to dynamics in firms' problems; here, firms face a static problem, but there is time-dependence coming from incomplete information and autocorrelation in the state over time.

lem. One special example is expected own-profit maximization. This rule satisfies

$$\sigma_f(\cdot) \in \arg \max_{\tilde{p}_f} \int \pi_{ft}(\tilde{p}_f, p_{(-f)t}; \xi_t, \omega_{ft}) d\rho_{ft}(\xi_t, p_{(-f)t}). \quad (10)$$

When firms share cost information with some rivals, their pricing rule can be measurable with respect to the cost shocks of those rivals, allowing firms to therefore directly incorporate rivals' expected profits into their objectives as well. While these static pricing rules do not accommodate directly modeling dynamic incentives, this allows a scope for collusive behavior. I return to this issue in Section 8.²⁷

To move towards an equilibrium concept, we must ensure that firms' conjectured price laws match reality. Fix a profile of pricing rules, signal functions, period-1 beliefs, and conjectured price kernels $\{\sigma_f, y_f, \rho_{f1}, F_f^p\}_{f \in \mathcal{F}}$. These, along with Bayes updates of beliefs in each period, induce an augmented process over $(\mathcal{J}_t, \theta_t, x_t, w_t, \xi_t, \omega_t, p_t, \rho_t)$. I call an initial belief vector $\{\rho_{f1}\}_{f \in \mathcal{F}}$ *regular* with respect to $\{\sigma_f, y_f, F_f^p\}_{f \in \mathcal{F}}$ if this process converges to a unique invariant distribution Π . Intuitively, this means that the impact of the initial distribution disappears over time.

Regularity implies that the notion of a long-run distribution of realized prices and primitives is well-defined. Let $\Pi^+(\cdot \mid p_t, \xi_t, x_{t+1}, w_{t+1})$ be the conditional distribution of p_{t+1} induced by the stationary joint distribution of consecutive observations under Π . Firm f 's conjectured price kernel is *long-run consistent* if, for every Borel set $A \subseteq \prod_{f' \neq f} \mathcal{P}_{f'}$ and Π -almost every conditioning state $(p_t, \xi_t, x_{t+1}, w_{t+1})$,

$$F_f^p(A \mid p_t, \xi_t, x_{t+1}, w_{t+1}) = \Pi^+(\mathcal{P}_f \times A \mid p_t, \xi_t, x_{t+1}, w_{t+1}). \quad (11)$$

In other words, the conjectured pricing law of motion must equal the actual long-run law of motion of prices that is induced by the pricing rules and those conjectured laws of motion.

Equilibrium definition. Given a set of initial beliefs $\rho_{f1}(\cdot)$, signal functions $y_f(\cdot)$ and pricing rules σ_f for each firm, and a data-generating process for primitives, a *Markov equilibrium in pricing rules* consists of a vector of prices p_t for each market, prior and posterior beliefs $\rho_{ft}(\cdot)$ and $\rho_{ft}(\cdot \mid y_{ft})$ for each firm in each market, and a conditional first-order Markov kernel $F_f^p(\cdot \mid p_t, \xi_t, x_{t+1}, w_{t+1})$ for each firm such that the initial beliefs are regular with respect to the kernel and the following hold:

1. *σ -optimality of prices.* Given prior beliefs ρ_{ft} , $p_{ft} = \sigma_f(\cdot)$ for all ft .

²⁷Focusing on static pricing rules precludes signal-jamming behavior in which firms set statically suboptimal prices in order to influence their rivals' beliefs in a manner beneficial to the firm's future profits. Capturing this dynamic signaling incentive would add significant complexity, so I focus on static expected profit maximization as the main competitive benchmark.

2. *Updating of beliefs.* For all ft , given signals $y_{ft} = y_f(p_t, s_t(p_t; \xi_t))$, $\rho_{ft}(\cdot \mid y_{ft})$ is derived from $\rho_{ft}(\cdot)$ using Bayes' rule, and $\rho_{f(t+1)}(\cdot)$ is derived from $\rho_{ft}(\cdot \mid y_{ft})$ using (9), with F_{fr} derived from F_r^ε and F_f^p as in (8).
3. F_f^p satisfies long-run consistency (11) for all f .

A special case is *Bertrand-Markov equilibrium*: a Markov equilibrium in pricing rules with the own-profit maximization pricing rule, (10).

5.2 Demand

As previously mentioned, I parameterize demand using a mixed logit model with linear utility in the spirit of [Berry et al. \(1995\)](#) for my empirical application. This is the standard demand framework in consumer packaged goods markets ([Nevo, 2001](#); [Miller and Weinberg, 2017](#); [Backus et al., 2021](#); [Bhattacharya et al., 2026](#)). Reintroduce the region subscript r . For a mass of consumers $i \in [0, I_{rt}]$ with unit demand over options in $\mathcal{J}_{rt}$ plus an outside option, which is purchasing a different meat, the utility enjoyed by consumer i from purchasing product j is

$$u_{ijrt} = \alpha_i p_{jrt} + x'_{jrt} \beta_i + \gamma_{\text{brand}(j)} + \gamma_r + \gamma_{\text{month}(t)} + \xi_{jrt} + \varepsilon_{ijrt} \quad (12)$$

where $\gamma_{\text{brand}(j)}$, γ_r , and $\gamma_{\text{month}(t)}$ are brand, region, and month-of-year fixed effects, and ε_{ijrt} is drawn independently and identically from a type-I extreme value distribution. θ_{rt} —the set of commonly observed consumer characteristics—includes these fixed effects and the distribution of (α_i, β_i) in market rt . The product characteristics are indicators for white meat, as well as an indicator for whole-smoked turkey in fresh turkey only. The outside option's utility is normalized to $u_{i0rt} = \varepsilon_{i0rt}$. (12) yields individual choice probabilities s_{ijrt} following the typical logit functional form, and integrating over (α_i, β_i) yields the demand functions s_{jrt} .

5.3 Application to Agri Stats

Before continuing, it is useful to describe how I will apply the model to test whether Agri Stats facilitated collusion. While I assume a stationary distribution of primitives throughout the sample period, I allow the periods before and after May 2019 correspond to two separate Markov equilibria in pricing rules. As is standard in the conduct testing literature, I assume that firms interacted competitively (here, according to a Bertrand-Markov equilibrium) while not sharing information ([Miller and Weinberg, 2017](#); [Calder-Wang and Kim, 2025](#)). The assumption of a Markov equilibrium in pricing rules before May 2019 allows me to recover

and propagate firms' beliefs from the realized price and share history without specifying the objective that generated these prices. Firms therefore enter June 2019 with beliefs over ξ that are carried over from the information sharing period, capturing the potential timing mismatch between conduct and information effects discussed in Section 4.

I also leverage the institutional details reviewed in Section 3 to specify signal functions with and without information sharing. I assume that Agri Stats shared the entire vector of past prices and quantities, as well as costs and inventories, with its clients: that is, for all $f \in AS$, $\mathcal{F}_f^{\text{obs}} = \{f' : f' \in AS \setminus \{f\}\}$ and $y_f(p_{rt}, s_{rt}) = \{(p_{f'rt}, s_{f'rt}) : f' \in AS\}$. Meanwhile, due to the relative lack of alternative data sources in poultry markets, I assume that all firms that are not sharing information do not observe costs or inventories for any other competitors, and only observe their own prices and quantities: $\mathcal{F}_f^{\text{obs}} = \emptyset$ and $y_f(p_{rt}, s_{rt}) = (p_{frt}, s_{frt})$. This applies both to nonsubscribing firms while the exchange is ongoing and to all firms after the exchange ends.

While it may seem stark to assume that firms do not directly observe rivals' prices or shares when not sharing information, or that firms use inventory information to perfectly predict their rivals' prices, I view these as *ex ante* conservative from a conduct testing perspective, as they allow for the largest information effect possible to potentially offset the reduced-form Agri Stats price effect. Nevertheless, I perform several robustness checks, described in Appendix C, to assess the sensitivity of the forthcoming conduct tests to these assumptions. First, I assume that firms not sharing information observe either average prices and total quantity in each market from government indices, or the full vector of past prices and quantities from retail data. Second, I relax the assumption that inventories enable firms to perfectly conjecture their rivals' prices, instead assuming that inventories give firms an unbiased but noisy signal of their rivals' prices. Third, rather than assuming that firms shift between equilibrium price kernels immediately in May 2019, I allow for a buffer period in which firms continuously transition between the pre-May 2019 and post-May 2019 kernels over a period of up to a year. The results appear robust to these assumptions.

Two additional assumptions warrant caveats. First, the mixed logit demand system may impose some parametric restrictions on marginal profits, hence (per Section 2) information effects. Mixed logit is the workhorse demand system in industrial organization, so characterizing information effects with it is important in itself; in addition, I verified by simulation that information effects are directionally ambiguous with even simple logit models, so the demand system does not assume the result. Regardless, incorporating nonparametric demand estimation procedures (e.g., [Compiani, 2022](#)) may be a fruitful avenue for future study. Second, my data contain retail prices and quantities, but the exchange was among upstream processors. Manufacturers setting retail prices is a standard assumption in studies of con-

sumer packaged goods (e.g., [Nevo, 2001](#); [Miller and Weinberg, 2017](#); [Backus et al., 2021](#)), and there is evidence this assumption fits scanner data well ([Duarte et al., 2024](#)). Moreover, bargaining power heterogeneity played a limited role in the Agri Stats litigation—Agri Stats did not circulate retailer-specific information, and all turkey purchasers were certified as one class for damages assessment—suggesting that retailer-processor bargaining is second-order in understanding the effects of Agri Stats.²⁸ However, I accept as a limitation that my model does not capture this vertical structure.

6 Identification, estimation, and results

6.1 Identification

I begin with identification of demand. When firms have complete information, the typical identification problem is that firms set prices knowing the realizations of the unobserved (to the econometrician) demand shocks ξ , causing price endogeneity. While I assume firms do not directly observe realizations of ξ , this endogeneity problem still arises: firms set prices having beliefs over ξ that are correct in expectation, causing correlation between prices and ξ . Thus, identification still requires instruments z that shift prices while being exogenous to demand shocks ([Berry and Haile, 2014](#)).

To resolve this endogeneity, I use three types of instruments. First, I use input and transportation costs: feed prices and diesel prices multiplied by the distance to the nearest poultry processing plant. I interact the latter with a store-brand indicator, reflecting that retailers may be differentially able to account for fluctuations in diesel prices by adding or reducing inventory at the store. Second, I exploit the quasi-experimental price variation from the end of the information exchange by using a pre-May 2019 indicator interacted with an Agri Stats indicator. This instrument is valid if the relative distribution of demand shocks is mean-independent of the end of the information exchange, while its relevance is demonstrated by the reduced-form results. Third, I use the differentiation instruments of [Gandhi and Houde \(2020\)](#), which shift prices by changing the degree of local competition in characteristics space analogously to traditional “BLP instruments.” For two characteristics x and x' (where x and x' could be the same characteristic), the instrument is

$$z_{jrt}^{GH,x,x'} = \sum_{k \in \mathcal{J}_{rt}} (x_{jrt} - x_{krt})(x'_{jrt} - x'_{krt}).$$

²⁸In addition, Appendix [C.4](#) provides some reduced-form evidence that vertical structure played a small role, if any, in explaining the Agri Stats effect.

I also estimate an exogenous “predicted price” characteristic conditional on all characteristics and instruments $\hat{p}_{jrt} = \mathbb{E}[p_{jrt} \mid x_{jrt}, z_{jrt}]$ using a random forest, and construct differentiation instruments based on $\hat{p}_{jrt}$.

Once demand is identified, the full distribution of (p, ξ, x, w) is identified. Hence, the joint distribution of $(p_{r(t+1)}, p_{rt}, \xi_{r(t+1)}, \xi_{rt}, x_{r(t+1)}, w_{r(t+1)})$ is identified for each region r , meaning the conditional distributions of $p_{r(t+1)} \mid p_{rt}, \xi_{rt}, x_{r(t+1)}, w_{r(t+1)}$ and $\xi_{r(t+1)} \mid \xi_{rt}$ are also identified. The equilibrium long-run consistency condition therefore identifies the equilibrium pricing laws of motion for each firm both before and after the exchange ends, along with the law of motion of ξ_r .

Supply-side identification, given the Bertrand-Markov equilibrium assumption after May 2019, rests on identifying the components of firms’ first-order conditions:

$$c_{jrt} = h(w_{jrt}) + \omega_{jrt} = p_{jrt} + \left(\mathbb{E}_{f_{rt}} [D_{p_f} s_{f_{rt}}(p_{rt}; \xi_{rt})'] \right)^{-1} \mathbb{E}_{f_{rt}} [s_{f_{rt}}(p_{rt}; \xi_{rt})] \quad \forall j \in \mathcal{J}_{f_{rt}}, \quad (13)$$

where $\mathbb{E}_{f_{rt}}[\cdot]$ is the expectation operator with respect to firm f ’s beliefs over $(\xi_{rt}, p_{(-f)rt})$ in market rt , and $D_{p_f} s_{f_{rt}}(\cdot)$ is the Jacobian of firm f ’s demand with respect to its own products’ prices. In the typical complete-information supply side model, the entire right-hand side of (13) is known after demand is identified, which identifies the vector of marginal costs, c_{jrt} . Here, however, we must know firms’ beliefs in order to evaluate the expectations on the right-hand side. This presents a fundamental identification problem: one observed price may be consistent with a firm having either high costs and pessimistic beliefs about their demand, or low costs and optimistic beliefs about their demand.

Nevertheless, I proceed to argue that given a set of initial conditions for firms’ beliefs, subsequent beliefs are indeed identified by the structure of the model, and that therefore costs are identified from the first-order conditions as usual. I argue by induction: if beliefs in period t are identified, then beliefs in period $t + 1$ are also identified. The first step is to identify posterior beliefs given prior beliefs in period t . With demand identified, $s_{rt}(p_{rt}; \xi_{rt})$ is known; hence, given a prior for firm f in market rt , all functions of s_{rt} and its inputs correspond to a known probability distribution. This includes the (known) signal function $y_f(p_{rt}, s_{rt})$. Second, since realized prices and shares come from data, the signal realization $y_{f_{rt}}$ itself is known. Hence, firm f ’s posterior distribution over $(p_{(-f)rt}, \xi_{rt})$ conditional on having received signal $y_{f_{rt}}$ is directly implied by Bayes’ rule, identifying posteriors. Firms then use the equilibrium laws of motion over $(\xi_{rt}, p_{(-f)rt})$ which, as discussed above, are identified because demand is identified, to move from the period- t posterior to the period- $t + 1$ prior. As such, firms’ prior beliefs in period $t + 1$ are identified from the updating rule, equation (9). In addition, notice that the assumption of a Markov equilibrium in pricing

rules allows us to identify beliefs in the exact same way before May 2019; this process did not rely on imposing the Bertrand-Markov pricing rule.

This leaves an initial conditions problem: period-zero beliefs must be specified in order to recover subsequent beliefs. To be internally consistent with a Markov equilibrium in pricing rules, they must be regular—that is, they must eventually converge in expectation to the stationary distribution implied by the estimated Π_r^+ . I take the natural approach of calibrating firms’ beliefs in the beginning of my sample (June 2017) to that stationary distribution itself, which is clearly regular. This is justified by the fact that the information exchange had been ongoing for more than ten years by this point, with a relatively stable membership of firms and minimal entry or exit in the market. Iteratively recovering beliefs in period $t + 1$ from beliefs in period t then identifies beliefs throughout the entire time period. Finally, with beliefs identified, post-May 2019 marginal costs, $h(\cdot)$, and ω are identified from firms’ first-order conditions (13) as usual.

6.2 Estimation

Estimation proceeds in four steps, closely following the identification argument but with additional parametric structure to account for the relatively high dimension of firms’ belief spaces. The estimation procedure is implemented separately for deli and fresh turkey. Additional details are available in Appendix B.2.

I begin with demand estimation, which uses the micro-BLP estimator of [Berry et al. \(2004\)](#). I parameterize the random coefficients $(\alpha_i, \beta_i) = (\alpha, \beta) + \Sigma \nu_i + \Pi d_i$, where Σ is diagonal, $\nu_i \sim \mathcal{N}(0, I)$ is conformable with Σ and independent of d_i , and d_i includes indicators for three household income groups (less than \$50k, between \$50k and \$100k, over \$100k) and whether children live in the household. Letting $\theta = (\alpha, \beta, \Sigma, \Pi)$ be the vector of structural demand parameters to estimate, I form two types of moments: instrument-exogeneity moments of the form $\mathbb{E}[z' \xi(\theta)]$, where $\xi(\theta)$ are the values that make predicted and actual shares match, and micro-moments of the form $\mathbb{E} \left[\sum_j s_{ijrt}(\theta) x_{jrt} \mid d_i = d \right]$ minus their empirical counterparts from the consumer panel data for each characteristic x and demographic d .²⁹ Finally, given the relatively large number of Gandhi-Houde instruments available, I reduce their dimension to principal components explaining 95% of their variance, a strategy also used in [Backus et al. \(2021\)](#). I use `pyblp` for estimation ([Conlon and Gortmaker, 2020; 2025](#)).

I next estimate the law of motion of ξ and firms’ equilibrium updating rules for p . I parameterize these as linear VAR(1) processes with product-region-specific means, common

²⁹To identify demographic interactions with the constant, I use micro-moments of the form $\Pr(d_i = d \mid j > 0)$.

Algorithm 1 Belief estimation

For each region $r = 1, \dots, R$:

1. *Initialize priors* for all firms as $(\xi_{r1}, p_{r1}) \sim \mathcal{N}(\mu_{r1}, \Sigma_{r1})$, where $\mu_{fr1} = (\bar{\xi}_r, \bar{p}_r + g(w_{r1}, z_{r1}^{GH}))$ and Σ_{fr1} is given by the stationary distribution implied by the timeseries estimates.
 2. For each time period $t = 1, \dots, T$ where t^* corresponds to May 2019, and each firm $f \in \mathcal{F}_{rt}$:
 - (a) *Adjust signals for information sharing.* If $t \leq t^*$ and f is an Agri Stats client, then define the signal function $y_f(p_{rt}, s_{rt}) = \{(p_{jrt}, s_{jrt}) : j \in \mathcal{J}_{f'rt}, f' \in AS\}$ and the set of firms sharing future price information $\mathcal{F}_f^{\text{obs}} = (\mathcal{F}_{rt} \cap AS) \setminus \{f\}$; otherwise, define $y_f(p_{rt}, s_{rt}) = \{(p_{jrt}, s_{jrt}) : j \in \mathcal{J}_{frt}\}$ and $\mathcal{F}_f^{\text{obs}} = \emptyset$.
 - (b) *Adjust priors given shared information about future prices.* Change elements of μ_{frt} corresponding to prices of firms in $\mathcal{F}_f^{\text{obs}} \cup \{f\}$ to the realized prices, and set all corresponding rows and columns of Σ_{frt} to zero.
 - (c) *Approximate posteriors* $\mathcal{N}(\mu_{frt}^{\text{post}}, \Sigma_{frt}^{\text{post}})$ given the realized signal using Laplace's method.
 - (d) *Compute next period's prior* $\mathcal{N}(\mu_{fr(t+1)}, \Sigma_{fr(t+1)})$ using the estimated timeseries process for (ξ, p) (14).
-

slopes across regions, and Gaussian innovations; that is,

$$\begin{aligned} \xi_{rt} &= \bar{\xi}_r + B^\xi \xi_{r(t-1)} + e_{rt}^\xi, & e_{rt}^\xi &\stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \Sigma^\xi), \\ p_{rt} &= \bar{p}_r + g(w_{rt}, z_{rt}^{GH}) + B^p(\xi_{r(t-1)}, p_{r(t-1)}) + e_{rt}^p, & e_{rt}^p &\stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \Sigma^p). \end{aligned} \quad (14)$$

Off-diagonal elements of B^ξ , B^p , Σ^ξ , and Σ^p reflect potential timeseries dependence and correlation of firms' demand and prices on one another. I include the Gandhi-Houde instruments in $g(\cdot)$, which nonparametrically shifts prices, as a reduced-form way of including the entire vector of product characteristics shifting each firm's price, and I assume a linear timeseries in ξ_{t-1} and p_{t-1} because it will result in conjugate priors, as I explain below.

To estimate (14), I use the system GMM estimator of [Blundell and Bond \(1998\)](#). This estimator is well-suited for so-called “small- T , large- N ” settings, which applies here because my sample has fewer time periods than regions, and is robust to unbalanced panels caused by the presence of regional products ([Roodman, 2009](#)). I use system GMM directly to estimate the law of motion of ξ using data throughout my sample period; to estimate the price timeseries, I first partial out $g(w_{rt}, z_{rt}^{GH})$ using a random forest before applying the system GMM following [Robinson \(1988\)](#), and I estimate separate timeseries before and after May 2019 to account for the change in equilibria when the information exchange ends.

Third, I estimate beliefs using the procedure outlined in Algorithm 1. I parameterize beliefs as jointly normal over (ξ, p) , starting in $t = 1$ at the stationary distribution implied by the estimate of (14). For each market rt and firm f , I define the signal functions and set of firms sharing forward prices with firm f based on whether the information exchange is ongoing and the firm’s membership in Agri Stats as described in Section 5.3. Then, given firm f ’s prior and realized signal, I approximate firm f ’s posterior using Laplace’s method, a classical Bayesian method for approximating complicated posteriors as Gaussians (Tierney and Kadane, 1986).³⁰ According to this method, the posterior is approximated as a joint normal with mean equal to the maximum *a posteriori* estimate of (ξ_{rt}, p_{rt}) conditional on the signal and covariance equal to the inverse Fisher information of the signal. Finally, given the linear-Gaussian laws of motion, firm f ’s prior in market $r(t + 1)$ is also normally distributed with a closed-form expression. Looping this process over all firms, periods, and regions results in belief estimates for all firms in all markets.

Finally, I back out costs from the first order conditions (13), evaluating expectations of shares and demand Jacobians with respect to each firm’s beliefs using quasi-Monte Carlo integration. I then estimate $h(w_{jrt})$ using a random forest, after residualizing product, region, and month fixed effects. w_{jrt} includes all cost shifters included as demand instruments—feed prices and diesel prices times distance to the nearest processing plant, with the latter interacted with store-brand indicators—although feed prices are absorbed by the month fixed effect. With the cost function estimated, I back out the residuals $\hat{\omega}_{jrt}$. I compute standard errors using a parametric bootstrap of demand parameters combined with a nonparametric bootstrap sampling over regions with replacement to account for asymptotic variance in the panel estimates, following Kapetanios (2008).

6.3 Results

Table 5 shows the estimated demand parameters. Mean price coefficients are negative and significant; the median own-price elasticity is -5.11 and -2.78 for fresh and deli turkey respectively, but there is substantial heterogeneity in products’ price sensitivity.³¹ The median aggregate elasticity, or the percentage change in inside share from a 1% increase in all prices, is around -0.6 for both categories, similar to the elasticity of turkey in the demand model used by the USDA (Maples et al., 2022). Demographic information is important in explaining taste heterogeneity: fresh turkey is most appealing to low-income consumers (echoing a USDA finding that turkey is an inferior good), and high-income consumers are less price-

³⁰This is an approximation because the demand function is nonlinear. It would be exact otherwise.

³¹The distribution of own-price elasticities is plotted in Figure A.4. I get somewhat more price-elastic results using feasible approximations to the Chamberlain (1987) optimal instruments; see Appendix C.3.

Parameter	Fresh turkey				Deli turkey		
	Constant	Price	Breast	WS	Constant	Price	Low fat
β		-2.904 (0.231)	3.574 (0.313)	-5.669 (0.875)		-0.260 (0.018)	0.061 (0.303)
diag(Σ)	3.153 (0.503)	0.617 (0.049)	0.975 (0.351)	0.695 (1.355)	0.466 (0.275)	0.000 (0.092)	0.202 (2.327)
II: Medium income	-2.436 (0.388)	1.031 (0.104)	-3.669 (0.219)	4.778 (1.240)	2.731 (0.178)	-0.278 (0.029)	-1.125 (0.120)
II: High income	-4.827 (0.340)	1.513 (0.099)	-3.578 (0.206)	5.052 (1.232)	1.077 (0.155)	-0.046 (0.025)	-1.041 (0.142)
II: Children	1.637 (0.250)	-0.029 (0.050)	-1.326 (0.112)	0.828 (0.239)	2.280 (0.118)	-0.259 (0.017)	-0.439 (0.049)
Number of observations		57,616				109,820	
Median own-price elasticity		-5.11				-2.78	
Median aggregate elasticity		-0.63				-0.55	
First stage R^2		0.70				0.83	
First stage F -statistic		3,197.7				408.5	

Table 5: Demand estimates

Note: Heteroskedasticity-robust standard errors in parentheses. “WS” refers to whole/smoked.

sensitive than medium-income consumers in both markets. Finally, the instruments predict prices well, with first-stage R^2 s of 0.7 or more and F -statistics above conventional thresholds for instrument relevance.

The timeseries results are summarized in Table 6. Both demand and prices are auto-correlated, as one might expect, with the autocorrelations not changing much between pre- and post-periods or between fresh and deli turkey. Off-diagonals of the slope matrix are on average zero but tend to be slightly positive among the same firm’s products, which could reflect common within-firm shocks or advertising. Shocks tend to be nearly independent across firms on average, with correlation coefficients near zero, except for demand shocks in fresh turkey, which have a small common-value component. Finally, recall that the price timeseries first residualizes a nonparametric function of observables; I find that this function predicts prices well, explaining 92.4% of the variation, but there is still some residual variation left to be predicted by the information firms observe.³²

To summarize the estimated beliefs, I assess how accurately firms perceive their own-price elasticity, and how that changes over time for Agri Stats and non-Agri Stats customers. For each firm in each market, I simulate draws from that firm’s beliefs over (ξ, p_{-f}) . Then, for

³²Figures A.5 and A.6 show the full distribution of diagonal elements and off-diagonal elements of B^ξ and B^p respectively, with off-diagonals broken out by whether the two firms have the same owner. Figure A.7 is a scatter plot of the predictions of prices made by $g(\cdot)$ against actual prices.

Slope matrix								
Diagonal			Off-diag. mean		Shock correlation			
Mean	SD	IQR	Same firm	Diff. firm	Mean	SD	IQR	
<i>A. ξ (demand shock)</i>								
Fresh	0.56	0.24	[0.44, 0.75]	0.033	-0.004	0.12	0.14	[0.03, 0.20]
Deli	0.57	0.26	[0.46, 0.76]	0.011	0.000	0.03	0.15	[-0.06, 0.11]
<i>B. Pre-period prices</i>								
Fresh	0.37	0.24	[0.18, 0.57]	0.020	-0.002	0.01	0.16	[-0.06, 0.10]
Deli	0.45	0.26	[0.32, 0.62]	0.011	0.003	0.02	0.16	[-0.05, 0.07]
<i>C. Post-period prices</i>								
Fresh	0.42	0.22	[0.26, 0.58]	-0.003	-0.003	0.03	0.16	[-0.04, 0.11]
Deli	0.49	0.27	[0.25, 0.68]	0.028	0.003	0.02	0.13	[-0.04, 0.07]

Table 6: Summary of timeseries results

Note: “Diagonal” and “Off-diag. mean” summarize the diagonal and off-diagonal elements of B^ξ and the price coefficients of B^p , respectively. “Shock correlation” summarizes the off-diagonal correlation coefficients in Σ^ξ and Σ^p , where $\text{Cor}_{jk}^\xi = \Sigma_{jk}^\xi / \sqrt{\Sigma_{jj}^\xi \Sigma_{kk}^\xi}$, only including the lower triangle as $\text{Cor}_{jk}^\xi = \text{Cor}_{kj}^\xi$, and similar for Σ^p . “Off-diag. mean” and “Shock correlation” columns only include pairs of products that compete against each other in at least six markets.

each draw, I compute the implied own-price elasticity of that firm’s products. This gives D draws that approximate firm f ’s beliefs over its products’ own-price elasticities, $\{\tilde{\epsilon}_{jrt}\}_{d=1}^D$. I then compute the scaled root mean squared error of the beliefs as

$$\text{RMSE}_{jrt} = \sqrt{\frac{D^{-1} \sum_{d=1}^D (\tilde{\epsilon}_{jrt} - \hat{\epsilon}_{jrt})^2}{\text{Var}(\hat{\epsilon}_{jr})}} \quad (15)$$

where $\hat{\epsilon}_{jrt}$ is the actual estimated own-price elasticity for product j in market rt , and $\text{Var}(\hat{\epsilon}_{jr})$ is its variance over time. I normalize the RMSE by the within-firm-region standard deviation because it provides a natural benchmark: if firms simply guessed that their own-price elasticity was equal to the time average in each market, then the average RMSE_{jrt} for each firm-region is 1 by construction.

Table 7 shows the average RMSE_{jrt} by year for Agri Stats clients and other producers. The averages are far below 1: both groups do better than simply guessing the average. While the information exchange was ongoing, Agri Stats clients more accurately perceived their own-price elasticity than other producers. The end of the exchange causes Agri Stats clients’ accuracy to worsen by approximately 1/5 of the within-product-region elasticity standard deviation. Echoing Table 3, the bulk of this effect occurs within the first year after the information exchange ends. To better understand what firms actually learn from the

Time period	Average RMSE _{jrt}	
	Agri Stats clients	Other producers
June 2017–May 2018	0.264	0.345
June 2018–May 2019	0.269	0.313
June 2019–May 2020	0.441	0.300
June 2020–May 2021	0.463	0.294
June 2021–May 2022	0.480	0.308

Table 7: Accuracy of beliefs about own-price elasticities

Note: RMSE_{jrt} is computed using a quasi-Monte Carlo with 1,000 draws per firm-market, then averaged across firm-markets within a year.

Time period	Average marginal cost		
	Agri Stats clients	Other producers	% difference
June 2019–May 2020	\$2.93	\$3.15	7.7%
June 2020–May 2021	\$2.99	\$3.05	2.1%
June 2021–May 2022	\$3.05	\$3.02	-0.9%

Table 8: Average post-May 2019 marginal costs for Agri Stats clients and other producers

Note: Average costs are unweighted within a year-product-region and weighted by total within-year-product-region quantity across product-regions within a year.

exchange, I compute the change in RMSE of Agri Stats clients’ beliefs over their own ξ , their rivals’ ξ , and their rivals’ prices separately, without aggregating them into elasticities. I find that firms do not lose much predictive power for their own demand—the RMSE for own- ξ is 5.4% lower with the information exchange than without—and that the bulk of the effect happens in rival ξ and rival prices, with post-May 2019 RMSE increases of 19.1% and 26.5% respectively. This result suggests that Agri Stats primarily informed firms about rival demand shocks and prices, previewing a finding in Section 8.

Finally, I find that supply-side changes not captured by information effects are at least in part responsible for the relative growth in Agri Stats clients’ prices toward the end of the sample period. As shown in Table 8, Agri Stats clients had 7.7% lower recovered marginal costs on average than other producers in the year after the information exchange ended, but were no more efficient by the end of the sample period. One possible interpretation is that, as claimed by Agri Stats, information sharing resulted in investment efficiencies, the loss of which caused long-run convergence of marginal costs. Of course, as mentioned in Section 3.2, this convergence could be due at least in part to input and labor markets becoming less monopsonistic after the exchange ended, and hence should not necessarily be credited as a positive effect of information sharing. Another interpretation is that firms

slowly began attempting to collude in the years following the end of the exchange, although collusion should in theory have been less effective without the monitoring provided by Agri Stats. Regardless of the ultimate explanation, potential long-run changes in marginal costs after the end of the exchange will be important to account for when testing for product market changes in conduct caused by Agri Stats’ exit.

7 Conduct testing

Are prices consistent with firms playing a Bertrand-Markov equilibrium with and without information sharing? Like complete-information conduct models, each vector of prices maps to a unique vector of marginal cost shocks, ω . Thus, following the conduct testing literature, we can exploit a restriction on how ω should behave under the true conduct model to formulate a statistic to test whether and to what extent conduct changed when the information exchange ended (Berry and Haile, 2014).

One option for such a restriction is to exploit exogeneity between cost shocks and a function $f(\cdot)$ of instruments z that shift the way markups depend on conduct without affecting ω , and form a moment restriction of the form $\mathbb{E}[f(z)'\omega] = 0$ (Backus et al., 2021). There are two reasons why this approach may not be ideal in this setting. First, unlike in the case considered by Backus et al. (2021) but similar to Miller and Weinberg (2017), here there is additional variation that can be exploited from the Bertrand-Markov conduct assumption after the end of information sharing, and a restriction based only on contemporaneous instrument-exogeneity does not use that variation. Second, to the extent that the information exchange caused long-run changes in unobservable costs (for example, due to investment efficiencies), the restriction may not hold if such changes are correlated with typical markup-shifting instruments (e.g., product characteristics).

To motivate an alternative restriction, consider the evolution of average cost shocks over time assuming that firms played a competitive Bertrand-Markov equilibrium throughout the sample period (Figure 7). The cost shocks for Agri Stats clients are on a relatively consistent trend up until May 2019, but they immediately drop after May 2019, while the drop of other producers’ costs is far less. However, like overall marginal costs, the Agri Stats client cost shocks converge back to the level of other firms’ shocks towards the end of the sample period, suggesting that these trends are not explained by changes to observables in the marginal cost function. Further, this immediate effect post-May 2019 is unlikely to have been caused by changes to unobservable marginal costs, as long as any foregone efficiencies would not have immediately affected costs once the information exchange ended.

I formalize this idea in the following moment condition, which states that there is no




Figure 7: Average cost shocks over time under Bertrand-Markov conduct

Note: Cost shocks are weighted by total product-region-level quantity.

discontinuity in the average difference between Agri Stats clients' cost shocks and other producers' cost shocks when the information exchange ended:

$$m := \mathbb{E} \left[\left(\lim_{t \uparrow t^*} \omega_{AS,rt} - \omega_{AS^C,rt} \right) - \left(\lim_{t \downarrow t^*} \omega_{AS,rt} - \omega_{AS^C,rt} \right) \right] = 0, \quad (16)$$

where t^* corresponds to May 2019, and AS and AS^C correspond to Agri Stats clients and other producers respectively. m is a “difference-in-discontinuities,” and can therefore be estimated using standard regression discontinuity techniques (Grembi et al., 2016). To do so, I regress $\hat{\omega}_{jrt}$ on the full set of pairwise interactions between a time trend, a post-May 2019 indicator, and an Agri Stats indicator, weighting observations by $k\left(\frac{t-t^*}{h}\right)$ for a kernel function k and positive scalar bandwidth h . $\hat{m}$ is the coefficient on the interaction between the Agri Stats indicator and the post-May 2019 indicator. Given a consistent estimate of m 's standard error $\hat{\sigma}$, which I compute using the bootstrap, the test statistic $t = \hat{m}/\hat{\sigma}$ is asymptotically standard normal under the correct conduct model.³³ I select the bandwidth by minimizing the out-of-sample estimated difference-in-discontinuities among periods outside the 3 months before to 6 months after the end of the exchange, which returns a bandwidth of 12 months.

The first row of Table 9 displays the estimates. Under the null hypothesis of Bertrand-Markov conduct throughout the sample period, the estimated discontinuity in Agri Stats customers' relative cost shocks at May 2019 is 0.220 (signifying higher costs with than without the information exchange), compared to a discontinuity of 0.003 for other producers.

³³In particular, the asymptotics here take the number of regions to infinity.

Conduct model	Discontinuity in ω		Difference	Test statistic	p -value
	AS clients	Others			
Bertrand-Markov	0.220 (0.082)	0.003 (0.042)	0.217 (0.101)	2.14	0.032
Bertrand-Nash	0.069 (0.036)	0.015 (0.029)	0.054 (0.048)	1.12	0.265

Table 9: Baseline conduct test

Note: Each row tests the null hypothesis that the specified conduct model was followed both before and after the information exchange ended, under the assumption that equation (16) holds. Bootstrap standard errors, computed with 500 bootstrap draws, are in parentheses. Discontinuities are estimated using a local linear regression before and after the bandwidth, with an Epanechnikov kernel and a bandwidth computed to minimize placebo discontinuities over June 2017-February 2019 and December 2019-May 2022, which results in a bandwidth of 12 months.

This results in an estimated moment value of $\hat{m} = 0.217$, which, when normalized by the standard error, results in rejecting the null hypothesis that conduct was Bertrand-Markov throughout ($p = 0.032$).

A question one might have had throughout the paper is whether information effects are empirically relevant enough to be worth modeling, given the theoretical subtlety of their impact on prices. If information effects turn out not to bias the result of conduct tests, then one could credibly use existing complete-information conduct models to test conduct in settings where information effects are present. To illustrate how much information effects matter, I reestimate marginal costs and recompute the conduct test under the standard complete-information Bertrand-Nash model, where firms still maximize their own static profits in every period but with complete information over demand shocks and rivals' prices. The estimated discontinuities and associated p -value assuming Bertrand-Nash conduct before and after May 2019 are listed in the second row of Table 9. While this results in more precise estimates due to the lack of timeseries estimation adding additional noise, the estimated moment value is much smaller than the Bertrand-Markov value, and Bertrand-Nash conduct results in a statistically insignificant estimated difference-in-discontinuities.³⁴

While the preceding single-hypothesis test of Bertrand-Markov conduct is valid under correctly-specified demand and cost functions, the conduct testing literature has emphasized making comparisons between different conduct models to add robustness to model misspecification (Backus et al., 2021; Duarte et al., 2024). In this spirit, I consider a parameterized family of Markov equilibria in pricing rules in which Agri Stats subscribers maximized a

³⁴Another potential test is whether there is a discontinuity in Agri Stats clients' cost shocks, without differencing other producers' cost shocks. These estimated discontinuities are shown in the first column of Table 9. This test results in an even lower p -value for the test of Bertrand-Markov conduct, while Bertrand-Nash conduct is again not rejected at 5% significance (although it is borderline).

Model	Estimated conduct parameter	95% confidence interval
Incomplete information	0.474	[0.210, 0.726]
Complete information	0.177	[0.000, 0.548]

Table 10: Conduct parameter estimates

Note: 95% confidence intervals are computed with 500 bootstrap draws. $\tau^* \in [0, 1]$ is estimated to minimize the square of the left hand side of equation (16), assuming Agri Stats subscribers maximized the objective in (17). Discontinuities are estimated using a local linear regression before and after the bandwidth, with an Epanechnikov kernel and a bandwidth computed to minimize placebo discontinuities over June 2017-February 2019 and December 2019-May 2022, which results in a bandwidth of 12 months.

weighted sum of their own expected profits and their fellow subscribers' expected profits while sharing information. That is, before May 2019, Agri Stats clients' prices solved

$$p_{f_{rt}}(\tau) = \arg \max_{p_{f_{rt}}} \mathbb{E}_{f_{rt}} \left[\pi_{f_{rt}}(p_{rt}; \xi_{rt}) + \tau \sum_{f' \in AS \setminus \{f\}} \pi_{f'_{rt}}(p_{rt}; \xi_{rt}) \right] \quad (17)$$

where the conduct parameter $\tau \in [0, 1]$ describes the fraction of other subscribers' profits internalized by Agri Stats clients. $\tau = 0$ corresponds to Bertrand-Markov, while $\tau = 1$ corresponds to joint expected profit maximization among subscribers.³⁵ Each candidate τ results in a different vector of cost shocks and therefore a different estimate of $\hat{m}(\tau)$, which can be used to compare different conduct models. I then estimate a conduct parameter by searching for the τ^* minimizing the distance between $\hat{m}(\tau^*)$ and zero.³⁶

As shown in Table 10, I estimate $\tau^* = 0.474$, demonstrating that Agri Stats clients acted as though they internalized nearly half their fellow subscribers' profits while sharing information. The 95% confidence interval is bounded away from zero. For comparison, I compute τ^* under the analogous complete-information model and find a conduct parameter of 0.177, with a 95% confidence interval including zero. These results suggest two broad takeaways. First, information effects cannot explain the price changes stemming from the end of the exchange. Even when considering alternative conduct models, the best-fitting model is

³⁵This family of conduct models was shown by [Corts \(1999\)](#) to not necessarily correspond to the reduced form of a complete-information dynamic game, and there is no reason to suspect that it does so in the present incomplete-information setting. This exercise should therefore be viewed as a quantification of conduct effects, rather than a fully specified model of how information sharing affects collusion. I discuss the challenges of modeling conduct effects of information exchanges further in Section 8.

³⁶In this case, there exists an interior τ^* such that the estimated discontinuity is exactly zero; I find that $\hat{m}(\tau)$ is monotone in τ , so τ^* is unique, and I search for it using bisection. An alternative approach to comparing conduct models is [Rivers and Vuong \(2002\)](#)-style pairwise testing based on differences in implied moment values. Such tests have different limiting distributions depending on whether the values of τ being considered result in moment values close to zero ([Hall and Pelletier, 2011](#)), and here there is a unique interior value of τ such that $\hat{m}(\tau) = 0$, making the limiting distribution of the Rivers-Vuong test statistic unclear. Hence, I adopt conduct parameter estimation as my main approach.

one in which firms exhibited some degree of coordination in prices while sharing information. Second, ignoring these information effects would have led to underestimating the extent to which firms were coordinating. This both emphasizes the importance of modeling these forces and points to the procompetitive nature of information effects in this context. I examine the latter in more detail in the following section.

8 Mechanisms and counterfactuals

The preceding section demonstrated that Agri Stats had economically significant information and conduct effects. This section quantifies and decomposes these effects to better understand the economic mechanisms behind Agri Stats’ effect on prices and welfare, including how each of the forces discussed in Section 2 contributes to the overall information effect and what firms learn from the exchange. I then turn to the potential effects of counterfactual information exchanges, a policy-relevant question given the prevailing regulatory strategy of pursuing restrictions on information sharing as a means of preventing collusion, both generally and in the settlement to the Agri Stats case specifically.

8.1 Economic mechanisms of Agri Stats

To assess the impact and mechanisms of the Agri Stats exchange, I simulate what would have happened had Agri Stats continued operating after May 2019. I first simulate a Bertrand-Markov equilibrium in which each firm maximizes their expected profits with the Agri Stats signal structure. Comparing these simulated outcomes to the realized outcomes gives the information effects of Agri Stats without any conduct effects. I then simulate a Markov equilibrium in pricing rules with the estimated conduct parameter, incorporating the conduct effects of Agri Stats.

The primary computational challenge in computing counterfactual Markov equilibria in pricing rules is simulating the counterfactual long-run law of motion of prices. In equilibrium, firms set prices to optimize their objective given beliefs, update these beliefs given the conjectured long-run price kernel, and conjecture a long-run price kernel that is consistent with the prices rivals actually set. For the realized information environment, the long-run distribution can be directly estimated from the data. Conversely, finding a counterfactual equilibrium requires finding a long-run price kernel Π^+ such that $\Pi^+ = \Phi(\Pi^+)$, where $\Phi(\Pi^+)$ gives the law of motion of the optimal prices given the beliefs generated by Π^+ .³⁷ To find this fixed

³⁷Unlike the standard supply-side model, this problem is not separable by market: equilibrium outcomes in one time period affect the long-run price distribution, thereby affecting equilibrium outcomes in all other periods.

Scenario	Average prices			Profits (\$m/year)			CS (\$m/year)	TS (\$m/year)
	AS	Others	Total	AS	Others	Total		
No sharing (actual)	4.24	4.10	4.19	210.2	81.6	291.8	396.1	687.9
Just direct effect	4.18	4.10	4.15	208.9	80.4	289.3	404.9	694.2
Direct + eqm effect	4.16	4.09	4.14	208.1	80.2	288.3	406.9	695.2
Information effect	4.16	4.09	4.14	208.0	80.3	288.2	407.2	695.5
Info + conduct effect	4.39	4.12	4.30	217.5	83.3	300.7	378.3	679.1

Table 11: Decomposing the impact of Agri Stats

Note: Average prices are weighted by no-sharing (actual) total quantity. In all scenarios except no sharing, Agri Stats clients share prices, quantities, costs, and inventories, and other producers just observe their own prices and quantities. In the “Information effect” scenario, firms maximize own-expected profits. In the “Just direct effect” scenario, firms update their beliefs using rivals’ no-sharing prices and shares, rather than the realized prices and shares. In both the “Just direct effect” and “Direct + eqm effect” scenarios, firms update beliefs using the estimated long-run law of motion of prices from the no-sharing equilibrium rather than the information effect equilibrium. In the “Info + conduct effect” scenario, Agri Stats clients maximize (17) using the estimated conduct parameter, while other producers continue to maximize own-expected profits.

point, I minimize the distance between $\Phi(\Pi^+)$ and Π^+ ; for more details, see Appendix B.3.³⁸

Table 11 shows Agri Stats’ impact on average prices, profits, and welfare through both the conduct and information channels. Comparing the first and last rows, I find that Agri Stats raises average prices by 2.5% (11 cents/lb) and lowers consumer surplus by 4.5% (\$17.8 million annually). This aggregate effect is the composition of a procompetitive information effect and an anticompetitive conduct effect: the information effect lowers prices by 1.3% and increases consumer surplus by 2.8%, and the conduct effect raises prices by 3.9% and lowers consumer surplus by 7.1%. The impact on Agri Stats clients’ prices alone lines up closely with the reduced form results (Table 2): the conduct effect of 5.4% is close to the short-run reduced-form estimate of 5.1%, and the overall effect of 3.5% to the long-run estimate of 3.4%. Agri Stats increases profits, primarily for Agri Stats clients (3.5%) but also for other producers (2.0%), which benefit from the softened competition and higher prices from Agri Stats clients. However, information effects alone lower Agri Stats clients’ profits, which may suggest a revealed preference argument that Agri Stats facilitates collusion: firms would not rationally pay for Agri Stats if there were no conduct effect. Agri Stats causes deadweight loss, reducing total surplus by 1.3%.

To dig into the mechanisms behind information effects, I decompose them into the forces described in Section 2: the direct effect, equilibrium effect, and signal distribution effect. Let $p(y^{NS}(p^{NS}, s^{NS}), \Pi^{+, NS})$ be the vector of expected profit-maximizing prices at beliefs given by the no-sharing signal functions $y^{NS}(\cdot)$ —evaluated at the no-sharing equilibrium

³⁸I keep marginal costs fixed at their estimated values in all simulations. If, as claimed by Agri Stats, information sharing directly causes marginal cost efficiencies, then estimated counterfactual prices with information sharing may be too high.

Module	Reduction in average RMSE for Agri Stats clients			
	Own-price elasticity	Own ξ	Rival ξ	Rival price
Fresh turkey	42.6%	11.7%	17.7%	41.6%
Deli turkey	36.0%	4.1%	32.5%	17.0%
Overall	41.1%	8.0%	25.4%	21.0%

Table 12: How much and what firms learn from Agri Stats

Note: RMSEs of own-price elasticities are normalized by the within-product SD of own-price elasticities before aggregating and taking percentage differences. For each of own ξ , rival ξ , and rival prices, RMSEs of beliefs over each applicable product-market combination are averaged within each firm-market combination, weighted by quantities, and then averaged across firm-markets.

outcomes (p^{NS}, s^{NS}) —and the no-sharing equilibrium price kernel $\Pi^{+,NS}$. The direct effect changes $y^{NS}(\cdot)$ to $y^S(\cdot)$, the information sharing signal function, giving Agri Stats clients more precise information. The equilibrium effect changes the inputs to the signal function to (p^S, s^S) , the equilibrium values with information sharing, updating firms’ beliefs based on the new prices and shares of their rivals. Finally, the signal distribution effect changes $\Pi^{+,NS}$ to $\Pi^{+,S}$, accounting for the new long-run law of motion of prices. $p(y^S(p^{NS}, s^{NS}), \Pi^{+,NS})$ incorporates just the direct effect, while $p(y^S(p^S, s^S), \Pi^{+,NS})$ incorporates the direct and equilibrium effects without the signal distribution effect.

The results of the decomposition are shown in the middle rows of Table 11. The direct effect accounts for 84.2% of the overall information effect on Agri Stats clients’ prices, with the equilibrium effect and signal distribution effect contributing 15.5% and 0.3% respectively. The direct effect on Agri Stats clients’ profits, while a smaller proportion at 55.7%, is still negative; the business-stealing externality caused by rivals responding to the signal outweighs firms being better able to maximize their profits due to the additional information provided by the signal. The direct effect on non-subscribers’ prices is zero by definition; the remainder of the effect (although small) is roughly evenly split between the equilibrium and signal distribution effects.

These results also provide insight into how much and what firms learn from the exchange, and how this learning depends on primitives. Table 12 shows how Agri Stats affects the accuracy of its clients’ beliefs, as measured by the reduction in average RMSE over own-price elasticities, own ξ , rival ξ , and rival prices. Agri Stats improves its clients’ beliefs over their own-price elasticities by 41%, with similar totals for both fresh and deli turkey. However, the two modules differ in *what* firms learn about. In both modules, firms have a relatively good sense of their own ξ without Agri Stats. However, Agri Stats provides relatively more information about rival demand for deli turkey, which makes sense: demand shocks are more common-valued in fresh turkey (Table 6), suggesting that fresh turkey producers can

use knowledge of their own shock to better forecast their rivals' shocks without Agri Stats. Conversely, Agri Stats reduces RMSE over rivals' prices more for fresh turkey producers than deli turkey producers. Table A.1 decomposes information and conduct effects by module and finds that while information effects lower prices by 2.1% in fresh turkey, they slightly raise prices in deli turkey, highlighting that these differences in primitives are ultimately economically important for information effects.

8.2 Bounds on collusion via information exchange

Simulating potential conduct effects of counterfactual information exchanges presents a conceptual challenge. While the estimated conduct parameter can credibly be used to simulate the effects of the Agri Stats exchange, it cannot be used to simulate the effects of other exchanges; different information structures likely correspond to different conduct parameters, and the lack of microfoundation of the conduct parameter model (Corts, 1999) makes selecting counterfactual parameters an *ad hoc* process at best. However, solving a fully microfounded dynamic collusive game with counterfactual information structures is difficult. Incomplete information exchanges lead to repeated games with private monitoring, a class of games for which finding optimal strategies is known to be intractable (Kandori, 2002). The problem stems from firms being able to condition their actions on their private history; this means firms must form beliefs over the entire signal history, rendering the recursive methods used in the analysis of public monitoring games inapplicable (Abreu et al., 1990; Fudenberg et al., 1994). Further, while Sugaya (2022) proved a folk theorem for private monitoring games, implying that infinitely patient firms can collude if signals have a sufficiently large support, this result is less useful for characterizing the potential for collusion in real-world situations with finitely patient firms.

While the theory literature has not characterized exact optimal strategies in repeated games with imperfect monitoring, it has derived necessary conditions that equilibrium distributions of actions must satisfy. Thus, my approach is to bound the distribution of prices under different information structures by solving a constrained optimization problem that maximizes average prices subject to these necessary conditions holding. Let $\alpha \in \Delta(\mathcal{P})$ be a stationary distribution over prices for all firms with support bounded by $\mathcal{P} = \prod_f \mathcal{P}_f$, and let $\Psi(\alpha) = \mathbb{E}_\alpha \left[\sum_j \psi_j p_j \right]$ be the weighted average price of a particular price distribution given weights ψ_j that sum to unity.³⁹ Suppose any equilibrium must satisfy the inequality $C(\alpha; y) \leq \mathbf{0}$, where $C(\cdot)$ is a vector of necessary conditions and $y = \{y_f\}_{f=1}^F$ is a counterfac-

³⁹Here I focus on bounds on average prices. One could also use the same procedure to bound expected profits, consumer surplus, or total surplus.

tual signal function of interest. A bound on average prices is found by solving

$$\max_{\alpha \in \Delta(\mathcal{P})} \Psi(\alpha) \quad \text{s.t.} \quad C(\alpha; y) \leq \mathbf{0}. \quad (18)$$

If the necessary conditions $C(\cdot)$ are also jointly sufficient, then this bound is sharp.

The immediate question is what necessary conditions are included in $C(\cdot)$. One standard condition is that it cannot be profitable for firms to deviate if rivals are able to punish them with full Nash reversion (Sugaya and Wolitzky, 2018a). Let p^C be the competitive price vector that simultaneously maximizes each firm's own expected profits, $\mathbb{E}_f[\pi_f(p)]$ be firm f 's expected profits given their beliefs over ξ and a price vector p , and δ be a common discount factor. Then for all prices $p \in \text{supp}(\alpha)$, all firms f , and all potential deviations p'_f , a necessary condition for α to be an equilibrium price distribution is

$$\overbrace{(1 - \delta) (\mathbb{E}_f[\pi_f(p'_f, p_{-f})] - \mathbb{E}_f[\pi_f(p)])}^{=\text{NashIC}_f(p, p'_f, \alpha)} + \delta \left(\mathbb{E}_f[\pi_f(p^C)] - \int \mathbb{E}_f[\pi_f(\tilde{p})] d\alpha(\tilde{p}) \right) \leq 0. \quad (19)$$

The first term is the payoff the firm gets from deviating to p'_f in the current period, when it is supposed to set p_f . The second term is the decline in expected future payoffs from rivals punishing that deviation with Nash reversion in all future periods. Of course, if firms cannot perfectly observe their rivals' actions, then they may not be able to perfectly punish any deviation with Nash reversion; that makes this a necessary but generally insufficient condition.⁴⁰

The Nash IC condition does not depend on the signal function, and hence will not characterize how changing the signal function affects possible collusive outcomes. To do so, I add a second set of conditions, derived by Sugaya and Wolitzky (2023; henceforth SW23). SW23's result characterizes signal functions by their detectability: how much the distribution of the signal changes when a firm changes its price. I briefly summarize the result below.⁴¹

Suppose that firms collectively observe signals $y = (y_1, \dots, y_F) \in \mathcal{Y}$. They depend on prices, but are stochastic; let $f_y(\cdot | p)$ be the density of the conditional distribution given a vector of prices $p \in \mathcal{P}$. The leading example is that the signals are functions of shares, which are stochastic conditional on prices due to demand uncertainty. The χ^2 -divergence, a

⁴⁰The condition is sufficient when all firms' signals include the full price vector. In this case, it reduces to the typical complete-information incentive compatibility constraint.

⁴¹Awaya and Krishna (2019) derive a similar restriction on equilibrium outcomes under alternative monitoring structures, which Awaya and Krishna (2020) use to bound collusive outcomes under counterfactual monitoring structures in a symmetric demand system with log-linear signals. The intuition behind their bound is similar—their result uses total variation distance, rather than χ^2 -divergence, as the measure of the detectability of a monitoring structure.

standard measure of statistical distance, between the conditional signal distributions under p and p' is

$$\chi^2(p' \parallel p; y) = \int_{\mathcal{Y}} \left(\frac{f_y(y \mid p') - f_y(y \mid p)}{f_y(y \mid p)} \right)^2 f_y(y \mid p) dy. \quad (20)$$

where if the distributions of $y \mid p$ and $y \mid p'$ have different supports (for example, if firms directly share prices), then $\chi^2(p' \parallel p; y) = +\infty$. Let a *manipulation* for firm f be a function $m_f : \mathcal{P}_f \rightarrow \Delta(\mathcal{P}_f)$; the interpretation is that when firm f is “supposed to” set price p_f , it sets price(s) according to $m_f(p_f)$ instead. The *detectability* of this manipulation is the expected χ^2 -divergence, while the *gain* from it is the expected difference in profits; that is,

$$\chi^2(m_f, \alpha; y) = \mathbb{E}_{p \sim \alpha} \left[\mathbb{E}_{p'_f \sim m_f(p_f)} [\chi^2(p'_f, p_{-f} \parallel p; y)] \right], \quad (21)$$

$$g_f(m_f, \alpha) = \mathbb{E}_{p \sim \alpha} \left[\mathbb{E}_{p'_f \sim m_f(p_f)} [(\pi_f(p'_f, p_{-f}) - \pi_f(p))] \right]. \quad (22)$$

Theorem 1 of SW23 states that a necessary condition for α to be an equilibrium distribution over prices is that for all firms f and manipulations m_f ,

$$g_f(m_f, \alpha) \leq \sqrt{\frac{\delta}{1 - \delta} \chi^2(m_f, \alpha; y) \text{Var}(\pi_f(\alpha))}. \quad (23)$$

where $\text{Var}(\pi_f(\alpha)) = \text{Var}_{p \sim \alpha}(\mathbb{E}_f[\pi_f(p)])$ is the variance in firm f ’s expected profits if the price distribution is α . The intuition is the following. Collusion is impossible if a firm can deviate profitably (high g_f) in a way that cannot be detected by rivals (low χ^2). In particular, the ratio between the two must be bounded by a function of how much firms weight future punishment relative to present gain and the size of that punishment, with the latter bounded by the variance of profits. This condition directly applies to the deviation in which firm f sets price $m_f(p_f)$ in one period, and then follows the action distribution implied by α in all periods that follow; this may not be the optimal deviation, which is why the condition is potentially insufficient.⁴²

To implement (18), I approximate α as a discrete distribution with N support points $\{p_n, a_n\}_{n=1}^N$. Thus, if there are J products, then there are $N(J + 1)$ choice variables: locations of N J -vectors of prices, and N weights summing to 1. I bound the support of this distribution by the competitive price vector p^C and joint expected profit-maximizing price vector p^M , which defines $\mathcal{P}$. To implement the necessary conditions, which must hold for all deviations, I write them as inner maximization problems enumerated over all possible

⁴²Theorem 2 of SW23 demonstrates that the condition is sufficient in the low discounting-low monitoring double limit.

deviations. The program is therefore

$$\begin{aligned}
\alpha^*(y) &= \arg \max_{\alpha=\{p_n, a_n\}} \sum_{n,j} a_n \psi_j p_{n,j} \\
\text{s.t.} \quad & \sum_n a_n = 1, \quad a_n > 0, \quad p_n \in [p^C, p^M] & \forall n, \\
& \max_{p'_f} \text{NashIC}_f(p, p'_f, \alpha) \leq 0 & \forall n, \forall f, \\
& \max_{m_f} \frac{g_f(m_f, \alpha)}{\sqrt{\chi^2(m_f, \alpha; y)}} \leq \sqrt{\frac{\delta}{1-\delta}} V_f(\alpha) & \forall f.
\end{aligned}$$

This problem is feasible; a point mass at the competitive price satisfies all constraints.⁴³ While this is a nonlinear and potentially nonconvex nested maximization problem, several factors make it tractable. First, the Nash IC maximization problem simply amounts to finding a best response for firm f to a rival price vector, hence is solvable using standard iteration techniques (Morrow and Skerlos, 2011). Second, all functions of p in the problem are just moments of expected profits, except for $\chi^2(\cdot)$, while signal distributions are well-approximated by additive-logistic normals, which admit a closed-form $\chi^2(\cdot)$.⁴⁴ This implies that all model objects have analytic gradients, while the gradient of the constraints with respect to elements of α can be derived using the envelope theorem.

This method provides an upper bound on average prices under a given information structure for one market, taking as inputs beliefs over ξ for each firm in that market. To compute an aggregate bound, I compute the bound separately in each region-module, using the most common product assortment in that region with beliefs at the stationary distribution implied by the timeseries estimates.⁴⁵ Each product has its average marginal costs in each region, and I calibrate $\delta = 1.0801^{-1/12}$, reflecting the monthly analogue of the food processing indus-

⁴³At the competitive price, there is no gain from deviation, trivially satisfying the Nash IC constraints. This also implies that the gain from any manipulation is bounded above by zero, while the right-hand side of (23) is bounded below by zero, so this solution satisfies the SW23 condition as well. I use the rival signal functions $y_{-f}(\cdot)$ as the signals in firm f 's SW23 constraint, as including the entire signal vector including firm f 's signal would result in a degenerate solution, as firm f observes its own price. Note that if firm f shares its price directly with another firm, then firm f 's SW23 constraint is vacuously satisfied. In this case, only the Nash IC constraints bind on firm f .

⁴⁴Signals are generally collections of shares, and the inside share vector from a logit demand system with a normally distributed ξ , conditional on a price vector, is exactly additive logistic-normally distributed. Here it is an approximation because of the random coefficients.

⁴⁵Because I collapse the time dimension, I refer to a region-module as a market for the remainder of the paper. I note three caveats about this strategy. First, the bound is computed separately in each region, meaning there is no scope for multimarket contact effects (Bernheim and Whinston, 1990). Second, the bound does not dynamically account for how expected future movement in ξ changes deviation constraints (Rotemberg and Saloner, 1986). Third, setting common stationary priors over ξ does not take into account the indirect effect that information sharing may have on collusive outcomes by homogenizing firms' beliefs themselves, affecting perceived deviation payoffs.

try’s 2021 weighted average cost of capital of 8.01% (Damodaran, 2021). Additional details, including gradients of all objects, are available in Appendix B.4.

8.3 Aggregation versus participation in information sharing

With these upper bounds, we can simulate the effects of counterfactual restrictions on data sharing. I focus on the tradeoff between two aspects of exchanges: *aggregation* (allowing information exchanges to share only aggregates, instead of disaggregated information) and *participation* (the fraction of the market that contributes to and receives the information exchange). Whether aggregation prevents collusion via information exchange has been a matter of broader policy debate: the FTC and DOJ allowed a “safe harbor” for sufficiently aggregated information exchanges until rescinding this policy in 2023, and have since submitted statements of interest in multiple cases stating that aggregated information exchanges can still harm competition.⁴⁶ Meanwhile, the DOJ’s 2025 settlement with RealPage, a rent-setting software, prevents it from using any nonpublic competitor information when determining a firm’s rental price “regardless of form, aggregation, anonymization, or age” (US v. RealPage, Inc., 2025). However, while the Agri Stats settlement prevents it from exchanging any price information, it allows sharing aggregates of other information, including quantities (US et al. v. Agri Stats, Inc., 2026). The settlement also mandates that Agri Stats allow anyone to purchase the reports, suggesting that participation may increase.

To investigate the tradeoff between aggregation and participation, I compute the upper bound under several different signal structures. I start with firms’ competitors not observing prices and shares, as is the case under no information sharing. I then compute bounds under the pre-settlement status quo: product-level quantities and/or prices are shared only among Agri Stats clients. Finally, I compute the bounds if all firms participate, but have access only to the total quantity—which complies with the Agri Stats settlement—and progressively add more disaggregated information. I scale average prices in each market to be between zero and one, where the competitive price is zero and the monopoly price is one.

Figure 8 shows the distribution of the scaled worst-case average price across markets for each of these signal structures. The first thing to note is that collusive prices can already be relatively high without information sharing: firms can attain 70.8% of monopoly prices on average. This suggests that even without observing rivals’ prices directly, firms can already detect deviation somewhat well using their own shares. There is a long left tail, however; the

⁴⁶For example, the DOJ’s statement of interest in *In Re Frozen Potato Products Antitrust Litigation* says that “to the extent that Defendants suggest that exchanges of aggregated, anonymized, or backward-looking information are incapable of causing anticompetitive effects, Defendants fail to grapple with the ways that such information can harm the competitive process” (Department of Justice, 2026, p. 6).




Figure 8: Distribution of upper bounds on average prices across markets

Note: I compute the upper bound on average prices within each market (region-module), then plot the distribution across markets. The whiskers are the 10th and 90th percentiles, the edges of the box are the 25th and 75th percentiles, the middle line is the median, and the point is the average.

10th percentile bound is 18.2%, suggesting that information sharing is necessary to sustain full collusion in many markets. Information sharing, as implemented by Agri Stats, makes a difference: Agri Stats clients sharing the full vector of prices and quantities among them results in an upper bound of 86.2% on average. Notably, once Agri Stats clients share the full vector of quantities, sharing prices does not matter much; price deviations are nearly perfectly detectable with precise information about quantities.

The key result, however, is that all firms sharing aggregate information results in the possibility of near-perfect collusion. Even if firms observe total quantity, the average upper bound increases to 95.1% of monopoly prices, and the 10th percentile increases to 90.5%. While sharing disaggregated quantities and prices does increase the bound under full participation, the increase is not as large—average bounds increase to 95.7% (disaggregated quantities) and 96.2% (disaggregated quantities and prices).⁴⁷

What explains why sharing aggregate information with all participants has such a large

⁴⁷Figure A.9(a) disaggregates the distribution into those who participate in any exchange (i.e., Agri Stats clients) and those who only participate in full-participation exchanges (other producers). It shows that full participation allows for the bound to increase for all firms, although (perhaps obviously) substantially more for firms that do not participate in the partial exchange.

impact? This computational result is similar to the analytical result of [Awaya and Krishna \(2020\)](#), who find that aggregate information is more useful for detecting deviation when shocks are uncorrelated among firms. The intuition, in their words, is that “whereas individual sales can be rather variable, aggregate sales are usually much less variable. Thus, if a firm cuts its price from an agreed upon level, this will cause, with high probability, an increase in aggregate sales observed by all” (pp. 422–423). Conversely, if aggregate sales are correlated across firms, then sharing aggregate information does not matter much—firms can already use their own sales information to get a sense for the common demand shock, with the total quantity not adding much information. Indeed, Figure [A.9\(b\)](#) plots the distribution of the bounds separately for fresh turkey, for which demand shocks are more correlated across products, and deli turkey, for which demand shocks are more independent. Without information sharing, fresh turkey is more susceptible to collusion than deli turkey, with average bounds of 87.1% and 55.6% respectively. This pattern continues under partial participation. However, under full participation, the fresh turkey bound is eclipsed by the deli turkey bound, at 94.0% and 96.1% respectively. While there are other differences between the modules—fresh turkey is more concentrated and more price-elastic than deli turkey, affecting deviation payoffs—this finding reinforces the notion that the structure and correlation of shocks is a key contributor to how information exchanges affect competition.

I close with several points about the interpretation of these results. First, these results are upper bounds, not exact equilibria. While one hopes that the extent to which these bounds are sharp is roughly consistent across counterfactual exchanges, there is no formal guarantee. In addition, information exchanges may be important for equilibrium selection, not just for broadening the set of equilibria. Second, the level of each bound is a function of the discount factor, and the effective discount factor in information exchanges that might be subject to antitrust scrutiny could be lower than that implied by firms’ cost of capital ([Harrington and Chang, 2009](#)). Thus, comparisons across counterfactuals may be more reliable than levels within a counterfactual. Third, the comparison between partial and full-participation information exchanges takes the set of participating firms as given. Endogenizing participation in information exchanges may be a worthwhile area for future study. Finally, one could interpret these results as suggesting that the Department of Justice erred in allowing Agri Stats to share aggregate quantity information and make its reports more broadly available for purchase. However, the intent behind the latter clause was so that poultry purchasers and workers could acquire the reports, which may have additional positive effects outside the model. In addition, the only protein for which Agri Stats currently produces reports is chicken, where its clients comprise nearly all of the market, limiting the scope of the unintended consequence of broadening participation in the exchange.

9 Conclusion

There is rising concern among antitrust enforcers, policymakers, and observers about whether firms use common vendors or software to collude with one another, sparking a multitude of cases alleging anticompetitive information sharing. This paper presents a new framework for empirically evaluating the effects of such exchanges, combining a workhorse industrial organization supply-and-demand model with an information structure that captures how information sharing changes firms' information sets, and therefore actions. The model can be used to test whether firms are using information exchanges to collude, as well as simulate the information effects of counterfactual exchanges. I also leverage recent theory results to bound counterfactual outcomes in order to assess the worst-case collusive scenario under different restrictions on information sharing.

I apply this machinery to the Agri Stats natural experiment, in which the largest turkey processors' prices fell by 5% in the year after they suddenly stopped sharing price, quantity, and inventory information with one another. This price change cannot be rationalized by firms continuing to maximize expected profits with more precise signals while sharing information; instead, I estimate that subscribing firms set prices as though they internalized 47% of fellow subscribers' profits while sharing information. Both conduct tests with complete-information models and counterfactual results suggest that information effects countervailed conduct effects; Agri Stats would have lowered prices by 1.3% had firms used the information to compete. Finally, the bounds suggest full participation in information exchanges substantially raises worst-case collusive prices, even if firms just share aggregate quantity information, while sharing disaggregated information within a smaller group has a smaller impact.

These results have several policy implications. First, while the conduct test and counterfactuals are estimated in the context of the poultry processing industry, the model can be used to test conduct and simulate information exchanges in other industries and contexts, as it captures the most common forms of information sharing. Second, conduct tests and counterfactuals reinforce theoretical results that complete-information conduct tests may be biased in the presence of changing information sets. Finally, the counterfactual results suggest that recent policy movement away from "safe harbors" for aggregated information exchanges is a good idea. Instead, information exchanges should be judged on a case-by-case basis with the economic fundamentals of the market in mind.

References

- ABREU, DILIP, DAVID PEARCE, AND ENNIO STACCHETTI. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5): 1041–1063.
- AGRI STATS. 2023. *DOJ Lawsuit Against Agri Stats Is Wrong on the Law and Bad for Consumers*. Available at <https://www.agristats.com/doj-lawsuit-against-agri-stats-is-wrong-on-the-law-and-bad-for-consumers/>. Accessed on January 13, 2026.
- ALBÆK, SVEND, PETER MØLLGAARD, AND PER B. OVERGAARD. 1997. "Government-Assisted Oligopoly Coordination? A Concrete Case." *Journal of Industrial Economics* 45 (4): 429–443.
- ASKER, JOHN, CHAIM FERSHTMAN, JIHYE JEON, AND ARIEL PAKES. 2020. "A Computational Framework for Analyzing Dynamic Auctions: The Market Impact of Information Sharing." *RAND Journal of Economics* 51 (3): 805–839.
- ASSAD, STEPHANIE, ROBERT CLARK, DANIEL ERSHOV, AND LEI XU. 2024. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market." *Journal of Political Economy* 132 (3).
- ATER, ITAI, AND OREN RIGBI. 2023. "Price Transparency, Media, and Informative Advertising." *American Economic Journal: Microeconomics* 15 (1): 1–29.
- ATHEY, SUSAN, AND KYLE BAGWELL. 2001. "Optimal Collusion with Private Information." *RAND Journal of Economics*, 428–465.
- . 2008. "Collusion with Persistent Cost Shocks." *Econometrica* 76 (3): 493–540.
- AWAYA, YU, AND VIJAY KRISHNA. 2016. "On Communication and Collusion." *American Economic Review* 106 (2): 285–315.
- . 2019. "Communication and Cooperation in Repeated Games." *Theoretical Economics* 14 (2): 513–553.
- . 2020. "Information Exchange in Cartels." *RAND Journal of Economics* 51 (2): 421–446.
- BACKUS, MATTHEW, CHRISTOPHER CONLON, AND MICHAEL SINKINSON. 2021. *Common Ownership and Competition in the Ready-to-Eat Cereal Industry*. Technical report 28350. National Bureau of Economic Research.
- BERNHEIM, B DOUGLAS, AND MICHAEL D WHINSTON. 1990. "Multimarket Contact and Collusive Behavior." *RAND Journal of Economics*, 1–26.
- BERRY, STEVEN, JAMES LEVINSOHN, AND ARIEL PAKES. 2004. "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market." *Journal of Political Economy* 112 (1): 68–105.
- BERRY, STEVEN T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics* 25 (2): 242–262.
- BERRY, STEVEN T., AND PHILIP A. HAILE. 2014. "Identification in Differentiated Products Markets using Market Level Data." *Econometrica* 82 (5): 1749–1797.
- BERRY, STEVEN T., JAMES A. LEVINSOHN, AND ARIEL PAKES. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–890.
- BHATTACHARYA, VIVEK, GASTÓN ILLANES, AVNER A. KREPS, JOSÉ D. SALAS, AND DAVID STILLERMAN. 2026. "A Large-Scale Evaluation of Merger Simulations." Working paper.
- BLUNDELL, RICHARD, AND STEPHEN BOND. 1998. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models." *Journal of Econometrics* 87 (1): 115–143.
- BONATTI, ALESSANDRO, GONZALO CISTERNAS, AND JUUSO TOIKKA. 2017. "Dynamic Oligopoly with Incomplete Information." *Review of Economic Studies* 84 (2): 503–546.

- BORENSTEIN, SEVERIN. 2004. "Rapid Price Communication and Coordination: The Airline Tariff Publishing Case (1994)." In *The Antitrust Revolution: Economics, Competition, and Policy*, edited by John E Kwoka and Lawrence J White. Oxford University Press.
- BOUSSIOS, DAVID. 2023. *Chicken Market Pricing*. Technical report. Accessed on October 6, 2025. Decision-Next.
- BRESNAHAN, TIMOTHY F. 1982. "The Oligopoly Solution Concept is Identified." *Economics Letters* 10:87–92.
- BYRNE, DAVID P., NICOLAS DE ROOS, A. RACHEL GRINBERG, AND LESLIE M. MARX. 2024. "Informed Sources and the Role of Platforms for Facilitating Anticompetitive Communication." In *Cartels Diagnosed: New Insight on Collusion*, edited by Joseph E. Harrington Jr. and Maarten Pieter Schinkel. Cambridge University Press.
- BYRNE, DAVID P., NICOLAS DE ROOS, MATTHEW S. LEWIS, LESLIE M. MARX, AND XIAOSONG WU. 2025. "Asymmetric Information Sharing in Oligopoly: A Natural Experiment in Retail Gasoline." *Journal of Political Economy* 133 (7): 2031–2088.
- CALDER-WANG, SOPHIE, MITSURU IGAMI, GI HEUNG KIM, AND TAKUO SUGAYA. 2025. "Information Sharing vs. Collusive Coordination: A Bayes Correlated Equilibrium Approach to Algorithmic Pricing in US Multifamily Rentals." Working paper.
- CALDER-WANG, SOPHIE, AND GI HEUNG KIM. 2025. "Coordinated vs. Efficient Prices: The Impact of Algorithmic Pricing on Multifamily Rental Markets." Working paper.
- CARLTON, DENNIS W., ROBERT H GERTNER, AND ANDREW M. ROSENFELD. 1997. "Communication among Competitors: Game Theory and Antitrust Application of Game Theory to Antitrust." *George Mason Law Review* 5:423.
- CHAMBERLAIN, GARY. 1987. "Asymptotic Efficiency in Estimation with Conditional Moment Restrictions." *Journal of Econometrics* 34 (3): 305–334.
- CLARK, ROBERT, AND JEAN-FRANÇOIS HOUDE. 2026. "Third-Party Information Platforms and Coordination: Evidence from the Collapse of a Cartel." Working paper.
- COMPIANI, GIOVANNI. 2022. "Market Counterfactuals and the Specification of Multiproduct Demand: A Nonparametric Approach." *Quantitative Economics* 13 (2): 545–591.
- CONLON, CHRISTOPHER, AND JEFF GORTMAKER. 2020. "Best Practices for Differentiated Products Demand Estimation with PyBLP." *RAND Journal of Economics* 51 (4): 1108–1161.
- . 2025. "Incorporating Micro Data into Differentiated Products Demand Estimation with PyBLP." *Journal of Econometrics*, 105926.
- CORTS, KENNETH S. 1999. "Conduct Parameters and the Measurement of Market Power." *Journal of Econometrics* 88 (2): 227–250.
- DAMODARAN, ASWATH. 2021. *Costs of Capital by Industry (in US \$)*. Available at <https://pages.stern.nyu.edu/~adamodar/>. Accessed on March 4, 2026.
- DEARING, ADAM, LORENZO MAGNOLFI, DANIEL QUINT, CHRISTOPHER SULLIVAN, AND SARAH WALDFOGEL. 2025. *Learning Firm Conduct: Pass-Through as a Foundation for Instrument Relevance*. Technical report. National Bureau of Economic Research.
- DEPARTMENT OF JUSTICE. 2026. *Statement of Interest of the United States of America*. In Re Frozen Potato Products Antitrust Litigation, No. 1:24-cv-11801, N.D. ill. February 27, 2026.
- DOYLE, MAURA P., AND CHRISTOPHER M. SNYDER. 1999. "Information Sharing and Competition in the Motor Vehicle Industry." *Journal of Political Economy* 107 (6): 1326–1364.
- DUARTE, MARCO, LORENZO MAGNOLFI, MIKKEL SØLVSTEN, AND CHRISTOPHER SULLIVAN. 2024. "Testing Firm Conduct." *Quantitative Economics* 15 (3): 571–606.
- ERICSON, RICHARD, AND ARIEL PAKES. 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies* 62 (1): 53–82.

- FEDERAL RESERVE BANK OF ST. LOUIS. *FRED*. Available at <https://fred.stlouisfed.org/>. Accessed on January 19, 2026.
- FERNÁNDEZ-LÓPEZ, JUANA, SENDRA-NADAL-ESTHER, AND ESTRELLA SYAS-BARBERÁ. 2010. "Slaughtering Equipment and Operations." In *Handbook of Poultry Science and Technology, Volume 1: Primary Processing*, edited by Isabel Guerrero-Legarreta. Hoboken: John Wiley & Sons, Inc.
- FUDENBERG, DREW, DAVID LEVINE, AND ERIC MASKIN. 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica* 62 (5): 997–1039.
- GANDHI, AMIT, AND JEAN-FRANÇOIS HOUE. 2020. *Measuring Substitution Patterns in Differentiated-Products Industries*. Technical report 26375. National Bureau of Economic Research.
- GARROD, LUKE, AND MATTHEW OLCZAK. 2018. "Explicit vs. Tacit Collusion: The Effects of Firm Numbers and Asymmetries." *International Journal of Industrial Organization* 56:1–25.
- GREEN, EDWARD J., AND ROBERT H. PORTER. 1984. "Noncooperative Collusion Under Imperfect Price Information." *Econometrica* 52 (1): 87–100.
- GREMBI, VERONICA, TOMMASO NANNICINI, AND UGO TROIANO. 2016. "Do Fiscal Rules Matter?" *American Economic Journal: Applied Economics*, 1–30.
- GRENNAN, MATTHEW, AND ASHLEY SWANSON. 2020. "Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining." *Journal of Political Economy* 128 (4): 1234–1268.
- GROSSEN, GRACE. 2023. *Avian Flu Losses in 2022 Affected Turkey Breast Prices More than Whole Hens*. <https://ers.usda.gov/data-products/charts-of-note/chart-detail?chartId=106639>. Accessed on January 19, 2026.
- HALL, ALASTAIR R, AND DENIS PELLETIER. 2011. "Nonnested Testing in Models Estimated via Generalized Method of Moments." *Econometric Theory* 27 (2): 443–456.
- HARJANI, SUNIL R. 2025. *Memorandum Opinion and Order*. In Re Turkey Antitrust Litigation, No. 1:2019-cv-08318, N.D. Ill. January 22, 2025.
- HARRINGTON, JOSEPH E. 2026. "A Critique of Recent Remedies for Third-Party Pricing Algorithms and Why the Solution Is Not Restrictions On Data Sharing." *Journal of Competition Law & Economics* 22 (1): 114–137.
- HARRINGTON, JOSEPH E., AND MYONG-HUN CHANG. 2009. "Modeling the Birth and Death of Cartels with an Application to Evaluating Competition Policy." *Journal of the European Economic Association* 7 (6): 1400–1435.
- HARRINGTON, JOSEPH E., AND ANDRZEJ SKRZYPACZ. 2007. "Collusion Under Monitoring of Sales." *RAND Journal of Economics* 38 (2): 314–331.
- . 2011. "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices." *American Economic Review* 101 (6): 2425–2449.
- HOLDHEIM, SACHIN, AND ZAAKIR TAMEEZ. 2022. "Danger in the Safety Zone: Information Sharing in the Meat Processing Industry." In *Reforming America's Food Retail Markets*, 11–23.
- IVALDI, MARC, BRUNO JULLIEN, PATRICK REY, PAUL SEABRIGHT, AND JEAN TIROLE. 2003. *The Economics of Tacit Collusion*. Technical report. DG Competition, European Commission.
- KANDORI, MICHIIHIRO. 2002. "Introduction to Repeated Games with Private Monitoring." *Journal of Economic Theory* 102 (1): 1–15.
- KAPETANIOS, GEORGE. 2008. "A Bootstrap Procedure for Panel Data Sets with Many Cross-Sectional Units." *Econometrics Journal* 11 (2): 377–395.
- KILTS CENTER. *NielsenIQ Consumer Panel and Retail Scanner Data*. Available at <https://www.chicagoboth.edu/research/kilts/research-data/nielseniq>. Accessed on January 19, 2026.
- KÜHN, KAI UWE, AND XAVIER VIVES. 1995. *Information Exchanges Among Firms and Their Impact on Competition*. Technical report. European Commission, Directorate-General for Competition.

- LEONARD, CHRISTOPHER. 2017. "Is the Chicken Industry Rigged? Inside Agri Stats, the Poultry Business's Secretive Info-Sharing Service." Available at <https://www.bloomberg.com/news/features/2017-02-15/is-the-chicken-industry-rigged>. Accessed on June 3, 2024. *Bloomberg*.
- LUCO, FERNANDO. 2019. "Who Benefits from Information Disclosure? The Case of Retail Gasoline." *American Economic Journal: Microeconomics* 11 (2): 277–305.
- MAPLES, WILLIAM E., B. WADE BRORSEN, WILLIAM F. HAHN, MATTHEW MACLACHLAN, AND LEKHNATH CHALISE. 2022. *Structure of the USDA Livestock and Poultry Baseline Model*. Technical report. USDA Economic Research Service Technical Bulletin No. 1956.
- MASKIN, ERIC, AND JEAN TIROLE. 2001. "Markov Perfect Equilibrium: I. Observable Actions." *Journal of Economic Theory* 100 (2): 191–219.
- MILLER, NATHAN H., AND MATTHEW C. WEINBERG. 2017. "Understanding the Price Effects of the Miller-Coors Joint Venture." *Econometrica* 85 (6): 1763–1791.
- MONTAG, FELIX, ALINA SAGIMULDINA, AND CHRISTOPH WINTER. 2026. "When Does Mandatory Price Disclosure Lower Prices? Evidence from the German Fuel Market." *Journal of Political Economy: Microeconomics* 4 (1): 106–146.
- MORROW, W ROSS, AND STEVEN J SKERLOS. 2011. "Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand." *Operations Research* 59 (2): 328–345.
- NEVO, AVIV. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica* 69 (2): 307–342.
- NOVSHEK, WILLIAM, AND HUGO SONNENSCHN. 1982. "Fulfilled Expectations Cournot Duopoly with Information Acquisition and Release." *Bell Journal of Economics* 13 (1): 214–218.
- PORTER, ROBERT H. 1983. "A Study of Cartel Stability: The Joint Executive Committee, 1880-1886." *Bell Journal of Economics* 14 (2): 301–314.
- RADNER, ROY, ROGER MYERSON, AND ERIC MASKIN. 1986. "An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria." *Review of Economic Studies* 53 (1): 59–69.
- RAITH, MICHAEL. 1996. "A General Model of Information Sharing in Oligopoly." *Journal of Economic Theory* 71:260–288.
- RIVERS, DOUGLAS, AND QUANG VUONG. 2002. "Model Selection Tests for Nonlinear Dynamic Models." *Econometrics Journal* 5 (1): 1–39.
- ROBINSON, PETER M. 1988. "Root-N-Consistent Semiparametric Regression." *Econometrica*, 931–954.
- ROODMAN, DAVID. 2009. "How to Do xtabond2: An Introduction to Difference and System GMM in Stata." *Stata Journal* 9 (1): 86–136.
- ROSSI, FEDERICO, AND PRADEEP K. CHINTAGUNTA. 2016. "Price Transparency and Retail Prices: Evidence from Fuel Price Signs in the Italian Highway System." *Journal of Marketing Research* 53:407–423.
- ROTEMBERG, JULIO J., AND GARTH SALONER. 1986. "A Supergame-Theoretic Model of Price Wars During Booms." *American Economic Review* 76 (3): 390–407.
- SCARCELLA, MIKE. 2026. "Data Company Agri Stats Settles Meat Workers' Wage-Fixing Lawsuit." Available at <https://www.reuters.com/legal/litigation/data-company-agri-stats-settles-meat-workers-wage-fixing-lawsuit-2026-01-09/>. Accessed on January 15, 2026. *Reuters*.
- STIGLER, GEORGE J. 1964. "A Theory of Oligopoly." *Journal of Political Economy* 72 (1): 44–61.
- SUGAYA, TAKUO. 2022. "Folk Theorem in Repeated Games with Private Monitoring." *Review of Economic Studies* 89 (4): 2201–2256.
- SUGAYA, TAKUO, AND ALEXANDER WOLITZKY. 2018a. "Bounding Payoffs in Repeated Games with Private Monitoring: n -Player Games." *Journal of Economic Theory* 175:58–87.
- . 2018b. "Maintaining Privacy in Cartels." *Journal of Political Economy* 126 (6): 2569–2607.
- . 2023. "Monitoring versus Discounting in Repeated Games." *Econometrica* 91 (5): 1727–1761.

- SWEETING, ANDREW, XUEZHEN TAO, AND XINLU YAO. 2024. “Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Horizontal Mergers.” *American Economic Journal: Microeconomics* 16 (3): 345–373.
- TIERNEY, LUKE, AND JOSEPH B KADANE. 1986. “Accurate Approximations for Posterior Moments and Marginal Densities.” *Journal of the American Statistical Association* 81 (393): 82–86.
- TUNHEIM, JOHN R. 2026. *Memorandum Opinion and Order Denying Agri Stats, Inc.’s Motion for Summary Judgment and Denying Motions to Exclude Expert Testimony*. US et al. v. Agri Stats, Inc., No. 0:23-cv-03009-JRD-JFD, D. Minn. February 24, 2026.
- US ET AL. V. AGRI STATS, INC. 2026. *Proposed Final Judgment*. No. 0:23-cv-03009-JRD-JFD, D. Minn. May 7, 2026.
- US V. REALPAGE, INC. 2025. *Proposed Final Judgment*. No. 1:24-cv-00710-WO-JLW, M.D.N.C. November 24, 2025.
- USDA. *Meat, Poultry and Egg Product Inspection Directory*. Available at <https://www.fsis.usda.gov/inspection/establishments/meat-poultry-and-egg-product-inspection-directory>. Accessed on January 19, 2026.
- . *Poultry Products Inspection Act*. Available at <https://www.fsis.usda.gov/policy/food-safety-acts/poultry-products-inspection-act>. Accessed on January 15, 2026.
- VIVES, XAVIER. 2002. “Private Information, Strategic Behavior, and Efficiency in Cournot Markets.” *RAND Journal of Economics*, 361–376.
- WELSHANS, KRISSA. 2017. “Georgia Ag Department Cancels New Poultry Price Index.” Available at <https://www.feedstuffs.com/livestock-and-poultry-market-news/georgia-ag-department-cancels-new-poultry-price-index>. Accessed on January 12, 2026. *Feedstuffs*.
- ZHENG, XIAOYONG, AND TOMISLAV VUKINA. 2007. “Efficiency Gains from Organizational Innovation: Comparing Ordinal and Cardinal Tournament Games in Broiler Contracts.” *International Journal of Industrial Organization* 25 (4): 843–859.

Online Appendix for “Information Sharing and Collusion: Theory and Evidence from Poultry Processing”

A Additional figures and tables

This section contains additional figures and tables that are referenced in the main text.

1. Figure A.1 displays direct and equilibrium effects of information sharing under different assumptions about primitives. Panel (a) shows direct effects for firm $f' \neq f$ using the same parameterization as in Section 2, where the only uncertainty is over firm f 's demand shock. In this case, firm f' increases its price in the low-demand state (which results in a higher residual demand for firm f') and decreases it in the high-demand state. Thus, the predictions of how prices and welfare are reallocated across demand states are opposite, although the CS asymmetry across states continues to be harmful to expected CS, as firm f' raises its price when its residual demand is higher.

Panel (b) shows the equilibrium effect for firm f under a different data-generating process for ξ . Now, without information sharing, firm f 's prior is that all values of ξ are low or all values of ξ are high, with uniform probability; while sharing information fully informs firm f of the realization. In this case, the direct and equilibrium effect reinforce each other: the direct effect incentivizes both firm f and its rivals to raise their prices, meaning q_f increases in the high-demand state and decreases in the low-demand state relative to no sharing, so the equilibrium effect causes firm f to raise its price by even more in the high-demand state and lower it by even more in the low-demand state.

2. Figure A.2 contains the evolution of prices and quantities of chicken products over time in the two modules I study.
3. Figure A.3 contains the results of the triple-differences event study that benchmarks the turkey event study against the placebo chicken event study. This figure shows the estimates and 95% confidence interval of β_t from the regression

$$\log(p_{jstm}) = \alpha_{jr} + \gamma_{mt} + \mu_t AS_{jm} + \beta_t AS_{jm} \text{Turkey}_m + \varepsilon_{jstm} \quad (24)$$

where the m subscript refers to the meat type (turkey or chicken), and Turkey_m is an

indicator for turkey products. β_t is normalized to zero in May 2019.

4. Figure A.4 shows a kernel density plot with the unweighted distribution of estimated own-price elasticities at realized prices, separately for fresh and deli turkey.
5. Figure A.5 shows the distribution of diagonal elements of the estimated timeseries slope matrices.
6. Figure A.6 shows the distribution of off-diagonal elements of the estimated timeseries slope matrices, separately by whether the entry corresponds to products made by the same firm. The vertical diagonal lines represent the means of these distributions.
7. Figure A.7 is a scatter plot containing predictions of $g(w_{jrt}, z_{jrt}^{GH})$ on the horizontal axis plotted against actual prices on the vertical axis, with a nonparametric line of best fit. A linear regression of actual on predicted prices yields an R^2 of 92.4%.
8. Figure A.8 binscatters the expectation of firms' beliefs about their products' own-price elasticity against the actual own-price elasticity given realized prices. As in the main text, I simulate the expected value of firms' beliefs about their own-price elasticity by taking draws from its beliefs over $(\xi, p-f)$, computing the elasticity, and taking the mean across draws.
9. Figure A.9 shows the distribution of average price upper bounds across markets separately for participants and non-participants of partial information exchanges, and by fresh and deli turkey.
10. Table A.1 presents the decomposition of Agri Stats' overall effects into information and conduct effects, separately for fresh and deli turkey.




Figure A.1: The effect of primitives on information sharing

Note: Both figures are based on logit duopolies, with indirect utilities given by $u_{if} = \xi_f - 3p_f + \varepsilon_{if}$ and both firms having zero marginal costs. Panel (a) shows the direct effect on firm 2's prices from learning about ξ_1 , when $\xi_1 \in \{2, 5\}$ with a uniform prior and $\xi_2 = 4$ with probability 1. Panel (b) shows the equilibrium effect on firm 1's prices from learning demand perfectly when $\xi \in \{(0, 0), (7, 9)\}$ with a uniform prior. The shaded regions are the changes in consumer surplus from the no-information benchmark in each state, with CS losses from information sharing shaded in red and CS gains shaded in green.




Figure A.2: Average chicken price and sales quantity over time

Note: Average prices are weighted by total within-UPC-store quantity throughout the time period.




Figure A.3: Triple-differences estimate of the Agri Stats effect

Note: Estimates come from a regression of log prices on UPC-region fixed effects, month-meat and month-Agri Stats indicator fixed effects, and a month-turkey indicator-Agri Stats indicator interaction, which are plotted along with their 95% confidence interval, with standard errors clustered by UPC-region. Regression is weighted by total within-UPC-region quantity throughout the time period. The vertical dashed line represents May 2019.




Figure A.4: Distribution of estimated own-price elasticities




Figure A.5: Distribution of diagonal elements of B^ξ and B^p




Figure A.6: Distribution of off-diagonal elements of B^ξ and B^p

Note: To be included, the off-diagonal must correspond to a pair of products with more than five markets in common. Off-diagonal elements of B^p only include coefficients of lagged prices on prices, not lagged demand shocks.




Figure A.7: Prices predicted by observables in timeseries model

Note: The red dashed line is the 45-degree line and the blue solid line represents the fitted predictions based on a generalized additive model with a cubic spline.






Figure A.8: Expected own-price elasticity against actual own-price elasticity

Note: Observations are divided into 25 equally sized bins based on the expected own-price elasticity implied by their producer's beliefs, with average expected and actual own-price elasticities then computed within bin. Expected own-price elasticities were computed using a quasi-Monte Carlo simulation with 1,000 draws for each firm-market. The red dashed line is the 45-degree line and the blue solid line is the line of best fit.



(a) Participants vs. non-participants



(b) Fresh vs. deli turkey

Figure A.9: Average price upper bounds by participation status and module

Note: I compute the upper bound on average prices within each market (region-module), then plot the distribution across markets. The whiskers are the 10th and 90th percentiles, the edges of the box are the 25th and 75th percentiles, the middle line is the median, and the point is the average. Participants are involved in all exchanges, while non-participants are only involved in those with full participation.

Scenario	Average prices			Profits (\$m/year)			CS (\$m/year)	TS (\$m/year)
	AS	Others	Total	AS	Others	Total		
<i>A. Fresh turkey</i>								
No sharing (actual)	3.54	3.61	3.57	129.0	53.5	182.5	265.6	448.1
Information effect	3.44	3.60	3.49	126.7	52.1	178.8	277.0	455.8
Info + conduct effect	3.68	3.64	3.66	135.5	54.6	190.1	254.2	444.3
<i>B. Deli turkey</i>								
No sharing (actual)	6.09	6.10	6.09	81.1	28.1	109.2	130.5	239.8
Information effect	6.10	6.10	6.10	81.2	28.1	109.4	130.3	239.7
Info + conduct effect	6.29	6.10	6.24	82.0	28.6	110.6	124.2	234.8

Table A.1: Agri Stats effects by module

Note: Average prices are weighted by no-sharing (actual) total quantity. In all scenarios except no sharing, Agri Stats clients share prices, quantities, costs, and inventories, and other producers just observe their own prices and quantities. In the “Information effect” scenario, firms maximize own-expected profits. In the “Info + conduct effect” scenario, Agri Stats clients maximize (17) using the estimated conduct parameter, while other producers continue to maximize own-expected profits.

B Data and computation details

B.1 Data processing

This section provides additional details of the data processing pipeline. I begin with scanner data from two NielsenIQ product modules: “fresh meat” and “lunchmeat-sliced-refrigerated.” I identify UPCs pertaining to chicken and turkey products using string searches on the native UPC descriptions; I exclude UPCs that contain multiple kinds of meats, such as ham and turkey deli combo packs. I also use the UPC descriptions to define NielsenIQ data fields in a consistent manner, since private-label UPCs typically do not have populated product description fields.

I aggregate prices and sales quantities by the pound. While most UPCs have their sales units denoted in pounds, some UPCs have units denominated in counts rather than in weight. I impute these products’ weight based on their price by first computing the average per-unit price of each weight of product with at least a 1% sales share in the same month-region-retailer combination, and then imputing each UPC’s weight as the one that minimizes the mean-squared difference in that product’s price and the average price of products of that size across month-region-retailers in which that UPC is found. I then aggregate the data to the UPC-month-store level by computing sales-weighted average prices. From there, I remove all UPCs made by firms that are not linked to a poultry processing plant, as well as all

stores that are not featured in the data in each year from 2017 to 2022, and deflate all prices to May 2019 values using the CPI-U. This results in the dataset used for the reduced-form analyses.

I next geocode all poultry processing plants and find the driving distance between them and each store in which a product from that poultry processing plant’s owner was sold. The most granular geographic information about stores in the NielsenIQ data is their three-digit ZIP code, or ZIP3; I compute distances to the centroid of the store’s ZIP3, unless the centroid does not lie close enough to a road for the Google Maps API to compute driving distances to it, in which case I manually adjust the point to the closest road to the centroid. I then compute driving distances between each plant-centroid pair using the Google Maps API, and compute the driving distance for each firm to each store under the assumption that firms supply stores using the nearest available plant. In keeping with the relatively fast spoilage of meat products, I assume that firms can only use the set of licensed plants in a given month to supply sales in that month; that is, if a firm opens or closes a new plant, that affects the driving distances contemporaneously. Occasionally, firms have sales before they have a processing plant or after their processing plants drop out of the data; in these cases, I take the set of available processing plants as their first or last one in the data respectively.

To move to the dataset used for structural estimation, I start by collapsing to the region-product-month level, where products are defined as in the main text, and removing all regions without at least one inside product purchase in each month. I then remove all region-products that never have an inside share of at least 1% in any month, and remove region-product-months with an inside share of less than 0.01%, as these are likely products that are either discontinued mid-month or not widely offered. I define Gandhi-Houde instruments as in the main text, using a random forest to compute the exogenous price variable $\mathbb{E}[p_{jrt} \mid x_{jrt}, z_{jrt}]$. I then identify inside purchases in the panel data and construct sample weights for computing micro-moments using NielsenIQ’s native projection weights, accounting for multiple panelist purchases in a given month by splitting a panelist’s weight across their purchases.

B.2 Estimation

B.2.1 Timeseries

Here I present more detail on my implementation of the [Blundell and Bond \(1998\)](#) estimator for the timeseries estimation of (14). I follow the [Sigmund and Ferstl \(2021\)](#) extension to panel VAR settings. The demand shock relation to estimate is

$$\xi_{rt} = \bar{\xi}_r + B^\xi \xi_{r(t-1)} + e_{rt}^\xi. \quad (25)$$

The [Blundell and Bond \(1998\)](#) estimator combines moments from two equations: a transformed (difference) equation instrumented by lagged levels of the outcome, and a level equation instrumented by lagged differences of the outcome. For any variable x , let Δx_{rt} denote its forward orthogonal deviation ([Sigmund and Ferstl, 2021](#)), and let $\mathbb{E}[\cdot]$ denote the expectation across regions r . Because ξ_{rt} and e_{rt}^ξ are J -vectors (where J is the number of products), each condition below is a $J \times J$ matrix of moments, with one scalar moment for each pairing of an instrument product and an outcome equation. The valid moments are then, for any period t and any lag $\ell \geq 1$,

$$\begin{aligned}\mathbb{E} \left[\xi_{r(t-\ell)} (\Delta e_{rt}^\xi)' \right] &= 0, \quad (\text{difference equation}) \\ \mathbb{E} \left[\Delta \xi_{r(t-\ell)} (e_{rt}^\xi)' \right] &= 0. \quad (\text{level equation})\end{aligned}\tag{26}$$

The instruments for the difference equation expand triangularly in the number of time periods, and many are weak ([Roodman, 2009](#)). I therefore restrict to the first two lags (i.e., $\ell \leq 2$) and collapse these moments following [Roodman \(2009\)](#), summing each lag's moment across the periods t at which it applies. The level equation instead uses only the most recent lagged difference as an instrument in each period, so its moments grow linearly rather than triangularly in the number of periods and require no such reduction. The final set of moments I use are the sample analogues of

$$\begin{aligned}\sum_t \mathbb{E} \left[\xi_{r(t-\ell)} (\Delta e_{rt}^\xi)' \right] &= 0, \quad \ell \in \{1, 2\} \\ \mathbb{E} \left[\Delta \xi_{r(t-1)} (e_{rt}^\xi)' \right] &= 0 \quad \forall t.\end{aligned}\tag{27}$$

I estimate $\bar{\xi}_r$ and B^ξ using two-step GMM. While the direct implementation requires a balanced panel, the presence of regional products (and, to a lesser extent, changes in product assortment over time) causes observation of an unbalanced panel. As recommended by [Roodman \(2009\)](#), I balance the panel by imputing missing values of ξ_{rt} as zeros, which results in valid moments.

To estimate Σ_{rt}^ξ , I first compute fitted residuals $\hat{e}_{rt}^\xi$ given the panel estimates, and compute pairwise covariances of $\hat{e}_{jrt}^\xi, \hat{e}_{krt}^\xi$ for all pairs of products (j, k) , treating the covariance as zero if the two products have fewer than two markets in common. As this may not result in a positive definite covariance matrix due to differences in product assortment across markets, I then project this matrix onto the space of positive definite matrices using the following procedure. Letting $\bar{\Sigma}^\xi$ be the matrix of pairwise covariances, I eigendecompose $\bar{\Sigma}^\xi = Q\Lambda Q'$. I then floor eigenvalues in Λ with a very small positive number, the result of which is Λ_+ ,

and back out $\widehat{\Sigma}^\xi = Q\Lambda_+Q'$. This gives the closest positive definite and symmetric covariance matrix to $\overline{\Sigma}^\xi$ in the Frobenius norm, which I then take into the remaining estimation steps.

Computation of the price timeseries follows similar steps, with three differences. First, the contemporaneous price depends on other observables nonlinearly through the function $g(w_{rt}, z_{rt}^{GH})$. I use the semiparametric two-step estimator of [Robinson \(1988\)](#) by residualizing p_{rt} and $\xi_{r(t-1)}$ against nonparametric functions of (w_{rt}, z_{rt}^{GH}) using a random forest and computing the [Blundell and Bond \(1998\)](#) estimator on the residuals. I also add a time trend and an Agri Stats indicator, although the latter is eventually absorbed by the product-DMA fixed effect $\bar{p}_r$. Second, there is the additional linear component in the slope matrix for $\xi_{r(t-1)}$. I instrument this regressor in the level equation with its lagged differences, adding the moment $\mathbb{E}[\Delta\xi_{r(t-1)}(e_{rt}^p)'] = 0$ for all t . Third, I estimate the relation separately before and after May 2019, as I assume two different Markov equilibria in pricing rules in each period.

B.2.2 Beliefs

To estimate beliefs, I first derive the stationary beliefs implied by the timeseries estimates. I treat (ξ, p) jointly as one random vector and write $B = \begin{pmatrix} B^\xi & 0 \\ B_\xi^p & B_p^p \end{pmatrix}$ as the matrix of VAR slope coefficients and $\Sigma = \text{BlockDiag}(\Sigma^\xi, \Sigma^p)$ as the covariance matrix of the innovations of that vector. The distribution at the very beginning of the sample is given by $(\xi_{r1}, p_{r1}) \sim \mathcal{N}(\mu_{r1}, \Sigma_{r1})$, where

$$\mu_{r1} = \begin{pmatrix} \bar{\xi}_r \\ \bar{p}_r + g(w_{r1}, z_{r1}^{GH}) \end{pmatrix}$$

$$\Sigma_{r1} = \{\Sigma_r : B\Sigma_r B' - \Sigma_r + \Sigma = 0\}$$

where Σ_{r1} solves the discrete-time Lyapunov equation. I seed initial beliefs for all firms in the first period as $\mu_{fr1} = \mu_{r1}$, $\Sigma_{fr1} = \Sigma_{r1}$. Then, fixing a region r and for all periods $t = 1, \dots, T$, where t^* corresponds to May 2019 and T corresponds to May 2022, I use the following procedure to iteratively compute beliefs.

1. For firm f producing in market rt , let $\mathcal{F}_{f,rt}^{\text{obs}} = \{f' : f' \in AS \setminus \{f\}\}$ if $f \in AS$ and $t \leq t^*$; otherwise, let $\mathcal{F}_{f,rt}^{\text{obs}} = \emptyset$.
2. For all $j = 1, \dots, J_{rt}$ such that $j \in \mathcal{J}_{f',rt}$ for an $f' \in \mathcal{F}_{f,rt}^{\text{obs}} \cup \{f\}$: replace the $(J+j)$ th element of $\mu_{f,rt}$ with $p_{j,rt}$ and the $(J+j)$ th row and column of $\Sigma_{f,rt}$ with zero. This corresponds to the firm knowing for sure that the price of good j will equal $p_{j,rt}$. (J is added because the first J components correspond to values of ξ .)

3. If $f \in AS$ and $t \leq t^*$, define the signal functions $y_f(\xi_{rt}, p_{rt}) = (p_{AS,rt}, s_{AS,rt}(p_{rt}, \xi_{rt}))$; otherwise, define $y_f(\xi_{rt}, p_{rt}) = (p_{f,rt}, s_{f,rt}(p_{rt}, \xi_{rt}))$. Define the realizations of the signal analogously.
4. Given $(\mu_{f,rt}, \Sigma_{f,rt}, y_{f,rt})$, compute $(\mu_{f,rt}^{\text{post}}, \Sigma_{f,rt}^{\text{post}})$ using Laplace's method, according to which

$$\begin{aligned}\mu_{f,rt}^{\text{post}} &= \arg \min_{(\xi, p)} ((\xi, p) - \mu_{f,rt})' \Sigma_{f,rt}^+ ((\xi, p) - \mu_{f,rt}) \quad \text{s.t.} \quad y_f(\xi, p_{-f}) = y_{f,rt} \\ \Sigma_{f,rt}^{\text{post}} &= U(U' \Sigma_{f,rt}^+ U)^{-1} U'\end{aligned}$$

where (with an obvious abuse of notation) $y_f(\xi, p_{-f}) = y_f(p_f, p_{-f}, s_{f,rt}(p_f, p_{-f}, \xi))$ is firm f 's signal as a function of ξ and p_{-f} , and U is an orthonormal basis of the tangent space of the constraint manifold $\{(\xi, p_{-f}) : y_f(\xi, p_{-f}) = y_{f,rt}\}$, evaluated at $\mu_{f,rt}^{\text{post}}$. Note that the pseudoinverse $\Sigma_{f,rt}^+$ is used because $\Sigma_{f,rt}$ is generally singular; all rows and columns corresponding to prices firm f knows in advance (which include at least its own prices) will be zero. Here, the pseudoinverse corresponds to taking the inverse of the submatrix associated with nondegenerate elements of firm f 's prior, while keeping the degenerate elements equal to zero.

5. Compute $(\mu_{f,rt(t+1)}, \Sigma_{f,rt(t+1)})$ according to the following:

$$\begin{aligned}\mu_{f,rt(t+1)} &= \begin{pmatrix} \bar{\xi}_r \\ \bar{p}_r + g(w_{r(t+1)}, z_{r(t+1)}^{GH}) \end{pmatrix} + B \mu_{f,rt}^{\text{post}} \\ \Sigma_{f,rt(t+1)} &= B \Sigma_{f,rt}^{\text{post}} B' + \Sigma\end{aligned}$$

If a new firm enters a particular region after period 1, I initialize that firm's beliefs at the stationary beliefs, unless that firm is an Agri Stats client and it is before May 2019, in which case I assign it the same beliefs as all other Agri Stats clients in that region in that period.

Finally, given that beliefs are persistent over time, one approximation error or particularly strange set of observed values can affect all future downstream beliefs. To lessen the chances of this happening, in each period I winsorize each firm's prior beliefs such that their mean beliefs over a particular variable x (where x is an element of either p or ξ) fall between $[\underline{x} - 3\sigma_x, \bar{x} + 3\sigma_x]$, where $\underline{x}$ and $\bar{x}$ are the minimum and maximum values of x in the data, and σ_x is its sample standard deviation. For prices, I additionally floor the mean belief at zero.

B.2.3 Marginal costs

Once firm f 's beliefs are computed, we can recover its costs under Bertrand-Markov from its first-order conditions, (13). For a given firm f in market rt , I first take M draws $(\xi_{rt}^m, p_{rt}^m)_{m=1}^M$ from $\mathcal{N}(\mu_{f,rt}, \Sigma_{f,rt})$ using quasi-Monte Carlo methods (shifted Halton draws with antithetic variates). I then find $c_{f,rt}$ using

$$c_{f,rt} = p_{f,rt} + \left(\frac{1}{M} \sum_{m=1}^M D_{p_f} s_{f,rt}(p_{rt}^m; \xi_{rt}^m)' \right)^{-1} \left(\frac{1}{M} \sum_{m=1}^M s_{f,rt}(p_{rt}^m; \xi_{rt}^m) \right) \quad (28)$$

where $D_{p_f} s_{f,rt}(p_{rt}^m; \xi_{rt}^m) = \left[\frac{\partial s_{j,rt}(p_{rt}^m; \xi_{rt}^m)}{\partial p_{k,rt}} \right]_{j,k \in \mathcal{J}_{f,rt}}$. Applying this formula for all f, rt such that $t > t^*$ gives all post-May 2019 costs. To compute pre-May 2019 costs under a particular value of τ for Agri Stats subscribers, I use the fact that Agri Stats clients have common priors while sharing information by solving for their costs jointly, according to

$$c_{AS,rt} = p_{AS,rt} + \left(\mathcal{H}_{AS,rt}(\tau) \odot \frac{1}{M} \sum_{m=1}^M D_{p_{AS}} s_{AS,rt}(p_{rt}^m; \xi_{rt}^m)' \right)^{-1} \left(\frac{1}{M} \sum_{m=1}^M s_{AS,rt}(p_{rt}^m; \xi_{rt}^m) \right) \quad (29)$$

where $\mathcal{H}_{AS,rt}(\tau)$ is the within-Agri Stats ownership matrix—element jk equals 1 if products j and k are produced by the same firm, and τ otherwise—and $\odot$ is the element-by-element multiplication operator. Other firms' costs follow (28).

Finally, to estimate the cost function $h(w_{j,rt})$ and recover the residuals $\omega_{j,rt}$, I first residualize all variables against product, region, and month (i.e., time) fixed effects, and then apply a random forest to the residuals. I do so jointly in the pre- and post-periods, separately for each value of τ .

B.2.4 Bootstrap

There are three sources of asymptotic variance in my estimation pipeline: the demand parameters, the timeseries parameters (which in part depend on the demand parameters, through ξ), and the cost function. To account for asymptotic variation in the demand parameters, I use a parametric bootstrap by pulling draws from the asymptotic distribution of the parameters implied by the two-step GMM robust covariance matrix, and then resolving for the values of ξ that make predicted and actual shares match for each draw of the parameters. Then, for each draw of the parameters, I account for other sources of asymptotic variation by drawing regions with replacement, following [Kapetanios \(2008\)](#)'s approach for

bootstrapping panel estimation routines with timeseries dependence. I construct 500 such joint draws of demand parameters and sets of regions after estimating demand, then run each draw through the preceding estimation steps and conduct testing routine. This process provides the standard errors I report in the main text.

B.3 Counterfactual Markov equilibria in pricing rules

For a given profile of signal functions $\{y_f\}$, computing a counterfactual Markov equilibrium in pricing rules requires finding a long-run price kernel Π_r^+ that is long-run consistent in the sense of (11): if firms conjecture that rivals' prices follow Π_r^+ , set their prices to maximize their expected profits, and update their beliefs given the signals those prices realize, then the induced long-run law of motion of prices must reproduce Π_r^+ . As described in Section 8, this is a fixed point of the map $\Pi_r^+ \mapsto \Phi(\Pi_r^+)$, where Φ returns the long-run law of motion of the expected-profit-maximizing prices generated by the beliefs that the conjectured Π_r^+ induces. I recover this fixed point by iterating on Φ : given a conjectured kernel, I simulate the post-May 2019 prices and beliefs it implies, re-estimate the kernel from those simulated prices, and repeat until the conjectured and re-estimated kernels coincide.

In practice, the parameterization of the price timeseries I use is defined by a tuple $(g(\cdot), \bar{p}_r, B^p, \Sigma^p)$. I use a random forest to estimate $g(\cdot)$, which is not continuous in the inputs, and hence is a bad approach for fixed-point methods. I therefore take the approach of only adjusting $g(\cdot)$ from the estimated function by a constant for each product-region combination. As any constant shift in $g(\cdot)$ is isomorphic to a shift in the product-region fixed effect $\bar{p}_r$, this is equivalent to keeping $g(\cdot)$ fixed in counterfactuals and only recomputing $(\bar{p}_r, B^p, \Sigma^p)$ for each counterfactual iteration. Holding $g(\cdot)$ fixed, I stack the remaining free parameters into a single vector

$$\Pi^+ = (\{\bar{p}_r\}_{r=1}^R, B_\xi^p, B_p^p, \Sigma^p), \quad (30)$$

where $B^p = (B_\xi^p \ B_p^p)$ is partitioned conformably with $(\xi_{r(t-1)}, p_{r(t-1)})$ as in Appendix B.2.2. One evaluation of Φ unpacks Π^+ , simulates the implied equilibrium prices and beliefs market by market, and re-estimates (14) on the simulated prices using system GMM to return $\Phi(\Pi^+)$. A counterfactual equilibrium is a zero of the residual $G(\Pi^+) = \Phi(\Pi^+) - \Pi^+$. Because B^p and Σ^p are pooled across regions, evaluating Φ couples all markets, so the fixed point cannot be solved region by region.

To compute equilibrium prices in a particular market rt given the vector of costs c_{rt} and a system of prior beliefs $\{\mu_{f,rt}, \Sigma_{f,rt}\}_{f \in \mathcal{F}_{rt}}$, I adapt the [Morrow and Skerlos \(2011\)](#) fixed-point iteration for this setting. Decompose $D_{p, s_{rt}}(p_{rt}; \xi_{rt}) = \Lambda_{rt}(p_{rt}, \xi_{rt}) - \Gamma_{rt}(p_{rt}, \xi_{rt})$, where Λ_{rt} is

a diagonal $J \times J$ matrix with the j th element of the diagonal equal to

$$\Lambda_{jrt}(p_{rt}, \xi_{rt}) = \int \alpha_i(\theta) s_{ijrt}(p_{rt}; \xi_{rt}, \theta) dF_{rt}(\theta) \quad (31)$$

and Γ_{rt} is a dense $J \times J$ matrix with jk th element equal to

$$\Gamma_{jkrt}(p_{rt}, \xi_{rt}) = \int \alpha_i(\theta) s_{ijrt}(p_{rt}; \xi_{rt}, \theta) s_{ikrt}(p_{rt}; \xi_{rt}, \theta) dF_{rt}(\theta) \quad (32)$$

where $F_{rt}(\theta)$ is the distribution of structural demand parameters (α_i, β_i) in market rt . One may then rewrite the first-order condition for firm f as $p_{frt} = c_{frt} + \zeta_{frt}(p_{frt})$, where

$$\zeta_{frt}(p_{frt}) = \mathbb{E}_{frt} [\Lambda_{frt}(p_{rt}, \xi_{rt})]^{-1} (\mathbb{E}_{frt} [\Gamma_{frt}(p_{rt}, \xi_{rt})'] (p_{frt} - c_{frt}) - s_{frt}(p_{rt}; \xi_{rt})). \quad (33)$$

Here, $\Lambda_{frt}(\cdot)$ and $\Gamma_{frt}(\cdot)$ are the submatrices of $\Lambda_{rt}(\cdot)$ and $\Gamma_{rt}(\cdot)$ associated with firm f 's products. The advantage of writing it in this way is that ζ_{frt} is a local contraction, and hence iterating on $p_{frt} \leftarrow c_{frt} + \zeta_{frt}(p_{frt})$ reliably converges to prices satisfying firm f 's first-order conditions. To apply the method, I compute expectations with respect to firm f 's beliefs using quasi-Monte Carlo, replacing the draws corresponding to firm f 's prices as appropriate. If a set of firms S share inventories, then I replace $\mathbb{E}_{frt}[\Lambda_{frt}(\cdot)]$ with $\mathbb{E}_{Srt}[\Lambda_{Srt}(\cdot)]$ and $\mathbb{E}_{frt}[\Gamma_{frt}(\cdot)']$ with $\mathbb{E}_{Srt}[\mathcal{H}_{Srt} \odot \Gamma_{Srt}(\cdot)']$, where $\mathcal{H}_{Srt}$ is the within- S ownership matrix and $\mathbb{E}_{Srt}[\cdot]$ is the operator that, in row j , computes the expectation of the matrix with respect to the beliefs of the firm that owns product j . Finally, given a set of equilibrium prices in market rt , I compute posterior beliefs using Laplace's method, and then compute priors in the following period using the estimated law of motion of ξ and the candidate Π^+ , as in Appendix B.2.2. Looping over all markets gives the full set of equilibrium prices, and then $\Phi(\Pi^+)$ is computed using system GMM (Appendix B.2.1).

My baseline solver is a damped fixed-point iteration $\Pi_{k+1}^+ = (1 - q)\Pi_k^+ + q\Phi(\Pi_k^+)$, iterated until the largest change across the $(B_\xi^p, B_p^p, \Sigma^p)$ blocks of Π^+ , or the change in equilibrium prices across iterates, falls below a tolerance. While I have found that dampening ($q < 1$) helps stabilize the iteration at the cost of convergence speed, I have also found that Φ is not guaranteed to be a contraction. When the fixed-point iteration does not result in convergence, I implement a Jacobian-free Newton-Krylov (JFNK) solver to minimize $\frac{1}{2} \|G(\Pi^+)\|^2$ (Knoll and Keyes, 2004). Newton's method computes each step δ_k from

$$(\Phi'(\Pi_k^+) - I) \delta_k = -G(\Pi_k^+) \quad (34)$$

where $\Phi'(\Pi_k^+)$ is the Jacobian of Φ at an iterate Π_k^+ . The following iterate is then $\Pi_{k+1}^+ =$

$\Pi_k^+ + a_k \delta_k$, where a_k is set using a global line-search method. Forming Φ' explicitly is prohibitive, as each of its columns would require a separate pass through the entire market-by-market equilibrium computation. Instead, I solve (34) with GMRES (Saad and Schultz, 1986), which accesses the Jacobian only through matrix-vector products. I approximate each such product with finite differences; for small ε ,

$$(\Phi'(\Pi_k^+) - I)v \approx \frac{\Phi(\Pi_k^+ + \varepsilon v) - \Phi(\Pi_k^+)}{\varepsilon} v. \quad (35)$$

Each GMRES iteration thus costs one additional evaluation of Φ and never requires the full Jacobian to be stored. The convergence criteria are the same as in the fixed-point iteration.

B.4 Upper bound

This appendix details the computation of the collusive upper bound described in Section 8.2. Throughout I condition on a single market and suppress the rt subscript. Each firm f enters with the belief $\xi \sim \mathcal{N}(\mu_f^\xi, \Sigma_f^\xi)$, derived from the stationary belief implied by the timeseries estimates. Thus,

$$\mathbb{E}_f[\pi_f(p)] = \mathbb{E}_{\xi \sim \mathcal{N}(\mu_f^\xi, \Sigma_f^\xi)} [(p_f - c_f)' s_f(p, \xi)] \quad (36)$$

is firm f 's expected profit given price vector p . As in Appendix B.2.2, I approximate all expectations against firm f 's beliefs using a quasi-Monte Carlo average over M draws $\{\xi_f^m\}_{m=1}^M$ from $\mathcal{N}(\mu_f^\xi, \Sigma_f^\xi)$, generated with shifted Halton sequences and antithetic variates. Finally, let $\kappa = \delta/(1 - \delta)$, and let p^C and p^M be the competitive and monopoly prices, which I solve for using the iteration scheme described in Appendix B.3.

B.4.1 The discretized problem

I approximate the occupation measure α by a finite mixture over N price nodes, $\alpha = \{(p_n, a_n)\}_{n=1}^N$, with $p_n \in \mathbb{R}^J$ and weights $a_n \geq 0$ summing to one. Given weights proportional to actual share $\psi_j \geq 0$ with $\sum_j \psi_j = 1$, the program is

$$\alpha^*(y) = \arg \max_{\{(p_n, a_n)\}_{n=1}^N} \sum_{n=1}^N a_n \sum_j \psi_j p_{n,j} \quad (37)$$

$$\begin{aligned} \text{s.t.} \quad & \sum_{n=1}^N a_n = 1, \quad a_n \geq 0, \quad p_n \in [p^C, p^M] \quad \forall n, \\ & \text{NashIC}_{n,f}(\alpha) \leq 0 \quad \forall f, \forall n, \end{aligned} \quad (38)$$

$$\text{SW}_f(\alpha) \leq \kappa \text{Var}_f(\alpha) \quad \forall f. \quad (39)$$

where $\text{Var}_f(\alpha) = \text{Var}_{p \sim \alpha}(\mathbb{E}_f[\pi_f(p)])$. The objective (37) is linear, with gradient $\partial/\partial a_n = \sum_j \psi_j p_{n,j}$ and $\partial/\partial p_{n,j} = a_n \psi_j$. The two constraints are defined below.

B.4.2 Nash IC

Let the punishment value be $\pi_f^C = \mathbb{E}_f[\pi_f(p^C)]$, which is independent of the node. If firm f is recommended price p_n but deviates, its best one-shot deviation profit is

$$\pi_f^D(p_n) = \max_{p_f^D} \mathbb{E}_f[\pi_f(p_f^D, p_{n,-f})], \quad (40)$$

the static best response to $p_{n,-f}$ under f 's belief. The value of not deviating at node n is $(1-\delta)\mathbb{E}_f[\pi_f(p_n)] + \delta \bar{\pi}_f(\alpha)$, where $\bar{\pi}_f(\alpha) = \sum_n a_n \mathbb{E}_f[\pi_f(p_n)]$ is the average equilibrium profit. The constraint is that for every firm f and node n ,

$$\text{NashIC}_{n,f}(\alpha) = (1-\delta) (\pi_f^D(p_n) - \mathbb{E}_f[\pi_f(p_n)]) + \delta (\pi_f^C - \bar{\pi}_f(\alpha)) \leq 0. \quad (41)$$

I solve (40) using a modified [Morrow and Skerlos \(2011\)](#) iteration scheme, where only firm f 's prices are updated in each iteration.

The derivative of expected profits with respect to prices is standard:

$$\frac{\partial \mathbb{E}_f[\pi_f(p)]}{\partial p_j} = \frac{1}{M} \sum_{m=1}^M \left[\mathbb{1}\{j \in \mathcal{J}_f\} s_j(p, \xi^m) + \sum_{k \in \mathcal{J}_f} (p_k - c_k) \frac{\partial s_k(p, \xi^m)}{\partial p_j} \right]. \quad (42)$$

The derivatives of π_f^D are

$$\frac{\partial \pi_f^D(p_n)}{\partial p_{n,j}} = \mathbb{1}\{j \notin \mathcal{J}_f\} \frac{\partial \mathbb{E}_f[\pi_f(p_f^D(p_n), p_{n,-f})]}{\partial p_j} \quad (43)$$

where $p_f^D(p_n)$ is the argmax of (40). The derivative when $j \in \mathcal{J}_f$ is zero by optimality of $p_f^D(p_n)$. Differentiating (41), the gradients of the Nash IC constraints are

$$\begin{aligned} \frac{\partial \text{NashIC}_{n,f}}{\partial a_{n'}} &= -\delta \mathbb{E}_f[\pi_f(p_{n'})], \\ \frac{\partial \text{NashIC}_{n,f}}{\partial p_{n',j}} &= \begin{cases} (1-\delta) \left(\frac{\partial \pi_f^D(p_n)}{\partial p_j} - \frac{\partial \mathbb{E}_f[\pi_f(p_n)]}{\partial p_j} \right) - \delta a_n \frac{\partial \mathbb{E}_f[\pi_f(p_n)]}{\partial p_j}, & n' = n, \\ -\delta a_{n'} \frac{\partial \mathbb{E}_f[\pi_f(p_{n'})]}{\partial p_j}, & n' \neq n. \end{cases} \end{aligned} \quad (44)$$

B.4.3 SW23

A manipulation is a map from each firm's recommended own-price to a deviation price; that is, it is a vector of deviation prices $p'_f = \{p'_{n,f}\}_{n=1}^N$ for each firm.⁴⁸ Because two nodes that share the same own-price $p_{n,f}$ must be manipulated identically, I parameterize $p'_{n,f}$ by one deviation block per distinct own-price among the nodes. Let $k(n)$ index the distinct own-price vector $p_{n,f}$, and let $p'_{k,f}$ be the deviation assigned to block k . Define the per-node and overall gain and detectability

$$\begin{aligned} g_{n,f}(p_n, p'_f) &= \mathbb{E}_f[\pi_f(p'_{k(n),f}, p_{n,-f})] - \mathbb{E}_f[\pi_f(p_n)], & G_f(p'_f, \alpha) &= \sum_{n=1}^N a_n g_{n,f}(p_n, p'_f), \\ c_{n,f}(p_n, p'_f) &= \chi^2(p'_{k(n),f}, p_{n,-f} \parallel p_n; y_f), & C_f(p'_f, \alpha) &= \sum_{n=1}^N a_n c_{n,f}(p_n, p'_f). \end{aligned} \quad (45)$$

The condition of Theorem 1 of SW23, applied to the binding manipulation for each firm, is

$$\max_{p'_f} \frac{G_f(p'_f, \alpha)}{\sqrt{C_f(p'_f, \alpha)}} \leq \sqrt{\kappa \text{Var}_f(\alpha)}. \quad (46)$$

To evaluate the condition as a constraint, I optimize the left-hand side of (46), holding α fixed. Let the solution be given by $p_f^*(\alpha)$. I then square the constraint on both sides, as this seems to perform better numerically. Writing the left-hand side as

$$\text{SW}_f(\alpha) = \frac{G_f(p_f^*(\alpha), \alpha)^2}{C_f(p_f^*(\alpha), \alpha)}, \quad (47)$$

and the right-hand side variance as

$$\text{Var}_f(\alpha) = \sum_{n=1}^N a_n (\mathbb{E}_f[\pi_f(p_n)])^2 - \left(\sum_{n=1}^N a_n \mathbb{E}_f[\pi_f(p_n)] \right)^2, \quad (48)$$

gives the elements of constraint (39).⁴⁹

The signal distribution and its χ^2 -divergence. Let $y_f(s)$ be the signal entering firm f 's SW23 constraint as a function of the vector of shares. Price signals are excluded because

⁴⁸The original result allows for manipulations to be distributions of prices. For computational reasons, I only consider deterministic manipulations. This weakly loosens the upper bound.

⁴⁹In practice, I floor G_f at zero when adding it to the constraint as a safety measure so that small negative values caused by numerical noise do not accidentally result in a positive value when squared. This does not matter in theory, as the solution to (46) is weakly positive.

if the firm shares prices then the SW23 constraint is vacuously satisfied and therefore not included in the program at all. I generally consider linear functions of the shares, so its Jacobian $D_y = \partial y_f / \partial s$ is constant.

I approximate the signal distribution as additive logistic-normal. I choose this approximation because the signals are selections and aggregates of shares, and the conditional distribution of the (plain) logit share vector given a normally distributed ξ is exactly additive logistic-normal. I transform signals by

$$v(y) = \log y - \log(1 - \mathbf{1}'y), \quad (49)$$

where $1 - \mathbf{1}'y_f$ is the residual mass of shares not in the signal. I treat $v \mid p \sim \mathcal{N}(\mu(p), \Sigma(p))$ and estimate its moments via simulation. Let $v^m = v(y(s(p, \xi^m)))$ be the transformed signal given the draw ξ^m . Then $\mu(p) = M^{-1} \sum_{m=1}^M v^m$ and $\Sigma(p) = M^{-1} \sum_{m=1}^M (v^m - \mu(p))(v^m - \mu(p))'$.

Let (μ_d, Σ_d) denote the moments at $(p'_{k(n),f}, p_{n,-f})$ and (μ_n, Σ_n) those at p_n . The χ^2 -divergence between two normals admits a closed form. Define $K = 2\Sigma_d^{-1} - \Sigma_n^{-1}$, $b = 2\Sigma_d^{-1}\mu_d - \Sigma_n^{-1}\mu_n$, $\mu^* = K^{-1}b$, and $q = 2\mu'_d \Sigma_d^{-1} \mu_d - \mu'_n \Sigma_n^{-1} \mu_n - b' \mu^*$. Then the divergence is given by

$$1 + \chi^2(p' \parallel p; y_f) = |\Sigma_n|^{1/2} |\Sigma_d|^{-1} |K|^{-1/2} \exp(-0.5q). \quad (50)$$

If K is not positive definite, then χ^2 diverges to $+\infty$.

Gradients. Let $u_d = \mu_d - \mu^*$ and $u_n = \mu_n - \mu^*$. Differentiating (50),

$$\begin{aligned} \frac{\partial \chi^2}{\partial \mu_d} &= -2(1 + \chi^2) \Sigma_d^{-1} u_d, \\ \frac{\partial \chi^2}{\partial \Sigma_d} &= (1 + \chi^2) \Sigma_d^{-1} (K^{-1} - \Sigma_d + u_d u_d') \Sigma_d^{-1}, \\ \frac{\partial \chi^2}{\partial \mu_n} &= (1 + \chi^2) \Sigma_n^{-1} u_n, \\ \frac{\partial \chi^2}{\partial \Sigma_n} &= 0.5(1 + \chi^2) (\Sigma_n^{-1} - \Sigma_n^{-1} (K^{-1} + u_n u_n') \Sigma_n^{-1}). \end{aligned} \quad (51)$$

The gradient of the ALR transformation is

$$D_v(y)_{k\ell} = \frac{\mathbb{1}(k = \ell)}{y_k} + \frac{1}{1 - \mathbf{1}'y}. \quad (52)$$

Let $\nabla_p s(p, \xi^m)$ be the share derivative from the demand system, and let $D_p v(p, \xi^m) = D_v(y(s(p, \xi^m))) D_y \nabla_p s(p, \xi^m)$. Then, the gradient of the moments entering χ^2 with respect

to prices is

$$\begin{aligned}
D_p \mu(p) &= \frac{1}{M} \sum_{m=1}^M D_p v(p, \xi^m), \\
\frac{\partial \Sigma(p)}{\partial p_j} &= \frac{1}{M} \sum_{m=1}^M \left[(D_p v(p, \xi^m))_{\cdot j} (v^m)' + v^m (D_p v(p, \xi^m))'_{\cdot j} \right] - \\
&\quad (D_p \mu(p))_{\cdot j} \mu(p)' - \mu(p) (D_p \mu(p))'_{\cdot j},
\end{aligned} \tag{53}$$

where $(A)_{\cdot j}$ denotes the j th column of A . Finally,

$$\begin{aligned}
\frac{\partial \chi^2}{\partial p_{n,j}} &= \frac{\partial \chi^2}{\partial \mu_n} \cdot \frac{\partial \mu_n}{\partial p_{n,j}} + \left\langle \frac{\partial \chi^2}{\partial \Sigma_n}, \frac{\partial \Sigma_n}{\partial p_{n,j}} \right\rangle, \\
\frac{\partial \chi^2}{\partial p_{d,j}} &= \frac{\partial \chi^2}{\partial \mu_d} \cdot \frac{\partial \mu_d}{\partial p_{d,j}} + \left\langle \frac{\partial \chi^2}{\partial \Sigma_d}, \frac{\partial \Sigma_d}{\partial p_{d,j}} \right\rangle,
\end{aligned} \tag{54}$$

where $\langle A, B \rangle = \sum_{k\ell} A_{k\ell} B_{k\ell}$ is the Frobenius inner product.

Let $c_{n,f}^*(\alpha) = c_{n,f}(p_n, p_{k(n),f}^*(\alpha))$ and $g_{n,f}^*(\alpha) = g_{n,f}(p_n, p_{k(n),f}^*(\alpha))$ be the components of the solutions to the LHS of (46). Then,

$$\begin{aligned}
\frac{\partial c_{n,f}^*(\alpha)}{\partial p_{n,j}} &= \frac{\partial \chi^2}{\partial p_{n,j}} + \mathbb{1}\{j \notin \mathcal{J}_f\} \frac{\partial \chi^2}{\partial p_{d,j}}, \\
\frac{\partial g_{n,f}^*(\alpha)}{\partial p_{n,j}} &= \mathbb{1}\{j \notin \mathcal{J}_f\} \frac{\partial \mathbb{E}_f[\pi_f(p_{k(n),f}^*(\alpha), p_{n,-f})]}{\partial p_{n,j}} - \frac{\partial \mathbb{E}_f[\pi_f(p_n)]}{\partial p_{n,j}}.
\end{aligned} \tag{55}$$

The gradients with respect to the weights $\{a_n\}$ are 0, as each inner problem only depends on the node locations.

Condition on $\{a_n, p_n\}$ and let $G_f^* = \sum_n a_n g_{n,f}^*$ and $C_f^* = \sum_n a_n c_{n,f}^*$. The gradients of SW_f and Var_f are

$$\frac{\partial \text{SW}_f}{\partial a_n} = 2 \frac{G_f^*}{C_f^*} g_{n,f}^* - \left(\frac{G_f^*}{C_f^*} \right)^2 c_{n,f}^*, \tag{56}$$

$$\frac{\partial \text{SW}_f}{\partial p_{n,j}} = a_n \left[2 \frac{G_f^*}{C_f^*} \frac{\partial g_{n,f}^*}{\partial p_{n,j}} - \left(\frac{G_f^*}{C_f^*} \right)^2 \frac{\partial c_{n,f}^*}{\partial p_{n,j}} \right], \tag{57}$$

$$\frac{\partial \text{Var}_f}{\partial a_n} = \mathbb{E}_f[\pi_f(p_n)]^2 - 2\bar{\pi}_f \mathbb{E}_f[\pi_f(p_n)], \tag{58}$$

$$\frac{\partial \text{Var}_f}{\partial p_{n,j}} = 2a_n (\mathbb{E}_f[\pi_f(p_n)] - \bar{\pi}_f) \frac{\partial \mathbb{E}_f[\pi_f(p_n)]}{\partial p_j}. \tag{59}$$

B.4.4 Information sharing and active constraints

I consider several types of signals. First, naturally, y_f includes $\{s_j\}_{j \notin \mathcal{J}_f}$; firm f 's rivals can always use their own shares to detect deviation from firm f . Second, if total quantities are shared, then y_f also includes $\sum_{j \in \mathcal{J}_f} s_j$; firm f 's rivals collectively observe total quantity and their own shares, letting them back out firm f 's total quantity but not how that quantity is split across firm f 's products. Third, if firm f shares its quantities with a subset of its rivals, then y_f also includes $\{s_j\}_{j \in \mathcal{J}_f}$; firm f 's SW23 constraint then includes the deviation on its own share. Fourth, if firm f shares its prices (even as an aggregate, such as an average price), then $\chi^2 = +\infty$ and the SW23 constraint is always vacuously satisfied, meaning firm f is only constrained by the Nash IC constraints. Fifth, I consider as a benchmark the case in which firms share future prices under the assumption that rivals can immediately punish deviation. In this scenario, the Nash IC also always holds, meaning the firms' prices are unconstrained. Thus, if all firms share future prices, the solution is a point mass on p^M .

C Robustness checks

This section discusses results under alternative assumptions. For the robustness checks described in Appendices C.1-C.3, I reestimate demand (if applicable), the timeseries process (if applicable), beliefs, and marginal costs under the Bertrand-Markov assumption. I then use the same procedure as in the main text to estimate the discontinuity in $\mathbb{E}[\omega_{jrt}]$ for Agri Stats clients and other producers at May 2019, and use this to compute the difference-in-discontinuities. Finally, I reestimate the conduct parameter that makes the estimated difference-in-discontinuities equal to zero. Table C.1 summarizes the results. While I discuss results specific to each sensitivity below, I ultimately conclude that the conduct test is robust to alternative assumptions about the information structure and equilibrium. In all sensitivities, the estimated discontinuity for Agri Stats clients is statistically distinguishable from zero at the 5% level, and not for other firms. The difference-in-discontinuities and the conduct parameter estimate are also statistically significant for all sensitivities except optimal demand instruments. Finally, Section C.4 discusses reduced-form evidence from my data that variation in vertical structure is inconsistent with it playing a large role in mediating the Agri Stats effect.

Sensitivity	Discontinuity in ω		Difference	Conduct parameter
	AS clients	Others		
Baseline	0.220 [0.082, 0.335]	0.003 [-0.104, 0.051]	0.217 [0.075, 0.378]	0.474 [0.210, 0.726]
<i>A. Information without sharing</i>				
Average price and total quantity	0.232 [0.118, 0.400]	-0.014 [-0.133, 0.050]	0.246 [0.134, 0.478]	0.489 [0.285, 0.741]
All prices and quantities	0.254 [0.109, 0.385]	-0.024 [-0.140, 0.032]	0.279 [0.132, 0.477]	0.547 [0.333, 0.798]
<i>B. Inventory signal with sharing</i>				
$\lambda = 0.1$	0.188 [0.059, 0.312]	0.002 [-0.100, 0.056]	0.186 [0.059, 0.350]	0.459 [0.180, 0.726]
$\lambda = 0.25$	0.173 [0.046, 0.295]	0.005 [-0.095, 0.061]	0.168 [0.043, 0.331]	0.454 [0.146, 0.747]
$\lambda = 0.5$	0.148 [0.026, 0.269]	0.011 [-0.091, 0.066]	0.136 [0.018, 0.304]	0.414 [0.072, 0.770]
<i>C. Transition between equilibria</i>				
3 months	0.222 [0.090, 0.331]	-0.005 [-0.097, 0.044]	0.226 [0.100, 0.361]	0.489 [0.257, 0.712]
6 months	0.177 [0.094, 0.312]	-0.013 [-0.097, 0.041]	0.190 [0.093, 0.331]	0.429 [0.231, 0.675]
12 months	0.139 [0.080, 0.292]	-0.020 [-0.094, 0.037]	0.160 [0.085, 0.328]	0.375 [0.208, 0.664]
<i>D. Optimal demand instruments</i>	0.065 [0.033, 0.143]	0.019 [-0.047, 0.058]	0.046 [-0.010, 0.159]	0.251 [0.000, 0.844]

Table C.1: Robustness of conduct results to alternative assumptions

Note: Brackets contain the 95% confidence intervals, computed with 500 bootstrap draws for the baseline and 200 bootstrap draws for sensitivities. The conduct parameter $\tau^* \in [0, 1]$ is estimated to minimize the square of the left-hand side of (16), assuming Agri Stats subscribers maximized the objective in (17). Discontinuities are estimated using a local linear regression before and after the bandwidth, with an Epanechnikov kernel and a bandwidth computed to minimize placebo discontinuities over June 2017-February 2019 and December 2019-May 2022 separately for each sensitivity.

C.1 Information structure

C.1.1 Information without sharing

The baseline assumption is that while not sharing information, firms only observe their own prices and shares. This assumption is defensible: according to litigation, firms paid millions of dollars for access to the Agri Stats reports and testified that their detail and timeliness are unmatched, while retailers attempted to purchase the reports and were denied (hence there must be additional information in the reports not available to retailers). However, it is possible that without Agri Stats, firms would be able to ascertain aggregates of past prices and shares from government indices, or more detailed data by purchasing retail scanner data.

I continue to assume that Agri Stats shared the full vector of prices and shares among its subscribers, and that firms observe their own prices and shares. In the first sensitivity, I

also assume that all firms have access to average prices and total inside shares; that is, they append to their own prices and shares the signal

$$y_{f_{rt}}^{\text{gov}}(p_{rt}, s_{rt}) = \left(\frac{1}{J_{rt}} \sum_{j \in \mathcal{J}_{rt}} p_{jrt}, \sum_{j \in \mathcal{J}_{rt}} s_{jrt} \right). \quad (60)$$

In the second sensitivity, I assume that all firms have access to the full vector of past prices and quantities; that is, $y_{f_{rt}}(p_{rt}, s_{rt}) = (p_{rt}, s_{rt})$ for all f_{rt} . In this case, the only informational benefit of Agri Stats is informing firms about rivals' costs and inventories.

Panel A of Table C.1 shows the results with these assumptions. These alternative assumptions strengthen the results, increasing the estimated difference-in-discontinuities and conduct parameter.

C.1.2 Inventory signal with sharing

The baseline analysis assumes that receiving inventory information ι_{jrt} enables perfect prediction of $p_{j_{r(t+1)}}$. This means that while sharing information, Agri Stats clients have no uncertainty about their fellow subscribers' prices when setting their own prices (which is the case in typical complete-information models). In this subsection, I weaken this assumption by assuming firms treat inventory information as an unbiased but noisy signal of future prices. In particular, firms conjecture that

$$\iota_{jrt} \mid p_{jrt} \sim \mathcal{N}(p_{jrt}, \lambda \Sigma_{jj}^p) \quad (61)$$

where Σ_{jj}^p is the jj entry of the estimated covariance matrix of price innovations. Firms update their priors accordingly. Then, $p_{jrt} = \iota_{jrt}$ is realized. The variance of the signal is proportional to the diagonal entries of Σ^p , so different values of λ can be interpreted as the additional precision provided by the Agri Stats inventory signal relative to the empirical variance of prices over time.

Suppose Agri Stats client f , while sharing information, enters period t with the prior $(\xi_{rt}, p_{(-f)rt}) \sim \mathcal{N}(\mu_{f_{rt}}, \Sigma_{f_{rt}})$. Let $\mu_{f_{rt}}^{AS \setminus f}$ and $\Sigma_{f_{rt}}^{AS \setminus f}$ be the subvector and submatrix associated with other Agri Stats subscribers' prices, let $\Sigma_{f_{rt}}^{o, AS \setminus f}$ be the covariance block between those prices and the rest of the state (that is, ξ_{rt} and non-subscribers' prices), and let $\Sigma_{AS \setminus f}^p$ be the diagonal of the submatrix of Σ^p corresponding to other Agri Stats clients' products. Define the gain $G_{f_{rt}} = \Sigma_{f_{rt}}^{AS \setminus f} \left(\Sigma_{f_{rt}}^{AS \setminus f} + \lambda \Sigma_{AS \setminus f}^p \right)^{-1}$. Then, before setting $p_{f_{rt}}$, firm f updates its

prior over $p_{AS\setminus f,rt}$ to $\mathcal{N}(\mu_{f,rt}^\lambda, \Sigma_{f,rt}^\lambda)$, where

$$\begin{aligned}\mu_{f,rt}^\lambda &= \mu_{f,rt}^{AS\setminus f} + G_{f,rt} \left(p_{AS\setminus f,rt} - \mu_{f,rt}^{AS\setminus f} \right), \\ \Sigma_{f,rt}^\lambda &= \Sigma_{f,rt}^{AS\setminus f} - G_{f,rt} \Sigma_{f,rt}^{AS\setminus f},\end{aligned}$$

and adjusts the covariance between these prices and the rest of the state to $\Sigma_{f,rt}^{o,AS\setminus f} (I - G_{f,rt}')$, holding the marginal belief over ξ_{rt} and non-subscribers' prices fixed.

The baseline model corresponds to $\lambda = 0$, where $G_{f,rt} = I$: the firm learns $p_{AS\setminus f,rt}$ exactly, and both $\Sigma_{f,rt}^\lambda$ and the covariances of these prices with the rest of the state collapse to zero. As λ increases, the inventory signal becomes weaker. As in the baseline, inventory information does not directly inform the firm about non-sharing firms' prices or ξ ; the firm does not revise the marginal beliefs over these objects, and the cross-covariances shrink only as much as is required to keep the firm's beliefs consistent with its sharpened belief over subscriber prices.

Panel B in Table C.1 displays the results with $\lambda = 0.1, 0.25$, and 0.5 . While observing less precise information slightly lowers the conduct parameter estimate, Bertrand-Markov conduct and $\tau^* = 0$ continue to be rejected.

C.2 Transition between equilibria

The moment condition is a difference-in-discontinuities at May 2019, meaning local conditions around May 2019 are important. While firms' beliefs vary continuously around this date, in the sense that June 2019 beliefs are derived from May 2019 beliefs despite them coming from separate equilibria, the binary assumption of two separate equilibria with different price Markov processes before and after May 2019 may be driving some of the difference. Although the distributions of the slope and innovation matrices appear similar before and after (Table 6), there may also be differences in the intercepts, or in particularly high-leverage slope coefficients.

To assess the sensitivity of the conduct tests to the transition between pricing processes, I smooth the transition between the estimated price timeseries. Let t^* correspond to May 2019, and $T > 0$ be an integer-valued transition length. For $t \in [t^* + 1, t^* + T]$, I assume that firms' conjectured Markov pricing rules are based on the linear mixture of the estimated pre-

and post-May 2019 timeseries:

$$p_{rt} = \frac{T + t^* - t}{T}(\bar{p}_r^{\text{pre}} + g^{\text{pre}}(w_{rt}, z_{rt}^{GH}) + B^{p,\text{pre}}(\xi_{r(t-1)}, p_{r(t-1)})) + \frac{t - t^*}{T}(\bar{p}_r^{\text{post}} + g^{\text{post}}(w_{rt}, z_{rt}^{GH}) + B^{p,\text{post}}(\xi_{r(t-1)}, p_{r(t-1)})) + e_{rt}^p,$$

$$e_{rt}^p \sim \mathcal{N}\left(\mathbf{0}, \frac{T + t^* - t}{T}\Sigma^{p,\text{pre}} + \frac{t - t^*}{T}\Sigma^{p,\text{post}}\right).$$

For product-regions that do not have pre-period observations, and hence do not have estimated fixed effects $\bar{p}_r^{\text{pre}}$, I just use $\bar{p}_r^{\text{post}}$ without mixing. The baseline analysis corresponds to $T = 1$; I reestimate beliefs, costs, and the conduct test with $T = 3, 6$, and 12 months.

Results are shown in Panel C of Table C.1. Bertrand-Markov conduct is rejected for all values of T . Notably, while the conduct parameter estimate does shrink for $T = 12$, the lower end of the 95% confidence interval is nearly identical to the baseline. This suggests that the mixing also has some precision benefits for the estimate as well. Ultimately, the assumption of an immediate transition to a new Markov equilibrium in pricing rules after May 2019 does not appear to be critical to the results.

C.3 Optimal demand instruments

Approximating the Chamberlain (1987) optimal instruments for use in BLP-style demand estimation was originally suggested by Berry et al. (1995), and recent papers have shown its ability to improve the performance of mixed logit demand models (Reynaert and Verboven, 2014; Conlon and Gortmaker, 2020). To test whether using optimal instruments affects my results, I follow the “approximate” optimal IV construction of Conlon and Gortmaker (2020), as initially conceived in Berry et al. (1999). I compute a pilot estimate of the structural parameters $\hat{\theta}^{1s}$ and residuals $\hat{\xi}^{1s}$ using one-step GMM and the baseline instruments. I then approximate $\mathbb{E}\left[\frac{\partial s(\theta; \xi)}{\partial \theta} \middle| z\right]$ with $\mathbb{E}\left[\frac{\partial s(\hat{\theta}^{1s}, \mathbb{E}[\hat{\xi}^{1s}])}{\partial \theta} \middle| z\right]$ and reestimate demand using two-step GMM with these as instruments.

Table C.2 shows the estimated demand parameters using the feasible approximation to the optimal instruments. The primary difference between these results and the baseline results is that estimates are made more price-sensitive, particularly in fresh turkey. The aggregate elasticity also moves further away from USDA’s estimates (Maples et al., 2022). Many of the other coefficients, particularly among the demographic interactions, remain similar.

I run these estimates through the remainder of the estimation pipeline and report the conduct results in Panel D of Table C.1. Notably, this is the only sensitivity where Bertrand-Markov conduct is not rejected at the 5% level, although it is at the 10% level ($p = 0.09$). If

Parameter	Fresh turkey				Deli turkey		
	Constant	Price	Breast	WS	Constant	Price	Low fat
β		-2.233 (0.149)	2.306 (0.504)	-3.131 (1.600)		-0.560 (0.050)	-0.532 (0.033)
diag(Σ)	0.152 (0.783)	0.165 (0.033)	1.357 (0.254)	0.350 (4.510)	0.000 (0.431)	0.131 (0.020)	0.000 (0.136)
II: Medium income	-1.653 (0.278)	0.665 (0.084)	-2.323 (0.194)	1.765 (0.243)	1.075 (0.060)	-0.074 (0.008)	-0.354 (0.041)
II: High income	-3.419 (0.261)	1.056 (0.081)	-2.398 (0.199)	2.232 (0.251)	0.052 (0.073)	0.064 (0.010)	-0.415 (0.039)
II: Children	0.996 (0.129)	-0.043 (0.033)	-0.725 (0.077)	0.003 (0.134)	1.134 (0.050)	-0.109 (0.006)	-0.044 (0.030)
Number of observations		57,616				109,820	
Median own-price elasticity		-6.41				-3.17	
Median aggregate elasticity		-1.06				-0.62	
First stage R^2		0.71				0.83	
First stage F -statistic		2,742.7				406.1	

Table C.2: Demand estimation results with approximate optimal instruments

Note: Heteroskedasticity-robust standard errors in parentheses. “WS” refers to whole/smoked.

we impose the alternative moment restriction of no discontinuity in Agri Stats clients’ costs, without differencing others’ costs, then Bertrand-Markov conduct is still rejected at the 5% level. If nothing else, these results demonstrate that the conduct test with the baseline supply-side model is falsifiable; the rejection of Bertrand-Markov conduct in the baseline analysis is not imposed by the supply-side assumptions. However, it does show that the results are at least somewhat sensitive to the demand estimates.

C.4 Vertical structure

This subsection discusses reduced-form evidence suggesting that processor-retailer bargaining was not a significant factor in the Agri Stats effect.

First, one could imagine that processors used the Agri Stats reports to learn which of their competitors were contracting with different retailers, and that therefore processors used the reports to segment the market. If this dynamic occurred, then the rate of churn of brand offerings at the retailer level might have been impacted by the end of the Agri Stats exchange. Figure C.1 plots churn — as measured by the share of production from producers that are new to a retail chain — by month over time. The first thing to notice is that churn is extremely small: less than 1% of production is from new producers to the chain. Second, this churn does not appreciably rise after May 2019, suggesting that there was no effect on




Figure C.1: Churn rates over time

Note: The vertical dashed line represents May 2019.




Figure C.2: Short-run Agri Stats effect by local retailer concentration

Each bubble is a market (region-module), and its size is proportional to total quantity in that market. The black line and shaded region are the predictions and 95% confidence intervals from a weighted regression of local short-run Agri Stats effect on the retail-level HHI, with standard errors computed using a Murphy-Topel correction.

churn of the information exchange ending. This suggests that we can limit the effects of the exchange to the intensive margin of prices, as opposed to the extensive margin of product assortment.

Second, I look for whether local retailer concentration explains heterogeneity in the short-run Agri Stats effect on prices. Since downstream bargaining power can countervail upstream collusion ([Chassang and Ortner, 2023](#)), one would expect that retailers with more market power would be less affected by any pre-May 2019 collusion, and therefore should have Agri Stats effects with a smaller magnitude. Figure [C.2](#) reveals that this is not the case (using local concentration as a proxy for market power): if anything, areas with more concentrated retail sectors had slightly larger-magnitude short-run Agri Stats effects, with a 100-point retail HHI increase being associated with a 0.04pp larger Agri Stats effect (although this is

not statistically significant at the 5% level). This suggests that heterogeneity in bargaining power is not a significant factor in explaining price effects of information sharing in this context.

References

- BERRY, STEVEN, JAMES LEVINSOHN, AND ARIEL PAKES. 1999. “Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy.” *American Economic Review* 89 (3): 400–431.
- BERRY, STEVEN T., JAMES A. LEVINSOHN, AND ARIEL PAKES. 1995. “Automobile Prices in Market Equilibrium.” *Econometrica* 63 (4): 841–890.
- BLUNDELL, RICHARD, AND STEPHEN BOND. 1998. “Initial Conditions and Moment Restrictions in Dynamic Panel Data Models.” *Journal of Econometrics* 87 (1): 115–143.
- CHAMBERLAIN, GARY. 1987. “Asymptotic Efficiency in Estimation with Conditional Moment Restrictions.” *Journal of Econometrics* 34 (3): 305–334.
- CHASSANG, SYLVAIN, AND JUAN ORTNER. 2023. “Regulating Collusion.” *Annual Review of Economics* 15:177–204.
- CONLON, CHRISTOPHER, AND JEFF GORTMAKER. 2020. “Best Practices for Differentiated Products Demand Estimation with PyBLP.” *RAND Journal of Economics* 51 (4): 1108–1161.
- KAPETANIOS, GEORGE. 2008. “A Bootstrap Procedure for Panel Data Sets with Many Cross-Sectional Units.” *Econometrics Journal* 11 (2): 377–395.
- KNOLL, DANA A, AND DAVID E KEYES. 2004. “Jacobian-Free Newton–Krylov Methods: A Survey of Approaches and Applications.” *Journal of Computational Physics* 193 (2): 357–397.
- MAPLES, WILLIAM E., B. WADE BRORSEN, WILLIAM F. HAHN, MATTHEW MACLACHLAN, AND LEKHNATH CHALISE. 2022. *Structure of the USDA Livestock and Poultry Baseline Model*. Technical report. USDA Economic Research Service Technical Bulletin No. 1956.
- MORROW, W ROSS, AND STEVEN J SKERLOS. 2011. “Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand.” *Operations Research* 59 (2): 328–345.
- REYNAERT, MATHIAS, AND FRANK VERBOVEN. 2014. “Improving the Performance of Random Coefficients Demand Models: The Role of Optimal Instruments.” *Journal of Econometrics* 179 (1): 83–98.
- ROBINSON, PETER M. 1988. “Root-N-Consistent Semiparametric Regression.” *Econometrica*, 931–954.
- ROODMAN, DAVID. 2009. “How to Do xtabond2: An Introduction to Difference and System GMM in Stata.” *Stata Journal* 9 (1): 86–136.
- SAAD, YUCEF, AND MARTIN H SCHULTZ. 1986. “GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems.” *SIAM Journal on Scientific and Statistical Computing* 7 (3): 856–869.
- SIGMUND, MICHAEL, AND ROBERT FERSTL. 2021. “Panel Vector Autoregression in R with the Package panelvar.” *Quarterly Review of Economics and Finance* 80:693–720.