

Reasoning About ‘When’ Instead of ‘What’ : Collusive Equilibria with Stochastic Timing in Repeated Oligopoly

Justin Grana ^{*†‡} James Bono ^{§¶} David Wolpert ^{*†}

Friday 4th August, 2017

Abstract

We analyze a continuous time game of Bertrand competition with private monitoring that includes asynchronous signals and asynchronous actions. Unlike existing models where firms observe a signal that is imperfectly correlated with demand parameters and firms’ prices, we assume firms observe demand and prices perfectly but at times governed by independent Poisson processes. This implies that any firm’s price change will go undetected for some time or may never be detected. Consequently, the model focuses on the strategic considerations that arise when firms reason about the *timing* of the signals rather than the content. Like traditional models, we find that if firms are patient enough in terms of time discounting, there exists a sequential equilibrium in which they collude. Unlike traditional discrete-time oligopoly models, we also show that for any discount rate, if firms receive signals about one another’s actions *fast enough*, there exists a sequential equilibrium in which they collude. This new result explicitly illustrates how parameters that govern the timing of

*Santa Fe Institute, 1399 Hyde Park Rd, Santa Fe, NM 87501

†Supported by Army Research Office Grant W911NF-12-R-0012-02

‡Corresponding Author

§Economists Inc., 101 Mission Street Suite 1000, San Francisco, CA 94105

¶Supported under NASA Grant NNH13ZEA001N-SSAT:B.2-AFCS-1.6

events are as important as the parameters that govern preferences when determining the existence of collusive equilibria. We then characterize the collusive equilibria in terms of the parameters that govern the timing of signals, actions and demand shocks.

1 Introduction

In many oligopoly settings, firms receive demand information asynchronously and at random times. Presumably, a firm that receives demand information faster would be able to act on such information to increase profits relative to its competitors. However, this feature is not clear in a collusive environment. For example, in 2015, the European Commission fined Express Interfracht—a cartel composed of Austrian railway company Österreichische Bundesbahnen and German railway company Schenker—€49M for “operating a cartel breach of EU anti-trust rules.” The unofficial press release noted that the companies limited competition by “exchanging confidential information on specific customer requests” and “sharing transport volumes contracted by downstream customers [European-Commission, 2015].” In another case among 11 producers of high-voltage underground cables, the producers agreed to geographically divide the market. In doing so, “whenever Japanese and Korean companies received requests from European customers, they would notify their European counterparts [European-Commission, 2014].” In both cases, the cartelists asynchronously received demand information but choose to share the information with their competitors. However, it is not immediately clear as to why the firms would not use the privately acquired information for their own profit, especially since any firm in the cartel is not aware of if and when other firms receive demand information.

Similarly, firms may not always know their competitors’ price but may learn such prices asynchronously and at random times. For example, there is a lack of transparency of price information in the health care industry [Crowe, 2017, Wong et al., 2016, Rosenthal et al., 2013]. However, independent organizations such as Clear Health Costs use consumer supplied data as well as independent market research to provide price information for specific procedures at specific hospitals. Due to variations in consumer reporting and market research efforts, such organizations do not always receive the most recent price information and thus can only update public price

information at asynchronous and random times. More generally, any market where price transparency relies on self-reporting consumers, firms will only learn about the price of their competitors asynchronously and at random times.

Many other real-world markets share the characteristic that the precise times when a player acquires information are random. For example, whenever a firm uses real-time market research, “now-casting”, social media mining, or crowd sourcing, information does not arrive at fixed intervals, but rather at times that are determined via stochastic processes. Those same firms who receive information at random times will often make their decisions at times that are determined stochastically, either in response to (randomly arriving) new information or due to stochastic internal deliberations. We refer to the types of strategic settings in which decision makers transmit, receive and react to information stochastically and asynchronously as **event-driven**. More generally, many economic scenarios besides oligopoly are event-driven. Prominent examples include investment decisions, innovation strategy and technological diffusion. Despite the ubiquity of real-world, event-driven scenarios, there has been minimal research dedicated to explicitly investigating the strategic consequences of an event-driven environment (a notable exception is [Keller et al., 2005]). In this paper, we highlight the strategic consequences of an event-driven environment in the context of Bertrand competition. In doing so, we intend to specifically determine how uncertainty in timing can impact the incentives for firms to collude. However, our broader goal is to highlight how including uncertainty in timing can lead to novel strategic behavior that remains hidden in standard discrete time games.

More specifically, we extend a traditional Bertrand model to a fully event-driven environment. We show how how uncertainty about the *timing* of events can play the same role that uncertainty about information content does in conventional Bertrand models. In particular, previous models have shown that for many demand, cost, and information set specifications, there exist collusive equilibria where firms’ per-period profits are greater than profits in the static Bertrand equilibrium. More generally, folk theorems relate the existence of collusive equilibria to the time-discounting of the decision-makers. Analogously, we establish and characterize the existence of collusive equilibria in an event-driven environment and highlight how the rates at which firms receive various types of information either facilitate or preclude collusion.

In our model firms receive demand as well as price information at (different) random times governed by independent Poisson processes. This is in contrast to typical

oligopoly models that assume firms receive signals synchronously at discrete, pre-fixed times. Upon receiving information, a firm can engage in (optional) cheap talk, and then set its price. A main feature of our model that distinguishes it from the literature is that each firm knows that its competitor will not detect any price change for a strictly positive (and possibly infinite) amount of time. As in traditional oligopoly models, we show that for any fixed values of the parameters, if firms are patient enough (in terms of the discount rate), then there exists a collusive equilibrium in which firms earn monopoly profits. However, we extend this intuition and show that if firms receive signals at a high enough rate, there exists a collusive equilibrium. So the parameters that govern the timing of an event-driven scenario are as important as the discount rate in determining the existence of collusive equilibria. Furthermore, increased transparency—in terms of the *frequency* of signals—can bolster the incentive to collude since it reduces the expected time it takes for a competitor to detect a deviation. From a policy-maker’s perspective, this implies that reducing the rate of price signals can benefit consumers.

We proceed as follows. In section 2 we compare our model with other oligopoly models in the literature. We also review folk theorems and highlight that none of the established folk theorems guarantee the existence of collusive equilibria in our model. In section 3 we provide a narrative description of our model followed by a formal definition. We also present our main results, establishing the existence of collusive equilibria in terms of the discount rate and the speed at which firms get signals. In section 4 we further characterize the collusive equilibria in terms of all of the model’s parameters. In section 5 we extend the basic model of section 3 and show that the results hold when firms also receive private profit signals. Finally in section 6 we suggest directions for future work in the study of event-driven games.

2 Repeated Games, Collusion and Folk Theorems

The literature on repeated games, cooperation and collusion is vast and expansive. Here, we outline the standard modeling assumptions of repeated games and review the literature most relevant to our work: repeated games and collusion models with private information and asynchronous signals and actions.

Existing dynamic models of collusion are best classified according to their mon-

itoring scheme.¹ The most basic monitoring scheme is perfect monitoring without private information. Under this scheme all firms' actions as well as game parameters are publicly and perfectly observed. For example, there exist collusive equilibria in a perfect information game when firms must consume a renewable resource to facilitate production [Colombo and Labrecciosa, 2015]

A slight relaxation of perfect monitoring with perfect information is the case of perfect monitoring with private information. One example is if firms' privately known per-unit marginal costs evolve according to a finite, discrete time Markov process, but all pricing decisions are perfectly observable by all firms. In this case, there exist at least two types of collusive equilibria, one in which firms do not reveal their private costs and one in which firms engage in cheap-talk and truthfully reveal their private costs [Athey and Bagwell, 2008]. Another example of perfect monitoring with private information is if firms receive a *private* signal that is correlated with demand conditions and then engage in simultaneous cheap talk and subsequently choose prices. Then, at the end of each period, firms' prices are publicly revealed. In this case the simultaneous and mandatory cheap talk can enforce a collusive equilibrium [Gerlach, 2009].

Oligopolies with perfect monitoring (with perfect or imperfect information) have the property that all deviations from a collusive agreement would be observable and therefore punishable. This property facilitates collusion in two ways. First, firms do not have to form beliefs about whether any of the other firms secretly undercut the market. So there is never uncertainty about—for example—if a period of low profits is due to low demand or a deviation. Instead, the firms know with certainty if they had been undercut. Second, perfect monitoring increases firms' ability to coordinate punishment. Since the entire history of past actions is common knowledge to all firms, if one firm were to deviate from the collusive agreement, *all* other firms would be aware of the deviation and could jointly punish. In contrast, without perfect monitoring, a punishing firm would not know if all other firms would enter the punishment phase since their signals may have differed. So it might not be optimal for a firm to enter the punishment phase after detecting a deviation if all other firms do not also enter the punishment phase.

Under imperfect public monitoring firms' past actions are imperfectly observable,

¹For a thorough review and rigorous definition of monitoring schemes, see [Mailath and Samuelson, 2006].

but all firms observe a common signal. The canonical example of public monitoring is a model in which firms in repeated Cournot competition are all subject to *the same* price shock but do not observe other firms' quantity choices. With a sufficiently high discount rate, there exists a collusive equilibrium in which firms' (per period) profits are higher than their equilibrium profits in the static game [Green and Porter, 1984]. Recently, imperfect public monitoring has schemes have been studied under heterogeneous discount rates [Sugaya, 2015].

The final monitoring scheme is known as private monitoring. Under private monitoring, firms do not observe the actions of other firms but receive a signal that is correlated with the actions of the other firms. The main distinction from public monitoring is that the signal is private information to the firm that receives the signal, *and* the firm is allowed to base its decision on the entire history of its private signals.² This severely complicates the mechanisms that facilitate collusion, because firms must form higher order beliefs. If a firm believes it has been undercut, it is only optimal for the firm to enter a punishment regime if it thinks that the deviating firm believes its deviation was detected, and so on.

Despite these complications, collusion is still possible under private monitoring when additional mechanisms are allowed. For example, with the existence of a mediator and cheap talk, for any arbitrary *fixed* period length, there exists a discount rate that supports a collusive equilibrium in which firms alternate being the monopolist [Rahman, 2014]. Similarly, when the private signals are correlated and firms engage in simultaneous cheap talk, there exists a discount rate that supports a collusive equilibrium [Awaya and Krishna, 2015]. Finally, when firms are allowed to make publicly observable side payments, a collusive equilibrium exists even with very limited information [Chan and Zhang, 2015].

While our work shares some characteristics with the models described above (for example, we use private demand information but public price signals like [Gerlach, 2009]), our event-driven model is distinct from the literature in several ways. The main distinguishing factor is that our model has asynchronous actions and asynchronous signals that occur at stochastic times and not at pre-fixed intervals. Also, our event-driven model is one of private monitoring since firms' pricing decisions are not observed by the other firm at the instant the decision is made. Additionally,

²As opposed to games solved using the perfect public equilibrium solution concept where firms may receive private information but do not base their decisions on the history of private information.

our results do not rely on the existence of any other mechanisms, such as a moderator, side payments or a public correlation device (see [Athey and Bagwell, 2008] and [Chan and Zhang, 2015] for such examples). In short, our model is the first to consider private information, private monitoring and asynchronicity of signals *and* actions.

Outside of the collusion literature, the work of [Keller et al., 2005] analyzes a multi-armed bandit game in which players receive payoffs at delayed times. However, that model does not analyze the impacts of cooperation due to the delay in timing. Additionally, the only delay is a delay in payoffs where our model allows for stochastic timing in signals and actions.

The analysis of [Abreu et al., 1991] concentrates on the impacts of timing in a prisoner’s dilemma game. However, the model we present here has several features that are not present in the model of [Abreu et al., 1991]. In [Abreu et al., 1991], players take actions simultaneously at prefixed and deterministic times and it is only the arrival of information that occurs at stochastic lags. Furthermore, in that model all players receive information simultaneously. In our model, we allow players to learn demand and take actions at asynchronous, stochastic and privately observed times. This is an important distinction because it means that players must reason about when and *if* other players receive information. Furthermore, in [Abreu et al., 1991] the main result of the “good news case” and “bad news case” rely on the assumption that the parameter governing the timing of signals depends on the all players’ previous actions. As a result, players detect deviations from cooperation by reasoning about the number of signals they receive. In our model, the frequency of signals is exogenous and thus players cannot use the frequency of signals to reason about other player’s past actions. Finally, our model shows that the relative rates at which nature changes, players observe nature, and players observe each other’s actions has important strategic consequences. Considering all three rates simultaneously is not provided in the model of [Abreu et al., 1991] but is common in several real world strategic environments in which different types of information arrive at various rates.

Other investigations into asynchronous repeated games are concerned with the limit of the frequency of actions, as the time between actions becomes small. For example, [Fudenberg and Levine, 2007] examines the impact of different asynchronous monitoring schemes under an assumption of frequent action. Similarly, [Sannikov and Skrzypacz, 2010] examines equilibria in games of frequent actions when players

obtain information continually through a stochastic process with a Brownian motion component and a Poisson component. However, in both of these models, actions occur synchronously at discrete time intervals, and the monitoring is restricted to public signals. Due to the synchronicity of actions and restriction to public monitoring scheme, neither [Sannikov and Skrzypacz, 2010] nor [Fudenberg and Levine, 2007] are directly related to our model in which players act asynchronously and receive private signals.

For some models of infinitely repeated oligopoly, folk theorems guarantee the existence of a collusive equilibrium. Such broad theorems establish that as long as players are patient enough, the only restriction on equilibrium payoffs in an infinitely repeated game is that they are feasible and above players' minimax payoffs in the static game. Folk theorems were originally applied to simple repeated games of perfect monitoring but recently “a large literature has extended the folk theorems under successively weaker assumptions about the monitoring structures that govern the signals players receive about one another's actions” [Fudenberg et al., 2014].

However, none of the folk theorems guarantee the existence of collusive equilibria in our event-driven oligopoly model. Even though our event-driven model includes cheap talk, the folk theorems that rely on cheap talk [Ben-Porath and Kahneman, 1996, Kandori and Matsushima, 1998, Obara, 2009, Compte, 1998] cannot be applied since those theorems do not consider asynchronicity. In fact, so-called “anti-folk theorems” show that the logic and main results of standard folk theorems do *not* apply when actions are asynchronous [see [Lagunoff and Matsui, 1997] and [Takahashi and Wen, 2003]]. The folk theorem with modeling assumptions closest to ours only considers asynchronous signals [Fudenberg et al., 2014]. In that treatment, players observe signals with either a stochastic lag or not at all. Firms can observe signals asynchronously, and the times at which signals are received are private information. Under these conditions, there is a folk theorem that proves the existence of multiple equilibria [Fudenberg et al., 2014]. However, this scenario differs from our model in that it assumes that players act simultaneously at fixed, discrete time intervals. Our model contains asynchronous actions, asynchronous signals and private information and thus far, there does not exist a folk theorem that characterizes the equilibrium set under such assumptions.

3 Model

In this section, we first present a broad overview of the event-driven Bertrand model followed by a formal specification. We next introduce the collusive information sharing scheme and give two propositions that characterize the existence of such an equilibrium. In order to present the fundamental components of our model, we make several simplifying assumptions in favor of expository clarity. However, we make special note as to where our assumptions can be relaxed at the cost of additional notation. We also point to other examples in the literature that adhere to similar assumptions as ours.

3.1 Model Overview

In our event-driven Bertrand duopoly model, demand is defined by a unit of consumers that all purchase a homogenous good as long as the price is below a common reserve price, as in [Athey and Bagwell, 2008].³ However, the common reserve price evolves stochastically over time. We assume that the reserve price jumps between two exogenously specified values where the jump times are governed by a Poisson process. Our model can easily be extended to N reserve prices at the expense of more notation. The jumps in the reserve price capture demand shocks and are a continuous time analog of a demand function described by a discrete time, discrete state Markov process, which are a common way to represent demand stochasticity over time (see, for example [Collard-Wexler, 2013].)

Firms do not learn about changes in demand immediately. Instead, there is an independent Poisson process for each firm that governs when that firm receives a private signal indicating the reserve price. That signal is noiseless. That is, although firms do not learn of a change in the reserve price at the instant it occurs, when they learn the reserve price, they learn the then-current price exactly. Since there are independent Poisson processes for each firm's demand signal, firms not only receive signals at different *times* but, in general, learn the reserve price at different *rates*.

Since the signal-generating Poisson process and the demand-changing Poisson process are independent, the reserve price can jump multiple times without a firm receiving a signal. Conversely, a firm can receive a signal about the reserve price multiple

³Our main results can be extended to the case of linear demand with a stochastic intercept or slope term, but we use the reserve price formulation for simplicity.

times during an interval without any change in the reserve price during that interval (so that it would be the same signal they receive each time). These delayed signals capture informational frictions due to uncertainty in timing.

Upon receiving a private demand signal, a firm makes two decisions. First, it decides whether to immediately engage in cheap talk with the other firm, and if so, what information to convey. A firm can tell the other firm the content of the signal, lie to the other firm or not tell the other firm anything. We assume that if a firm decides to engage in cheap talk, the message it sends is immediately received by the other firm. A key distinction in the event-driven model is that cheap talk is optional, and players do *not* engage in cheap talk simultaneously as they do in [Awaya and Krishna, 2015] and the folk theorems of [Fudenberg et al., 2014]. Second, the firm chooses a price. We assume that firms can only change their price upon the receipt of new information. This assumption is consistent with discrete time games of Bertrand competition considered in the literature which assume price changes, signals and cheap talk all happen in the same period as observations. While this assumption may seem limiting, it is mainly for analytical convenience. Later, we argue that allowing firms to change their prices at random times when they do not receive information does not change the main result.

If the firm that receives a signal from the market decides to send a message to the other firm, both firms will have the opportunity to simultaneously adjust their price (messages are sent and received instantaneously). If they do not, then only the firm that received the signal from the market can change its price, not the other firm.

Firms also receive a public signal, revealing both firms' current price. The public signal arrives at random times governed by another independent Poisson process.⁴ The assumption that firms learn their rival's price with a stochastic lag is a relaxation of a standard assumption in models of perfect monitoring that firms immediately learn about their rival's price whenever it changes. Introducing this stochastic lag means that firms have uncertainty about when they might learn the other firm's price. This uncertainty is distinct from the standard source of uncertainty – an imperfect (but immediate) signal about the other firm's price. Just as firms can adjust their own price upon receiving a private demand signal or message, they can adjust their own

⁴In some markets, these signal generating processes may be correlated. Furthermore, firms may endogenously invest in increasing the rate at which they receive signals. While we acknowledge these complexities, we leave such an analysis for future work.

price upon receiving the public price signal.

The results and intuition do not require that both firms' prices be publicly revealed simultaneously. At the expense of a more complicated model, there could be two independent Poisson processes each governing the times at which one of the firms' prices are publicly revealed. The only requirement for the main results is that for both firms, at random times there is a public signal that reveals the firm's current price.

Finally, while firms learn about demand and prices, they do not explicitly observe their own profits. The broad motivation for such a modeling choice are markets in which demand information as well as pricing information arrives much faster than profit information. For example, many large firms only compute their profits as part of a quarterly report but must make many pricing decisions in the interim without knowing their most up-to-date profits. However in section 5, we analyze an augmented form of the initial model that includes private profit signals. We show that doing so does not preclude the existence of an information-sharing cartel.

3.2 Model and Main Result

In this section, we present the model in formal detail. We begin by defining the firms and the reserve price. Next, we describe the distribution of the signals and action times and then define histories, strategies and expected profits. We conclude this section with our main result that characterizes the collusive equilibrium in terms of time discounting and the frequency of signals.

3.2.1 The Firms and the Market

The model involves two firms indexed by $i = 1, 2$ that choose prices in \mathbb{R}_+ at discrete times in $[0, \infty]$. We write p_i^t for firm i 's price at time $t \in [0, \infty]$. The firms can change their price upon receiving a signal or message as described below.

The reserve price, ρ , at time t is given by $\rho_t \in \{\rho^L, \rho^H\}$ where $\rho^L < \rho^H$. We define $\mathcal{R} = \frac{\rho^H}{\rho^L}$ as the ratio of reserve prices. We assume that the times at which ρ_t switches between its values are governed by a Poisson process with rate γ .

3.2.2 Private Demand Signals

At times $\tau_i^1, \tau_i^2 \dots \in \mathbb{R}_+$ (with $\tau_i^j > \tau_i^k$ iff $j > k$), firm i receives a signal $y \in Y = \{y^L, y^H\}$ that indicates the current reserve price. Those times at which firm i receives a signal from the market are randomly distributed according to $\tau_i^k - \tau_i^{k-1} \sim \exp(\lambda_i)$. Formally, let $y_{\tau_i^k}$ be the signal that i receives at time τ_i^k . Then, conditional on firm i receiving its k 'th signal at time τ_i^k , $y_{\tau_i^k} = y^L$ iff $\rho_{\tau_i^k} = \rho^L$, and similarly $y_{\tau_i^k} = y^H$ iff $\rho_{\tau_i^k} = \rho^H$. This means that at the instant firm i receives a message indicating the reserve price, it knows the reserve price exactly.

3.2.3 Public Price Signals

Let $Z = \mathbb{R}_+^2$ with generic element z be the set of public signals both firms simultaneously receive that indicates each firm's current price in the market. Firms have perfect recall so each firm always knows its own price. Therefore, the only information z would give to firm i is firm j 's price. The public signal z is observed at times $\tilde{\tau}^1, \tilde{\tau}^2 \dots \in \mathbb{R}_+$ with $\tilde{\tau}^j > \tilde{\tau}^k$ iff $j > k$. Formally, let $z_{\tilde{\tau}^k}$ be the public signal at time $\tilde{\tau}^k$. Then $z_{\tilde{\tau}^k} = (p_1^{\tilde{\tau}^k}, p_2^{\tilde{\tau}^k})$. That is, the signal z tells both firms the most recent price chosen by each firm. The times of the public signal are such that $\tilde{\tau}^k - \tilde{\tau}^{k-1} \sim \exp(\mu)$. The signal z is not indexed by firms because the signal is public. This distinguishes the public signal z regarding price information from the private signal y regarding demand information.

3.2.4 Messages

Firms are allowed to engage in “cheap talk” upon receiving a private signal, y . That is, upon receiving a private signal firm i chooses a price in \mathbb{R}_+ and sends a message $x \in X = \{x^L, x^H, \emptyset\}$ to firm j . Intuitively, firm i can choose to tell firm j the reserve price it actually observed, the opposite reserve price, or nothing at all. If firm i chooses \emptyset , firm j does not receive a message and is therefore unaware that firm i received a signal indicating the reserve price.

3.2.5 Histories and Strategies

Firms have the opportunity to adjust their price whenever they receive a demand signal, price signal or a message. Let $h_i^k = (\omega_1, a_1), (\omega_2, a_2) \dots (\omega_k, a_k)$ be a history for firm i that entails k pairs of observations and actions. Each observation, ω_m , is

a tuple that is an element of $\mathbb{R}_+ \times (Y \cup Z \cup X)$ where the first element of ω_m is the time of the m 'th observation and the second element is the observation of firm i .⁵ Each action, a_m , is a tuple in $\mathbb{R}_+ \times Z$ that indicates the action that firm i took upon receiving signal ω_m . The first element of a_m is the price that firm i set upon receiving the m 'th signal, and the second element is the message firm i sent to firm j upon receiving the m 'th signal.⁶ We write the set of all of i 's possible histories of length k as H_i^k and define $H_i = \cup_{k=0}^{\infty} H_i^k$ as all of firm i 's possible histories with $H_i^0 = \emptyset$. We sometimes write h_i^t for firm i 's history at *time* t , with the context making the meaning clear. We let H be the set of full game histories.

Similarly, define an *actionable history* of firm i of length k as $\tilde{h}_i^k = ((\omega_1, a_1), (\omega_2, a_2) \dots (\omega_{k-1}, a_{k-1}), \omega_k)$ which indicates that firm i has previously observed signals $\omega_1, \omega_2 \dots \omega_{k-1}$ and taken associated actions $a_1, a_2 \dots a_{k-1}$, has just received a signal ω_k , and now must decide how to set its price and the message (if any) to send to the other firm. For each firm i , we write \tilde{H}_i^k for the set of all actionable histories of length k and \tilde{H}_i for the set of all actionable histories of any length. So each firm's strategy is a function $S_i : \tilde{H}_i \rightarrow (\mathbb{R}_+, Z)$. Again, we sometimes write \tilde{h}_i^t to be an actionable history at *time* t , with the context making the distinction clear.

3.2.6 Profits

Firms accrue profits continuously. Normalizing constant per-unit marginal cost to 0, at time t firm i accumulates reward at the rate

$$r_i^t(p_i^t, p_j^t, \rho_t) = \begin{cases} p_i & \text{if } p_i < p_j \text{ and } p_i \leq \rho_t \\ \frac{p_i}{2} & \text{if } p_i = p_j \leq \rho_t \\ 0 & \text{otherwise.} \end{cases}$$

For given values of $p_i^t, p_j^t, \rho_t \forall t$ firm i 's discounted profits are given by

$$\Pi_i = \int_0^{\infty} e^{-\delta t} r_i^t(p_i^t, p_j^t, \rho_t) dt \quad (1)$$

⁵An "observation" refers to either a signal from the market or a message from another firm

⁶If ω_v indicates a public price signal, then the second element of a_v is restricted to be \emptyset , which captures the fact that there is no cheap talk after public price signals.

where δ is a discount rate such that higher values of δ indicate *higher* preference for profits in the present relative to profits in the future, i.e., less patience.

Given a strategy profile, $s = (s_1, s_2)$, firm i 's expected profits are given by

$$U_i(s) = \mathbf{E}_{s,\rho}[\Pi_i]. \quad (2)$$

The subscript on the expectation operator indicates that the expectation is taken with respect to each firms' prices—which depend on s —and ρ which evolves independently of s .

3.3 Collusive Scheme

This section presents the collusive strategy profile and proves that it is supported in equilibrium. For reference, a description of key parameters is given in table 1.

Parameter	Definition
ρ^L	Low reserve price
ρ^H	High reserve price
\mathcal{R}	ρ^H / ρ^L
λ_1	Rate at which firm 1 receives private demand signals
λ_2	Rate at which firm 2 receives private demand signals
γ	Rate at which the reserve price changes
μ	Rate at which firms receive public price signals
δ	Common discount rate

Table 1: Parameter Dictionary

Definition 1. A Truthful Information Sharing Collusive (*TISC*) scheme is a strategy profile in which firm $i = 1, 2$ choose actions as follows:

1. When firm i receives a signal from the market indicating the reserve price, firm i informs firm j of the reserve price and firm i sets its price equal to the reserve price.
2. When firm i receives a message from firm j indicating the reserve price, firm i sets its price to the price indicated by firm j .

3. When firm i receives a public price signal that indicates both firms are charging the same price, firm i does nothing. If the public price signal indicates that the firms are charging different prices, firm i sets its price to 0 forever.

For the TISC scheme to be an equilibrium, firms' incentives must satisfy three conditions. The first is that upon learning the reserve price, firms do not have the incentive to undercut one another. This is the notion of collusion that is typical in the literature, and we call it the **collusion incentive**. Second, under the TISC scheme firms must have an incentive to truthfully convey the reserve price in their cheap talk. We call this the **truth-telling incentive**. Finally, firms must have the incentive to adjust their prices in response to changes in the reserve price. For example, if at a given instant both firms are charging ρ^L and firm i receives a signal from the market that the reserve price is ρ^H , it must benefit firm i to charge ρ^H given that it tells firm j that ρ^H is the reserve price. We call this the **price-adjusting incentive**. These three incentives are not necessarily distinct. For example, upon learning the reserve price is ρ^L , firm i can capture the entire market by changing its reserve price to ρ^L and not telling firm j anything. Firm i can achieve the same result by telling firm j that the reserve price is ρ^L and undercutting ρ^L slightly.⁷

Proposition 1. For any $\mu, \gamma, \lambda_1, \lambda_2 > 0$ and $1 + \frac{\gamma}{\gamma + \lambda_1 + \lambda_2} < \mathcal{R} < 2 + \frac{\lambda_1 + \lambda_2}{\gamma}$, there exists a $\underline{\delta} > 0$ such that for any value of $\delta < \underline{\delta}$, there exists a sequential equilibrium where equilibrium strategies, s^* , are specified according to the TISC scheme. We call this the **TISC equilibrium**.

Even though all proofs are given in the appendix, the proof of Proposition 1 is worth summarizing here as it discusses how the model reduces to a continuous time Markov game and how to apply to one-stage deviation principal to determine the existence of a collusive equilibrium. Intuitively, it is driven by two main properties of the TISC scheme. First, we show that under the TISC scheme, the firms' rate of rewards are governed by a *continuous time Markov chain*. This allows us to employ the value function approach for analyzing expected payoffs in a Markov Chain. Second, we illustrate that firms' continuation values in equilibrium only depend on whether they have observed a deviation in the past and the most recent signal they received.

⁷Since prices are real numbers, we adopt the convention, and terminology, that a firm i can undercut firm j "slightly". That is, suppose firm j 's price is ρ^L . Then firm i can "undercut slightly" and capture the entire market but accrue rewards at a rate of ρ^L .

This allows us to adopt a finite automaton representation of strategy profiles [Mailath and Samuelson, 2006]. This reduction to a finite automaton representation reveals that all possible one-shot deviations by a firm yield one of five possible continuation values. Consequently, we prove that for a δ small enough and appropriate values of \mathcal{R} , none of the deviations are profitable. Furthermore, since firms (endogenously) truthfully share demand information, it is trivial to define the beliefs that characterize the sequential equilibrium.

The condition $1 + \frac{\gamma}{\gamma + \lambda_1 + \lambda_2} < \mathcal{R} < 2 + \frac{\lambda_1 + \lambda_2}{\gamma}$ represents upper and lower bounds on the ratio of high to low reserve prices to ensure that the price adjusting incentive condition is met. If the ratio is less than the left hand side, then firms will not charge ρ^H upon learning the reserve price is ρ^H because of the risk of earning 0 profits when the reserve price jumps back down to ρ^L without their knowing. On the other hand, if the ratio is greater than the right hand side, firms will not charge ρ^L upon learning the reserve price is ρ^L because of the risk of missing out on greater profits when the reserve price jumps back up to ρ^H without their knowing. However, if \mathcal{R} is outside of the bounds, firms might still be willing to collude, they would just collude at a fixed price, never communicate and ignore demand signals. Since we are interested in the case where demand information is profitable to firms, we will thus only focus on the case where $\mathcal{R} = 2$. That is, we will only focus on the case where for all values of the parameters, firms have an incentive to adjust their prices in response to changes in demand.

The same reasoning that establishes proposition 1 establishes the following:

Proposition 2. *For any $\delta, \gamma, \lambda_1, \lambda_2, \rho^L, \rho^H > 0$ such that $1 + \frac{\gamma}{\gamma + \delta + \lambda_1 + \lambda_2} < \mathcal{R} < 2 + \frac{\delta + \lambda_1 + \lambda_2}{\gamma}$, there exists a $\bar{\mu}$ such that for any value of $\mu > \bar{\mu}$, there is a TISC equilibrium.*

Proposition 2 says that for any value of δ and appropriate ratio \mathcal{R} , as long as the public price signal (and therefore the possibility of punishment) happens "fast enough," then the TISC scheme is supported in equilibrium. While mathematically similar to proposition 1, proposition 2 presents a new condition to support a collusive equilibrium. Namely, firms do not need to be "patient enough" to guarantee the existence of the collusive equilibrium. Instead, it is only required that detection will likely occur fast enough after a deviation. In other words, increased monitoring quality through an increase in the *frequency* of monitoring can enforce collusion in

the same way as increased monitoring as a result of increased signal accuracy. The next section provides more detail as to how, *ceteris paribus*, δ and μ determine the existence of a TISC equilibrium.

It is important to distinguish the importance of proposition 2 from the more traditional results in collusion models that characterize the collusive equilibrium in terms of the discount rate. First, it is true that the punishment strategy under the TISC scheme is the same as in traditional collusion models: common knowledge of a deviation leads to Nash reversion. However, traditional collusion models often relate the discount rate with the informativeness of the signal to characterize the collusive equilibrium. In general, the more noise in the monitoring, the harder it is for firms to collude. However, proposition 2 relates the discount rate with the parameters that govern the *timing* of signals rather than the content of the signals. So while the qualitative characteristics of the punishment scheme in our model are similar to traditional models of collusion, our model relates the discount rate with the temporal uncertainty rather than the contextual uncertainty inherent in the monitoring. For more discussion of the qualitative differences between the event-driven model and discrete time models of variable period length, see the discussion at the end of section 4.

We also mentioned that allowing firms to change their price at random times when they do not receive information would not change the main results. We now informally argue such a claim. It is true that in the case of a monopoly with only one firm, that the firm may change its price at times in between the times that it receives information. More formally, the monopolist would hold beliefs about the underlying reserve price and then choose its price such that it is acting optimally, given its beliefs. In the two firm case, this is also true; firms *would* choose to change their prices at times when they don't receive information. However, we claim that although firms would choose to change their prices at times where they do not receive information, for sufficiently high values of μ , the firms would still have no incentive to break the collusive scheme when they are uncertain about the underlying reserve price. That is, if firms had the opportunity to change their price between the times they receive information, they may choose to do so but would also truthfully tell the other firm that they are changing their price and the other firm would then also change its price.

To see why, note that the continuation value of a deviation from the collusive scheme when a firm is uncertain about the reserve price is upper bounded by the

continuation value of the most profitable deviation when a firm does know the reserve price exactly. Analogously, the continuation value to a firm for colluding when the firm is uncertain about the reserve price is lower-bounded by the lowest continuation value along the equilibrium path (the exact continuation values are given in the appendix). Importantly, the lowest continuation value along the equilibrium path is strictly greater than 0 for all values of the parameters. However, as $\mu \rightarrow \infty$, the continuation values for all deviations from the collusive scheme approach zero. This means that for sufficiently high values of μ , the continuation value for a deviation when a firm is uncertain about the price can be made arbitrarily close to 0 while at the same time, the continuation value for colluding remains bounded away from 0. Therefore as long as the rate of public price signals is high enough, firms would not have an incentive to break from the collusive scheme when they have the opportunity to act at times where they do not receive information.

4 Equilibrium Characterization and Comparative Statics

A central feature of any event-driven game are the stochastic processes that govern it. Those processes combine to determine the incentives faced by the decision makers. Adjusting the underlying parameters can change the characteristics of the resulting game in extreme ways. For example, in the model considered here, if μ is such that the expected wait between public price signals is very long relative to time-discounting, then the role of public monitoring is diminished, thereby threatening the truth-telling and collusion incentives. Similarly, if γ is such that the reserve price changes very quickly relative to the private signals the firms receive, then the price-adjusting incentive is diminished. So the model analyzed here represents a broad spectrum of incentive regimes. In this section we demonstrate how changes in the stochastic processes combine to support or destroy the TISC equilibrium. To begin, we formalize the notion that the existence of the TISC equilibrium only depends on $\Lambda = \lambda_1 + \lambda_2$:

Proposition 3. *The existence of the TISC equilibrium only depends on the sum $\Lambda = \lambda_1 + \lambda_2$ and not the individual values of $\lambda_1 + \lambda_2$.*

This result may seem counter-intuitive since it is plausible that firms that receive

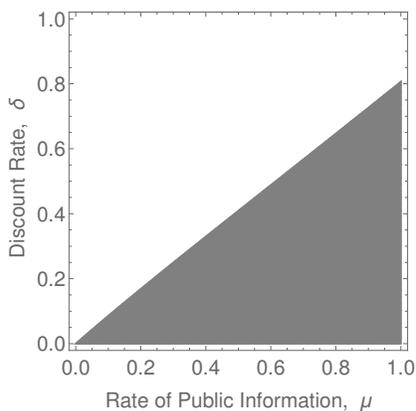


Figure 1: Collusive region in (δ, μ) space, $\gamma = 1$, $\Lambda = 2$, $\mathcal{R} = 2$

information faster should have more of an incentive to deviate in order to capitalize on superior information. This is not the case under the TISC equilibrium because the firm that receives information slower can hold information “hostage.” That is, firm j can make a credible threat to firm i that if i does not share information with j , then j will revert to marginal cost pricing. This is true even in the case when firm j *never* receives private information, i.e., $\lambda_j = 0$. Therefore, informational asymmetries in terms of the difference in rates do not affect the existence of a collusive equilibrium.

The next proposition formally characterizes the parameter region in which the TISC equilibrium exists in terms of the parameters μ and δ : A discussion of the proposition and an intuitive graphical representation follows.

Proposition 4. *Let $\mathcal{R} = 2$ and let $T(\delta, \gamma, \Lambda)$ be a mapping that gives the (possibly multiple) values of μ such that firm i is indifferent between participating in the TISC scheme and deviating. Then:*

1. T is a function
2. T is a bijection in δ when holding γ and Λ fixed
3. T is strictly increasing in δ
4. $\forall \gamma$ and Λ , $\lim_{\delta \rightarrow 0} T = 0$
5. $\forall \gamma$ and Λ , $\lim_{\delta \rightarrow \infty} T = \infty$
6. $\forall \gamma$ and Λ , $\lim_{\delta \rightarrow \infty} \frac{\partial T}{\partial \delta} = 1$

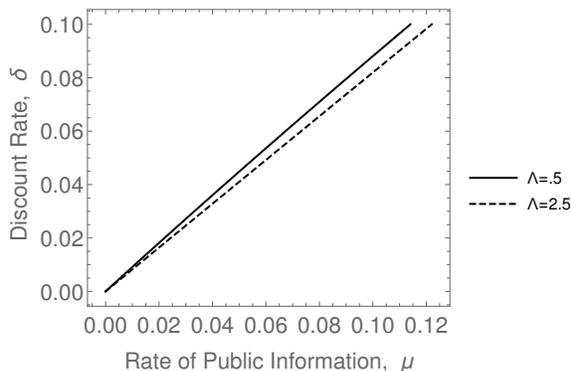


Figure 2: Change in TISC equilibrium region in (δ, μ) space when Λ changes. $\gamma = 1$ and $\mathcal{R} = 2$. Each line defines the TISC region where the area below the line is where the TISC equilibrium is supported.

Proposition 4 says that any increase in the discount rate that would potentially preclude the existence of the collusive equilibrium can be offset by an increase in the rate of public price signals. In other words, more impatient firms must be threatened by more frequent detection in order to sustain the collusive equilibrium. Figure 1 illustrates the central concepts of proposition 4. For fixed values of γ and Λ , the shaded region in figure 1 indicates the values of μ and δ in which the TISC equilibrium is supported. From the figure, it is clear that relatively more patient firms (lower δ) and relatively more frequent monitoring (high μ) support the TISC equilibrium. Furthermore, as the figure as well as item (6) in proposition 4 suggest, the tradeoff between the discount rate and the monitoring frequency approaches linear for some values of the parameters. Under that regime, an increase in δ can always be offset by a proportional increase in μ to maintain the TISC equilibrium.

Finally, we analyze the effects of a change in the rate of private demand signals via a proposition and a conjecture in which we present numerical evidence.⁸ We then discuss the economic interpretation of the effects of a change in Λ .

Proposition 5. *Let $\mathcal{R} = 2$ and fix γ . Then $\forall \Lambda > 0 \lim_{\delta \rightarrow 0} \frac{\partial^2 T}{\partial \delta \partial \Lambda} > 0$.*

Conjecture 1. *Let $\mathcal{R} = 2$. Then, for any value of γ and Λ , there exists a δ^* such that $\frac{\partial T}{\partial \Lambda} \leq 0$ if and only if $\delta > \delta^*$.*

Since there is a monotonic mapping between firms' patience and the rate of public price signals that define the collusive region, proposition 5 says that if price signals

⁸We characterize as conjectures those propositions for which we provide only a numerical proof.

are relatively infrequent, then increasing the rate of private demand signals *reduces* the incentive to collude. This result is illustrated in figure 2. In the figure, the dashed and solid lines define the region in (δ, μ) space that supports the TISC equilibrium. Since the solid line—which represents a slower rate of demand signals—is above the dotted line, increasing the rate of demand signals reduces the parameter region that supports the TISC scheme *when price signals are relatively infrequent*.

Conjecture 1 says that if the public price signals happen fast enough, then increasing the rate of private information *increases* the collusive region. The numerical evidence for conjecture 1 is given in figure 3. For various values of γ , figure 3 plots the value of δ that determines whether increasing the rate of private demand signals expands or contracts the collusive region. For example, fix $\gamma = 2$. Then, figure 3 says that when $\Lambda = .01$ —represented by the solid line—if δ is such that the firms are indifferent between participating in the TISC scheme and deviating *and* δ is greater than (approximately) 5, an increase in the rate of demand signals would incentive the firms to participate in the TISC scheme. On the other hand, if $\delta < 5$, then an increase in the rate of demand signals would incentive the firms to undercut one another. A complete example of proposition 5 and conjecture 1 is given in figure 4. In the figure, the blue area represents the gain in the collusive region as a result of an increase in the rate of demand signals and the red region represents the loss in the collusive region as a result of the same change in the rate of demand signals. The white areas are where the existence of the collusive region remained unchanged.

To understand how the existence of a TISC equilibrium depends on the rate of demand signals, it is necessary to separate how such an increase in Λ changes expected profits along the equilibrium path as well as to a potential deviator. Along the equilibrium path, increasing the rate of demand signals allows firms to more quickly adjust their prices in response to changes in the reserve price. As an extreme example, if firms learn about changes in the reserve price immediately they can always capture the benefit of higher reserve prices and never risk earning zero profits due to an unobserved jump to a low reserve price.

On the other hand, lemma 4 (in the appendix as part of the proof of 4) says that the most profitable deviation is the one in which a firm undercuts the market slightly upon learning that the reserve price is ρ^H . Under such a deviation the firm risks undercutting the market (and therefore risks detection) while earning zero profits since the reserve price may have jumped without the deviating firm receiving

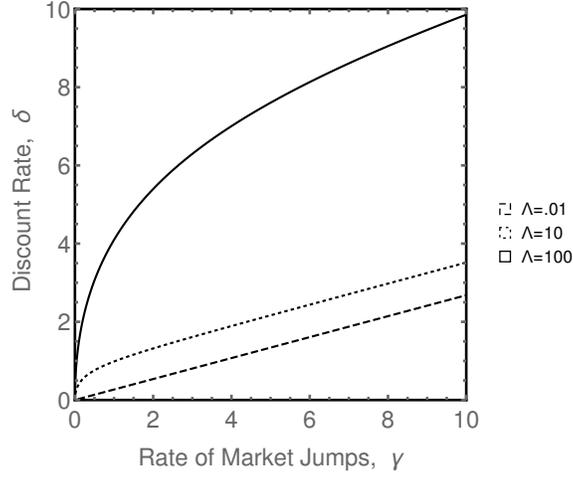


Figure 3: Regions in which $\frac{\partial T}{\partial \Lambda} < 0$. For each value of Λ —represented by either a dashed, dotted or solid line—the region in which $\frac{\partial T}{\partial \Lambda} < 0$ is the region *above* the line.

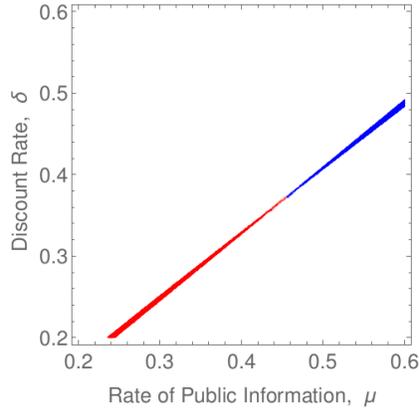


Figure 4: Change in TISC equilibrium region in (δ, μ) space when Λ changes from .5 to 2.5, $\gamma = 1$ and $\mathcal{R} = 2$. The blue region is the gain in the TISC equilibrium region as a result of the jump in Λ and the red region is the loss of the collusive region as a result of the jump. The unshaded area represents parameter values where the existence of a TISC equilibrium is unaffected.

a private demand signal. However, when the rate of demand signals increases, the average amount of time a potential deviator must endure earning zero profits while risking detection decreases. Therefore, increasing the rate of demand signals provides an additional incentive for firms to undercut since it reduces the amount of time a deviator spends "at risk" while earning zero profits.

In total, increasing the rate of demand signals increases the expected utility of firms in a TISC scheme because the firms can better capture demand fluctuations. On the other hand, increasing the rate of demand signals increases the expected utility of a potential deviator by reducing the amount of time the deviator risks detection while earning zero profits. The relative magnitudes of these two effects determine the change in the collusive region as a result of a change in the rate of demand signals. What proposition 5 and conjecture 1 say is that when public price signals are "fast", the reduction in the probability of detection to a potential deviator is relatively small. Therefore, the increase in expected utility from capturing demand fluctuations under the TISC scheme outweighs the affect of a reduction in the probability of detection. On the other hand, when price signals are relatively "slow," increasing the rate of demand signals has a relatively large reduction on the probability of detection and therefore increases a firm's incentive to undercut its opponent.

In addition to the direct insight provided by the comparative static analysis, the results in this section highlight how our results in an event-driven environment differ from simply adjusting the period length in a discrete time environment. For example, in a simple discrete time oligopoly setting firms have a per-period discount rate. If we assume that a period is δt units of time, then changing δt effectively modifies the *per-period* discount rate. As shown in the simplest version of the model of [Abreu et al., 1991], the strategic consequence of reducing the period length are *the same* as that of increasing firm's patience. Our model has a similar attribute; increasing firm's patience can have a similar consequence as decreasing the average time between public price revelations. However, our model has several differences in both the formulation and the results. In our formulation, the time between monitoring times is stochastic so firms need to account for the randomness in signal times. More substantially, an element that is not covered in many other variable length oligopoly models is the possibility for a firm undercuts its competitor and then reverts back to the collusive scheme without its competitor ever receiving any signal that it was undercut. Although we prove that in equilibrium, this does not occur, discrete time

models that just adjust the period length do not allow for this phenomenon.

Another crucial difference between the event-driven formulation and discrete time models with variable period length is that in our event-driven model, there are additional Poisson process that evolve along with the Poisson process that governs the public price revelations. For this reason, a change in the discount rate may not have the same effect as an equal change to the average rate of monitoring. This is best illustrated in figures 2 and 4 which show that the relative changes to the discount rate and the rate of monitoring required to preserve the TISC region depends on the rate at which firms receive demand information. More technically, the slope of the curve at any given point (the appendix shows that the relationship between δ and μ is non-linear) depends on the other rate parameters of the model. Yet in other terms, the rate of substitution between the average time between monitoring and the discount rate to preserve the TISC equilibrium is not 1, is not constant and depends on the values at which firms observe demand. To our knowledge, none of the oligopoly models that consider variable period lengths have uncovered such a result and our event-driven model is the first to highlight such considerations. This is an important result because many other oligopoly models argue that decreasing player's patience by changing the discount rate is equivalent to increasing the time period, often proportionally. However, our model shows that the relationship is not proportional and actually depends on the rate at which the underlying environment changes.

5 Robustness to Private Profit Signals

In the model presented above, the only profit information firms receive is through the public price signal, giving each firm an imperfect indication of profit rates. However, in many real-world markets, firms receive profit information in addition to demand and price information. This section shows that with mild additional assumptions, the TISC scheme is supported when firms also *privately* observe profits at random times. To illustrate the existence of a TISC equilibrium with private profit signals, we make the following three modifications to the original model presented in section 3.2:

1. We introduce a third reserve price, ρ^M . At random times, the reserve price jumps between ρ^H, ρ^M and ρ^L with $\rho^L < \rho^M < \rho^H$. It jumps from ρ^L to ρ^M at rate γ and ρ^H to ρ^M at rate γ and jumps from ρ^M to ρ^L and from ρ^M to ρ^H at rate $\frac{\gamma}{2}$. Such a rate specification implies that the rate at which demand

changes does not depend on the reserve price. The price never switches directly between ρ^L and ρ^H without traversing ρ^M . And the asymmetry going from ρ^L to ρ^H implies that the jump rate does not depend on the reserve price, but this is not necessary for our results. It is possible to specify any rate at which the reserve price jumps among prices but a constant jump rate simplifies the analysis. The set of private signals a firm can receive from the market indicating the reserve price is also expanded to include a signal for each reserve price.

2. $\rho^H, \rho^M > 0$ and $\rho^L < 0$. This says that two of the reserve prices are above the firms' marginal cost (normalized to zero) while the lowest reserve price is below marginal cost. The negative reserve price is not an issue if ρ_t is simply interpreted as price net costs. Of course, this is standard in competition with constant marginal cost.
3. At rates η_1, η_2 , firms receive a private signal indicating their profit rate.
4. Firm's message spaces are expanded so that they have the opportunity to completely correlate their messages with their observations, if they choose.

The main effect of these three additions is that the reserve price now jumps between three values. However, the values are such that a monopolist would not change its price to capture the market when the reserve price is ρ^L . The implications of such a modeling choice in a strategic scenario are discussed after proposition 6.

To facilitate analysis we also make one additional assumption:

Assumption 1. *Upon receiving a private profit signal, a firm can only change its price and/or send a message if the signal indicates that the firm is earning 0 profits.*

This assumption seems more restrictive than it is. Recall that if both firms are charging ρ^M but the reserve price jumps to ρ^H , any private profit signal to a firm will indicate that the firm's profits are $\frac{\rho^M}{2}$ and doesn't indicate whether the reserve price is ρ^H or ρ^M (though it would guarantee that the reserve price is not ρ^L). Therefore, the only reason firms would be willing to change their prices upon receiving a positive private profit signal is if they believe that demand has jumped up since their last opportunity to act. To make such a decision, a firm would have to form beliefs over trajectories of ρ_t to determine the probability that ρ_t equals ρ^H at the time it received its profit signal. However, this decision is *not* strategic in that a monopolist

would also have to form such beliefs. Furthermore, allowing this would unnecessarily comingle the role of the profit signal with that of the demand signal. Hence, we make this small simplification.

We now define a collusive equilibrium with private profit signals:

Definition 2. *A TISC scheme with private profit signals is a strategy profile in which:*

1. *When firm i receives a signal indicating the reserve price is ρ^H or ρ^M , firm i informs firm j of the reserve price and firm i sets its price equal to the reserve price.*
2. *When firm i receives a signal indicating the reserve price is ρ^L , firm i informs firm j of the reserve price and firm i sets its price to ρ^M .*
3. *When firm i receives a message from firm j indicating the reserve price is ρ^H or ρ^M , firm i sets its price to the reserve price. When firm i receives a message from firm j indicating the reserve price is ρ^L , firm i sets its price to ρ^M .*
4. *When firm i receives a price signal that indicates both firms are charging the same price, firm i does nothing. If the price signal indicates that the firms are charging different prices, firm i sets its price to 0 forever.*
5. *When firm i receives a signal indicating that its profits are zero, the recipient firm informs the other firm and both firms charge ρ^M .*

Definition 2 expands the TISC scheme in two main ways. First, firms do not adjust their price to all reserve prices but only to ρ^H and ρ^M . Instead of adjusting to ρ^L upon receiving a signal from the market, firms engage in cheap talk and set their price to ρ^M . Secondly, upon earning 0 profits, firms engage in cheap talk and both firms set their price to ρ^M . Proposition 6 establishes the existence of the TISC scheme with private profit signals:

Proposition 6. *Let $\Lambda = \lambda_1 + \lambda_2$ and $\eta = \eta_1 + \eta_2$. For any $\mu, \gamma, \Lambda, \eta > 0$ and $1 + \frac{2\gamma(\gamma+\Lambda+\eta)}{\gamma^2+4\gamma(\eta+\Lambda)+2(\eta+\Lambda)^2} < \frac{\rho^H}{\rho^M} < 3 + \frac{2\Lambda}{\gamma}$, there exists beliefs β^* and a $\underline{\delta}$ such that for any value of $\delta < \underline{\delta}$, the strategy profile s^* that specifies that firms participate in a TISC scheme with private profit signals is a sequential equilibrium.*

A version of proposition 2 can also be established for the case of private profit signals.

The proof of proposition 6 is exactly parallel to that of proposition 1. The element that makes this possible is that the low reserve price less than zero means that firms receive private signals indicating that they are earning zero profits *along the equilibrium path*. So in equilibrium, a firm that receives a signal indicating zero profits does not punish because the firm (correctly) believes it is a result of demand jumping to the low reserve price rather than being undercut by the other firm. This assumption is similar to what is known as a “full support” assumption that is standard in many models of repeated oligopoly (see [Mailath and Samuelson, 2006] for more details regarding the technical importance of full support assumptions). In contrast, suppose there were only two reserve prices, ρ^H and ρ^L , and that both prices were above the firms’ marginal cost. If firm j undercuts firm i while firm i is charging ρ^L and firm i receives a profit signal indicating zero profits, then firm i knows with certainty that firm j has deviated from the collusive equilibrium. However, it is not optimal for firm i to punish by adjusting to marginal cost pricing because firm j does not know that firm i discovered the deviation and therefore firm j will not enter the punishment phase. So the collusive equilibrium *based around grim trigger punishment* unravels, and it is the asynchronicity of private profit signals which does the unraveling, even though the signal is noiseless. Nevertheless, there may be other mechanisms (such as a carrot-stick type scheme) that enforce collusion when firms punish on private profit information. We leave a further discussion of such an extension to section 6.

Despite not punishing on private information, novel strategic considerations arise when firms receive private profit information. Specifically, consider the case where firm i deviates by undercutting slightly upon learning the reserve price is ρ^H . Of course, firm i risks being detected as a result of a public price signal. However, firm i faces the additional risk that firm j receives a private profit signal indicating that it is earning 0 profits and therefore changes its price to ρ^M . In this case, firm i would also have to set its price to ρ^M even though the reserve price might still be ρ^H . This element has nothing to do with firm i weighing the benefit of increased profits with the probability of being detected (since firm j does not punish upon a profit signal indicating 0 profit). Instead firm i faces the risk that firm j receives a signal and thinks that demand has actually decreased when it hasn’t, causing both firms to lower their prices and forego the benefits of high demand. As a result, the addition

of private profit signals increases the support of a TISC equilibrium.

6 Future Work

In this paper we extended the standard model of Bertrand competition to an environment in which firms receive private information and take actions at asynchronous and random times. We established that if players are patient enough *or* if price signals arrive fast enough, there exist collusive equilibria in which firms truthfully share private demand information and share monopoly profits. A central contribution was to show how the parameters that govern timing can be used to determine the existence of collusive equilibria similar to the way that parameters that govern preferences (i.e. the discount factor) determine the existence of collusive equilibria. We then showed how varying the frequency of signals can affect the collusive region. If players are patient, then increasing the rate of private information increases the incentive to collude. We then expanded our model to include private profit signals and showed there still exists a truthful information-sharing collusive equilibrium.

There are two obvious ways to extend such a model. The first is to add noise to firms' signals, i.e., a model with asynchronous and noisy signals. This would combine the event-driven formulation of Bertrand competition with the more common definition of imperfect private monitoring. A more general route for future work involves examining coordination mechanisms in models with asynchronous private information so that we can remove the assumption of synchronous price signals. As mentioned in section 5, the existence of a TISC equilibrium relies on firms not being able to privately detect a deviation along the equilibrium path. However, in real world markets it might be the case that a firm can privately and asynchronously detect a deviation. Intuition suggests that such a detection might lead to a “price war” through which firms would return to marginal cost pricing. But even if a firm could privately and asynchronously detect deviations, there may be an alternative punishment that deters deviations from occurring in the first place. However, such a punishment mechanism—for Bertrand games or repeated asynchronous games in general—is yet to be established.

There are at least two approaches to allowing private and asynchronous detection. One possibility is to introduce a strict full-support condition on the private profit signal, while removing the public price signals. For example, consider the case of

linear demand in which the intercept parameter evolves according to some stochastic (not necessarily Poisson) process with full support on the positive real numbers. Then the distribution of private profit signals has full support over the entire positive reals for any joint strategy profile. Therefore, firms never know with certainty whether or not they have been undercut and therefore never punish. Instead, the mechanism that would facilitate collusion would be similar to that in section 5. If firm 1 undercuts firm 2, then firm 2 believes that demand has decreased. So firm 2 would then lower its price, thus reducing firm 1's profits and its incentive to undercut in the first place. A second approach would be to examine alternative coordination mechanisms. In this paper, we exclusively considered the case of a "grim trigger" punishment strategy. However, there are other punishment schemes such as the "carrot-stick" scheme that might help to facilitate collusion [Athey and Bagwell, 2008].

In broader terms, the driving motivation of this work is to shed light on the strategic considerations that arise when players must reason about timing. Indeed, the results showed that when firms interact in an event-driven environment, the relative frequency of signals and actions have important strategic consequences. Since event-driven phenomena occur in many economic scenarios, future work should consider the strategic implications of stochastic timing across the spectrum of dynamic games.

Appendix A

All references to an online appendix can be found at http://jtgrana.com/index.php/supm_1/

Proposition 1 (formally): *For any actionable history $\tilde{h}_i \in \tilde{H}_i$ where there has not been a public price signal that revealed a deviation, let $m_{\tilde{h}_i}$ be the last element of the actionable history (which recall is the signal or message the player just received). Then define s^* as:*

$$s^*(h) = \begin{cases} (\rho^L, \emptyset) & \text{if } m_{\tilde{h}_i} = x^L \\ (\rho^H, \emptyset) & \text{if } m_{\tilde{h}_i} = x^H \\ (\rho^L, x^L) & \text{if } m_{\tilde{h}_i} = y^L \\ (\rho^H, x^H) & \text{if } m_{\tilde{h}_i} = y^H \\ (p_i, \emptyset) & \text{if } m_{\tilde{h}_i} = (p_i, p_j) \text{ and } p_i = p_j \\ (0, \emptyset) & \text{if } m_{\tilde{h}_i} = (p_i, p_j) \text{ and } p_i \neq p_j \end{cases} \quad (3)$$

and for any actionable history \tilde{h}_i where a public signal has revealed a deviation, $s^*(\tilde{h}_i) = (0, \emptyset)$. Then, if $1 + \frac{\gamma}{\gamma + \lambda_1 + \lambda_2} < \frac{\rho^H}{\rho^L} < 2 + \frac{\lambda_1 + \lambda_2}{\gamma}$, there exists $\underline{\delta}$ and set of beliefs β^* such that for all $\delta < \underline{\delta}$, then each player having assessment (s^*, β^*) defines a sequential equilibrium. That is, there exists a sequence $\Sigma = (s^1, \beta^1), (s^2, \beta^2) \dots \rightarrow (s^*, \beta^*)$ such that no firm can change its action at any actionable history and improve its expected utility given its beliefs and $\beta^\epsilon, \epsilon = 1, 2, \dots$ are derived using Bayes rule.

Proof. To prove the proposition, first we establish that along the equilibrium path, profits evolve according to a continuous time Markov Chain (CTMC). Then, we present the automaton representation of strategies under s^* . We then illustrate the continuation values in equilibrium and the possible deviations and show that s^* is a Perfect Bayesian (Nash) equilibrium. Finally, we present a sequence of strategies and consistent beliefs that converge to (s^*, β^*) and show that s^* is sequentially rational given β^* .

We begin with three lemmas:

Lemma 1. *Under s^* , the triple (p_1^t, p_2^t, r_t) is governed by a first order CTMC.*

Proof. The value of r_t jumps at times governed by a Poisson process with rate γ .

Under s^* the values of p_1^t, p_2^t change when either of the firm's receives a signal from the market. The signal is generated according to a Poisson process with rate $\lambda_1 + \lambda_2$ and the prices firms set only depends on the most recent signal and is independent of past private histories. Since changes in the state (r_t, p_1^t, p_2^t) are governed by Poisson processes, the inter-event times are distributed exponentially. As a result of exponentially distributed event times and history independent state transitions, (r_t, p_1^t, p_2^t) is governed by a first order CTMC. \square

Lemma 2. *For any strategy for firm i , let β^* be firm i 's belief of full game histories when firm j 's strategy is s^* . Then β^* is consistent if it places probability 1 on one and only one h_j^t at any actionable history \tilde{h}_i where a public signal has not revealed a deviation.*

Proof. Since according to s^* , firm j shares all private information through cheap talk, firm i knows the time and content of j 's observations as well as j 's actions. Therefore, i knows h_j^t with probability 1 (which implies firm i know p_j^t for all t). \square

Lemma 3. *For any strategy for firm i , let β^* be firm i 's belief of full game histories when firm j plays s^* and no public history has revealed a deviation. Suppose at time t , firm i observes anything except a public price revelation. Then β^* is consistent if it places probability 1 on a specific value of r_t .*

Proof. At history \tilde{h}_i^t where $\omega_t \neq x^L, x^H$, firm i either received a noiseless signal from the market or a truthful message from firm j . In either case, Bayes rule prescribes that β_i assigns probability 1 to one value of r_t at \tilde{h}_i^t . \square

Lemmas 2 and 3 state that at any actionable history, firm i knows firm j 's private history and the current reserve price exactly. Assuming no public signal has revealed a deviation, the strategy s^* and lemmas 1-3 imply that

1. Firm i 's continuation value at any history only depends on the current value of r_t and not on the entire history of r_t .
2. At any actionable history, firm i knows r_t via the message it receives.
3. Firm i knows what observation j received at any actionable history as well as firm j 's current price.

Consequently, when firm j plays s^* with consistent beliefs β^* as defined in lemmas 2 and 3, it is possible to adopt a finite automaton representation of firm i 's strategies where i 's automaton states are defined by its observation at an actionable history *and* whether or not a public signal has revealed a deviation.⁹ The continuation values under s^* are expressed in equation set 4. The notation $V_k(m)$ indicates the continuation value for firm i at automaton state defined by the observation m . $V_k(\emptyset)$ represents a state in the CTMC where firms do not take an action and is only expressed for clarity.

$$\begin{aligned}
V_1(m \in \{x^L, y^L\}) &= \mathbf{E}_{J_1} \left[\int_0^{J_1} \frac{\rho^L}{2} e^{-\delta t} dt + e^{-\delta J_1} V_2 \right] && \text{(Learns reserve price is } \rho^L \text{)} \\
V_2(\emptyset) &= \mathbf{E}_{J_2} \left[\int_0^{J_2} \frac{\rho^L}{2} e^{-\delta t} dt + \right. \\
&\quad \left. e^{-\delta J_2} \left(\frac{\gamma}{\gamma + \lambda_1 + \lambda_2} V_1 + \frac{\lambda_1 + \lambda_2}{\lambda_1 + \lambda_2 + \gamma} V_3 \right) \right] && \text{(Reserve Price jumps to } \rho^H \text{)} \\
V_3(m \in \{x^H, y^H\}) &= \mathbf{E}_{J_3} \left[\int_0^{J_3} \frac{\rho^H}{2} e^{-\delta t} dt + e^{-\delta J_3} V_4 \right] && \text{(Learns reserve price is } \rho^H \text{)} \\
V_4(\emptyset) &= \mathbf{E}_{J_4} \left[0 + \right. \\
&\quad \left. e^{-\delta J_4} \left(\frac{\lambda_1 + \lambda_2}{\lambda_1 + \lambda_2 + \gamma} V_1 + \frac{\gamma}{\lambda_1 + \lambda_2 + \gamma} V_3 \right) \right] && \text{(Reserve price jumps to } \rho^H \text{)}
\end{aligned} \tag{4}$$

The operator \mathbf{E}_{J_i} represents the expectation taken over the amount of time the system is in the state indicated by V_i . The parenthetical comments briefly describe the event that has just occurred that gives the continuation value indicated by the equation.

By applying the Laplace-Stieltjes transform, the above equations can be rewritten

⁹See [Mailath and Samuelson, 2006] for a detailed treatment of automaton representations.

as¹⁰:

$$\begin{aligned}
V_1(m \in \{x^L, y^L\}) &= \frac{\rho^L}{2(\gamma + \delta)} + \frac{\gamma}{\gamma + \delta} V_2 \\
V_2(\emptyset) &= \frac{\rho^L}{2(\gamma + \lambda_1 + \lambda_2 + \delta)} + \frac{\gamma}{\gamma + \lambda_1 + \lambda_2 + \delta} V_1 + \frac{\lambda_1 + \lambda_2}{\lambda_1 + \lambda_2 + \gamma + \delta} V_3 \\
V_3(m \in \{x^H, y^H\}) &= \frac{\rho^H}{2(\gamma + \delta)} + \frac{V_4}{\gamma + \delta} \\
V_4(\emptyset) &= 0 + \frac{\lambda_1 + \lambda_2}{\lambda_1 + \lambda_2 + \gamma + \delta} V_1 + \frac{\gamma}{\lambda_1 + \lambda_2 + \gamma + \delta} V_3
\end{aligned}$$

Holding firm j 's strategy fixed, the one-stage deviation principal says that s^* is optimal for firm i if there is not a one-time deviation for firm i at V_1 , V_3 or any other state of the CTMC reachable from V_1 or V_3 that increases firm i 's continuation value.¹¹ However, the only automaton state reachable from V_1 and V_3 not along the equilibrium path is the one that includes a public revelation of prices where the firms charged different prices. In this case it is (weakly) dominant for firms to set their price to 0. Therefore, the only deviations that must be evaluated are the various one-shot deviations at V_1 and V_3 . All possible deviations can be condensed into 1 of 5 possible continuation values, which are summarized in box 1. We now analyze all possible one-shot deviations at actionable histories in turn.

¹⁰For a complete review of value functions and cost/reward computation in continuous time Markov chains, see the excellent treatment given in chapter 6 of [Kulkarni, 2009].

¹¹The use of the one stage deviation principal in this context is formally justified by the one-stage deviation principal for semi-markov processes which is given in [Stone, 1973].

Case 1: Undercutting ρ^L slightly

1. Both firms are charging ρ^H . Firm i receives message y^L and does not tell firm j that the market price has jumped and charges ρ^L
2. Firm i receives message y^L and tells firm j that the market price has jumped to ρ^L and undercuts ρ^L slightly.
3. Firm i receives message x^L from firm j and undercuts ρ^L slightly.
4. Firm i receives message y^L and tells firm j that the market price has jumped to ρ^H and undercuts ρ^L slightly.

Case 2: Undercutting ρ^H slightly

1. Firm i receives message y^H and tells firm j that the market price has jumped to ρ^H and undercuts ρ^H slightly.
2. Firm i receives message x^H from firm j and undercuts ρ^H slightly.
3. Both firms are charging ρ^H . Firm i receives message y^H and undercuts ρ^H slightly.

Case 3: Undercutting ρ^H by charging ρ^L

1. Firm i receives message y^H and tells firm j that the market price has jumped to ρ^H and undercuts ρ^H at ρ^L (or slightly less than ρ^L).
2. Firm i receives message x^H from firm j and undercuts ρ^H at ρ^L (or slightly less than ρ^L).

Case 4: Preference to collude at ρ^L

1. Firm i receives message y^H and tells firm j that the market price is ρ^L and charges ρ^L .
2. Firm i receives message y^H , knows that firm j 's price is ρ^L and does not say anything and charge ρ^L .

Case 5: Preference to collude at ρ^H

1. Firm i receives message y^L and tells firm j that the market price is ρ^H and charges ρ^H .
2. Firm i receives message y^L , knows that firm j 's price is ρ^H and does not say anything and charge ρ^H .

Box 1: Categorization of all possible deviations by continuation value

Case 1: Undercutting ρ^L slightly (deviating at V_1) The continuation values when i deviates at V_1 by undercutting slightly are given by:

$$\begin{aligned}
V_1'(m \in \{x^L, y^L\}) &= \mathbf{E}_{J_1} \left[\int_0^{J_1} \rho^L e^{-\delta t} dt + \right. \\
&\quad \left. e^{-\delta J_1} \left(\frac{\lambda_1 + \lambda_2}{\lambda_1 + \lambda_2 + \gamma + \mu} V_1 + \frac{\gamma}{\lambda_1 + \lambda_2 + \gamma + \mu} V_5 + \frac{\mu}{\lambda_1 + \lambda_2 + \gamma + \mu} V_6 \right) \right] \\
V_5(\emptyset) &= \mathbf{E}_{J_2} \left[\int_0^{J_2} \rho^L e^{-\delta t} dt + \right. \\
&\quad \left. e^{-\delta J_2} \left(\frac{\lambda_1 + \lambda_2}{\lambda_1 + \lambda_2 + \gamma + \mu} V_3 + \frac{\gamma}{\lambda_1 + \lambda_2 + \mu + \gamma} V_1' + \frac{\mu}{\lambda_1 + \lambda_2 + \mu + \gamma} V_6 \right) \right] \\
V_6(m = (p_i, p_j), p_i \neq p_j) &= 0
\end{aligned}$$

Note the new action state V_6 , which is the actionable state when a public price revelation indicates that i has deviated. In that case it is weakly dominant for firm i to charge 0. By again applying the Laplace-Stieltjes transform and solving the system of equations for V_1 and V_1' , we get the necessary condition for i not to undercut upon learning the reserve price ρ^L :

$$V_1 - V_1' > 0$$

or equivalently

$$\sum_{\alpha=0}^5 C_\alpha \delta^\alpha > 0 \tag{5}$$

where

$$C_0 = \gamma \mu \Lambda (2\gamma + \mu + \Lambda) (\gamma (\rho^H + 2\rho^L) + (\rho^H + \rho^L) \Lambda)$$

and $\Lambda = \lambda_1 + \lambda_2$ and $C_\alpha, \alpha \neq 0$ are given in the online supplementary materials.

Case 2: Undercutting ρ^H slightly (Deviation at V_3) The continuation values when i undercuts the market slightly upon learning the reserve price is ρ^H are given by:

$$\begin{aligned}
V_3'(m \in \{x^H, y^H\}) &= \mathbf{E}_{J_3} \left[\int_0^{J_1} \rho^H e^{-\delta t} dt + e^{-\delta J_3} \left(\frac{\mu}{\mu + \gamma} V_6 + \frac{\gamma}{\gamma + \mu} V_7 \right) \right] \\
V_7(\emptyset) &= \mathbf{E}_{J_5} \left[0 + \right. \\
&\quad \left. e^{-\delta J_5} \left(\frac{\lambda_1 + \lambda_2}{\lambda_1 + \lambda_2 + \gamma + \mu} V_1 + \frac{\gamma}{\lambda_1 + \lambda_2 + \gamma + \mu} V_3' + \frac{\mu}{\lambda_1 + \lambda_2 + \mu + \gamma} V_6 \right) \right]
\end{aligned}$$

The necessary condition for i not to deviate is then given by:

$$V_3 - V_3' > 0$$

or equivalently

$$\sum_{\alpha=0}^5 D_\alpha \delta^\alpha > 0 \tag{6}$$

where

$$D_0 = \gamma \mu \Lambda (2\gamma + \mu + \Lambda) (\gamma (\rho^H + 2\rho^L) + (\rho^H + \rho^L) \Lambda) = C_0 \tag{7}$$

and $D_\alpha, \alpha \neq 0$ are given in the online supplementary materials.

Case 3: Undercutting ρ^H by pricing at ρ^L (deviating at V_3) The continuation values when i undercuts the market by charging ρ^L upon learning the price is ρ^H and tells j the price is ρ^H are given by:

$$\begin{aligned}
V_3''(m \in \{x^H, y^H\}) &= \mathbf{E}_{J_3} \left[\int_0^{J_1} \rho^L e^{-\delta t} dt + e^{-\delta J_3} \left(\frac{\mu}{\mu + \gamma} V_6 + \frac{\gamma}{\gamma + \mu} V_8 \right) \right] \\
V_8(\emptyset) &= \mathbf{E}_{J_5} \left[\int_0^{J_5} \rho^L e^{-\delta t} + \right. \\
&\quad \left. e^{-\delta J_5} \left(\frac{\lambda_1 + \lambda_2}{\lambda_1 + \lambda_2 + \gamma + \mu} V_1 + \frac{\gamma}{\lambda_1 + \lambda_2 + \gamma + \mu} V_3'' + \frac{\mu}{\lambda_1 + \lambda_2 + \mu + \gamma} V_6 \right) \right]
\end{aligned}$$

The necessary condition for i not to deviate is then given by

$$V_3 - V_3'' > 0$$

or equivalently

$$\sum_{\alpha=0}^5 E_{\alpha} \delta^{\alpha} \quad (8)$$

where

$$E_0 = \gamma \mu \Lambda (2\gamma + \mu + \Lambda) (\gamma (\rho^H + 2\rho^L) + (\rho^H + \rho^L) \Lambda) = C_0 = D_0 > 0 \quad (9)$$

and $E_{\alpha}, \alpha \neq 0$ are given in the online supplementary materials.

Case 4: Firm i charges ρ^L , tells firm j that the reserve price is ρ^L upon learning the reserve price is ρ^H .) The continuation values to firm i by making a one-shot deviation in which it tells firm j that the reserve price is ρ^L and charges ρ^L although firm i knows the reserve price is ρ^H are given by:

$$V_3'''(m = y^H) = \mathbf{E}_{J_3} \left[\int_0^{J_2} \frac{\rho^L}{2} e^{-\delta t} dt + e^{-\delta J_2} \left(\frac{\gamma}{\gamma + \lambda_2} V_1 + \frac{\lambda_1 + \lambda_2}{\lambda_1 + \lambda_2 + \gamma} V_3 \right) \right] \quad (10)$$

Applying the Laplace-Stieltjes transform gives the necessary condition for i not to have an incentive to deviate, which is given by:

$$V_3 - V_3''' > 0$$

or equivalently

$$\mathcal{R} > 1 + \frac{\gamma}{\gamma + \delta + \Lambda} \quad (11)$$

Case 5: Firm i charges ρ^H , tells firm j that the reserve price is ρ^H upon learning the reserve price is ρ^L .) The continuation value for a one time deviation by firm i by telling firm j that the reserve price is ρ^H when it is actually ρ^L is given

by:

$$V_1''(m \in \{x^L\}) = \mathbf{E}_{J_1} \left[0 + e^{-\delta J_1} \left(\frac{\gamma}{\gamma + \lambda_1 + \lambda_2} V_3 + \frac{\lambda_1 + \lambda_2}{\gamma + \lambda_1 + \lambda_2} V_1 \right) \right] \quad (12)$$

Applying the Laplace-Stieltjes transform gives the necessary condition for i not to have an incentive to deviate, which is given by:

$$V_1 - V_1'' > 0$$

or equivalently

$$\mathcal{R} < 2 + \frac{\delta + \lambda_1 + \lambda_2}{\gamma} \quad (13)$$

As δ goes to 0, the left hand side of equations 5, 6 and 8 go to $C_0 = D_0 = E_0 > 0$, equation 11 goes to $\gamma\lambda_1\Lambda(2\gamma + \Lambda)(\gamma(\rho^H - 2\rho^L) + \Lambda(\rho^H - \rho^L))$ and equation 13 goes to $\gamma\lambda_1\Lambda(2\gamma\Lambda)(-\gamma\rho^H + (2\gamma + \Lambda)\rho^L)$. Since the left hand side of equations 5, 6, 8, 11 and 13 are all continuous in δ , there exists a $\underline{\delta}$ such that for all $\delta < \underline{\delta}$ the conditions in equations 5, 6, 8, 11 and 13 are met as long as $\frac{2\gamma + \Lambda}{\gamma + \Lambda} < \frac{\rho^H}{\rho^L} < \frac{2\gamma + \Lambda}{\gamma}$. It is not necessary to check if firm i should deviate upon a public revelation of prices since the continuation value of such a deviation would be a convex combination of continuation values of deviating upon learning the reserve price exactly. This implies that for a low enough δ , s^* and β^* as described in lemmas 2 and 3 is a Perfect Bayesian Equilibrium (i.e. s^* is a best response to s^* and beliefs along the equilibrium path are derived through Bayes rule).

To establish that s^* is also a sequential equilibrium, it is necessary to show that s^* is sequentially rational off of the equilibrium path. Define $\Sigma, (s^1, \beta^1), (s^2, \beta^2) \dots$ indexed by ϵ to be a sequence of strategy profiles and associated consistent beliefs (for both firms) such that at any actionable history:

1. Firms play the action prescribed by s^* with probability $1 - \frac{1}{\epsilon}$.
2. Firms randomize uniformly over all actions (prices and messages) with probability $\frac{1}{\epsilon}$

As $\epsilon \rightarrow \infty$, $s^\epsilon \rightarrow s^*$ and along the equilibrium path β_i^ϵ converges to placing

probability 1 on values of r_t and \tilde{h}_j^t (since the probability of firm j doing anything other than what is prescribed by s^* at or before time t goes to 0 as $\epsilon \rightarrow \infty$). Therefore, s^ϵ converges to s^* and *along the equilibrium path*, β_i^ϵ converges to β^* .

Lastly, it is necessary to verify that s^* is sequentially rational off of the equilibrium path given beliefs β^* . The only actionable histories for firm i off of the equilibrium path are ones that include an undetected deviation by firm i and ones that have revealed a deviation by either firm.

Consider any actionable history \tilde{h}_i^t where i has deviated but not been detected. By lemmas 2 and 3, β^* places probability 1 on one value of \tilde{h}_j^t and r_t . Therefore, firm i 's continuation value from playing s^* after an undetected deviation can be represented by either V_1 or V_3 , both of which firm i would not choose to deviate from s^* . Therefore s^* is sequentially rational at any \tilde{h}_i^t in which firm i has deviated but not been detected.

For any actionable history \tilde{h}_i^t that has revealed a deviation, β_i^* specifies some distribution over past histories. However, since s^ϵ converges to s^* , upon realizing a deviation, it is sequentially rational for firm i to play 0 regardless of the belief of past histories.

□

Proof of Proposition 2. The proof is exactly the same as the proof of 1 with the exception that equations 5,6,8 can be written as quadratics in μ and equations 11 and 13 can be rearranged to establish bounds on \mathcal{R} in terms of δ .

□

Proof of Proposition 3. After performing the variable substitution $\Lambda = \lambda_1 + \lambda_2$ in equations 5,6,8, 11 and 13, no λ_1 or λ_2 terms remain. See the online supplementary material for a full definition of the terms.

□

Proof of Proposition 4. To prove the 6 statements in proposition 4 that establish the equilibrium characteristics, we first state a lemma:

Lemma 4. *Suppose $\mathcal{R} = 2$ so that $1 + \frac{\gamma}{\gamma + \lambda_1 + \lambda_2} < \mathcal{R} < 2 + \frac{\lambda_1 + \lambda_2}{\gamma}$ is true for all values of the parameters. Then, the TISC equilibrium exists if firm i does not have an incentive to undercut firm j slightly upon receiving a message from firm j that the reserve price is ρ^H . Formally, if $V_3 - V_3' > 0$ then $V_3 - V_3'' > 0$ and $V_1 - V_1' > 0$.*

Proof. Assume that $V_3 - V_3' > 0$. Then

$$(V_1 - V_1' - (V_3 - V_3')) = \delta(\delta + 2\gamma - \mu)\rho^L(\delta + \Lambda)(\delta + 2\gamma + \Lambda)(\delta + \mu + \Lambda).$$

This implies that $V_1 - V_1'$ is less than $V_3 - V_3'$ only if $\mu > \delta + 2\gamma$. Since $V_1 - V_1'$ is increasing in μ (see the online supplementary material for the full representation), it suffices to show that at $\mu = 2\gamma + \delta$, $V_1 - V_1' > 0$. Making the appropriate substitution yields $V_1 - V_1' = \gamma(\delta + 2\gamma)\rho^L(\delta + \Lambda)(4(\delta + 2\gamma)^2 + (9\delta + 16\gamma)\Lambda + 3\Lambda^2) > 0$. To show that $V_3 - V_3' > 0$ implies that $V_3 - V_3'' > 0$, it can be shown that $V_3 - V_3'' - (V_3 - V_3') = 2\delta(\delta + 2\gamma)\rho^L(\delta + \Lambda)(\delta + 2\gamma + \Lambda)(\delta + \mu + \Lambda) > 0$. Therefore, whenever $V_3 - V_3' > 0$, $V_3 - V_3'' > 0$. \square

Lemma 4 simplifies the analysis since it says that a necessary and sufficient condition for the existence of the collusive region is that $V_3 - V_3' > 0$ since this condition implies that all other necessary conditions are met. In other words, $V_3 - V_3' = 0$ defines the collusive region. We can now prove claims 1-6 in proposition 4 in order.

1. **T is a function:** Since $V_3 - V_3'$ can be written as a quadratic in μ , applying the quadratic formula to $V_3 - V_3' = 0$ yields:

$$T(\delta, \gamma, \Lambda) = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

where

$$\begin{aligned} a &= \rho^L(\delta + \Lambda)(2\delta^2 + 2\delta(3\gamma + \Lambda) + \gamma(4\gamma + 3\Lambda)) \\ b &= \rho^L(\delta + \Lambda)(2\delta^2\Lambda + \gamma(2\gamma + \Lambda)(4\gamma + 3\Lambda) + 2\delta(2\gamma^2 + 4\gamma\Lambda + \Lambda^2)) \\ c &= -2\delta(\delta + 2\gamma)\rho^L(\delta + \Lambda)(\delta + \gamma + \Lambda)(\delta + 2\gamma + \Lambda). \end{aligned} \quad (14)$$

Since $a > 0$, $b > 0$ and $c < 0$, there exists exactly one positive root for all values for a fixed value of δ, γ and Λ . Therefore, T is a function since for fixed values of δ, γ and Λ , T returns only one value of μ .

2. **T is a bijection in δ when holding γ and Λ fixed:** We need to show that T is one-to-one and onto. To show that T is onto, we need to show that for any fixed μ , there exists a value of δ that makes $V_3 - V_3' = 0$. Since all parameters are greater than 0, $V_3 - V_3'$ can be written as:

$$\begin{aligned}
V_3 - V'_3 = & -\rho^H \delta^5 - (5\gamma\rho^H + 3\Lambda\rho^H)\delta^4 + 2\rho^L (-8\gamma^2 + \mu^2 - 12\gamma\Lambda + \mu\Lambda - 3\Lambda^2) \delta^3 \\
+ & \left[-8\gamma^3\rho^L + 4\gamma^2\rho^L(\mu - 7\Lambda) - \right. \\
& \left. 2\rho^L\Lambda(-2\mu^2 - 2\mu\Lambda + \Lambda^2) + \gamma(6\mu^2\rho^L + 8\mu\rho^L\Lambda - 18\rho^L\Lambda^2) \right] \delta^2 \\
+ & \left[8\gamma^3\rho^L(\mu - \Lambda) + 2\mu\rho^L\Lambda^2(\mu + \Lambda) + \right. \\
& \left. 2\gamma^2(2\mu^2\rho^L + 7\mu\rho^L\Lambda - 6\rho^L\Lambda^2) + \gamma\Lambda(9\mu^2\rho^L + 11\mu\rho^L\Lambda - 4\rho^L\Lambda^2) \right] \delta \\
+ & \gamma\mu\Lambda(2\gamma + \mu + \Lambda)(4\gamma\rho^L + 3\Lambda)\rho^L \tag{15}
\end{aligned}$$

For a fixed μ , as $\delta \rightarrow \infty$, $V_3 - V'_3 \rightarrow -\infty$ and as $\delta \rightarrow 0$, $V_3 - V'_3 \rightarrow \gamma\mu\Lambda(2\gamma + \mu + \Lambda)(4\gamma\rho^L + 3\Lambda)\rho^L > 0$. Since $V_3 - V'_3$ is continuous, by the intermediate value theorem there is at least one value of δ such that $V_3 - V'_3 = 0$ for any fixed μ . This proves T is onto. To show that T is one-to-one, we need to show that for a fixed value of μ , there are not multiple values of δ that make $V_3 - V'_3 = 0$. To show that this is the case, we remind the reader of a well-known result:

Definition 3 (Descartes Rule of Sign Changes). *For any polynomial arranged in decreasing degree, the number of positive roots is upper bounded by the number of sign changes of the coefficients. Furthermore, the number of positive roots is either equal to the number of sign changes or less than the number of sign changes by an even amount.*

Note that the sign of the coefficients on δ^5, δ^4 are unambiguously negative and the constant in equation is unambiguously positive in equation 15. Therefore, the number of positive roots of equation 15 is either 1 or 3. There are three roots if the signs of the coefficients of δ^3, δ^2 and δ^1 are one of $' + - +'$, $' + - -'$, $+ + -$ or $' - + -'$ (other combinations of signs yield 1 positive root). Let D_3, D_2 , and D_1 be the coefficients on δ^3, δ^2 and δ^1 respectively. We now show that none of the three possible sign patterns are possible.

- (a) Assume $D_3 > 0$ then it must be that $\mu > \frac{-\Lambda + \sqrt{32\gamma^2 + 48\gamma\Lambda + 13\Lambda^2}}{2}$. Since D_2 is increasing in μ for positive values of μ , for it to be the case that $D_3 > 0$ and $D_2 < 0$ it must be that at $\mu = \frac{-\Lambda + \sqrt{32\gamma^2 + 48\gamma\Lambda + 13\Lambda^2}}{2}$, $D_2 < 0$. Plugging in for μ in D_2 gives

$$D_2 = \rho^L (2\gamma + \Lambda) \left(20\gamma^2 + 10\Lambda^2 + \gamma \left(27\Lambda + \sqrt{32\gamma^2 + 48\gamma\Lambda + 13\Lambda^2} \right) \right) > 0.$$

Therefore, it is impossible for $D_3 > 0$ and $D_2 < 0$. This proves that it is impossible for the coefficients to have signs $' + --'$ and $' + -+'$.

- (b) Assume $D_2 > 0$, then it must be that $\mu > \frac{-\gamma^2 - 2\gamma\Lambda - \Lambda^2 + \sqrt{13\gamma^4 + 54\gamma^3\Lambda + 61\gamma^2\Lambda^2 + 25\gamma\Lambda^3 + 3\Lambda^4}}{3\gamma + 2\Lambda}$. Since D_1 is increasing in μ for positive values of μ , for it to be the case that $D_1 < 0$ and $D_2 > 0$, it must be that at $\mu = \frac{-\gamma^2 - 2\gamma\Lambda - \Lambda^2 + \sqrt{13\gamma^4 + 54\gamma^3\Lambda + 61\gamma^2\Lambda^2 + 25\gamma\Lambda^3 + 3\Lambda^4}}{3\gamma + 2\Lambda}$. Plugging this value in for μ in D_1 gives

$$D_1 = \frac{\rho^L}{(3\gamma + 2\Lambda)^2} \left[32\gamma^6 + 180\gamma^5\Lambda + 4\Lambda^6 + \gamma^4(413\Lambda^2 + 16\phi) + \gamma^3(415\Lambda^3 + 14\Lambda\phi) + \gamma^2(194\Lambda^4 + 13\Lambda^2\phi) + \gamma(42\Lambda^5 + 2\Lambda^3\phi) \right] \quad (16)$$

where $\phi = \sqrt{13\gamma^4 + 54\gamma^3\Lambda + 61\gamma^2\Lambda^2 + 25\gamma\Lambda^3 + 3\Lambda^4}$. Since all parameters are greater than 0, $D_1 > 0$ and therefore it is impossible for $D_1 < 0$ while $D_2 > 0$ and the signs of the coefficients cannot be $' - +-'$ or $' + +-'$.

Since for a fixed μ there is only 1 positive δ such that $V_3 - V_3' = 0$, T is one-to-one.

3. **T is strictly increasing in δ :** Since T is a continuous bijection it must be strictly monotone. It is straightforward to show that T is monotone increasing in δ .
4. **$\lim_{\delta \rightarrow 0} T = 0$:** Setting $\delta = 0$ in equation 14 and simplifying shows that $\lim_{\delta \rightarrow 0} T = 0$.
5. **$\lim_{\delta \rightarrow \infty} T = \infty$:** The highest order δ term under the radical in equation 14 is δ^8 and it has a positive coefficient. The highest order δ term in b is δ^3 . Therefore, $\lim_{\delta \rightarrow \infty} T = \infty$.

6. $\lim_{\delta \rightarrow \infty} \frac{\partial T}{\partial \delta} = 1$: The expression for $\frac{\partial T}{\partial \delta}$ is given in the online supplementary material. Taking the limit as $\delta \rightarrow \infty$ gives the result. □

Proof of Proposition 5. See online supplementary material. □

Proof of Proposition 6. As in proposition 1, the market and prices evolve according to a continuous time Markov Chain under the TISC scheme. The only actionable automaton state not reachable along the equilibrium path is one in which a public price signal has revealed a deviation. Therefore, the same argument posed in 1 about beliefs (both on-path and off-path) apply in this case. Therefore, it is only necessary to check on-path deviations. See the online supplementary material for a fully detailed treatment. □

References

- [Abreu et al., 1991] Abreu, D., Milgrom, P., and Pearce, D. (1991). Information and timing in repeated partnerships. *Econometrica*, 59(6):1713–1733.
- [Athey and Bagwell, 2008] Athey, S. and Bagwell, K. (2008). Collusion with persistent cost shocks. *Econometrica*, 76(3):493–540.
- [Awaya and Krishna, 2015] Awaya, Y. and Krishna, V. (2015). On communication and collusion. *The American Economic Review*.
- [Ben-Porath and Kahneman, 1996] Ben-Porath, E. and Kahneman, M. (1996). Communication in repeated games with private monitoring. *Journal of Economic Theory*, 70(2):281 – 297.
- [Chan and Zhang, 2015] Chan, J. and Zhang, W. (2015). Collusion enforcement with private information and private monitoring. *Journal of Economic Theory*, 157:188–211.
- [Collard-Wexler, 2013] Collard-Wexler, A. (2013). Demand fluctuations in the ready-mix concrete industry. *Econometrica*, 81(3):1003–1037.
- [Colombo and Labrecciosa, 2015] Colombo, L. and Labrecciosa, P. (2015). On the markovian efficiency of bertrand and cournot equilibria. *Journal of Economic Theory*, 155:332 – 358.
- [Compte, 1998] Compte, O. (1998). Communication in repeated games with imperfect private monitoring. *Econometrica*, 66(3):597–626.
- [Crowe, 2017] Crowe, J. (2017). Reform, revolution and disruption in mental health care: a consumer’s perspective. *Public health research and practice*, 27(2).
- [European-Commission, 2014] European-Commission (2014). Antitrust: Commission fines producers of high voltage power cables €302 million for operating a cartel. *European Commission Press Release Database*.
- [European-Commission, 2015] European-Commission (2015). Antitrust: Commission fines cargo train operators €49 million for cartel. *European Commission Press Release Database*.

- [Fudenberg et al., 2014] Fudenberg, D., Ishii, Y., and Kominers, S. D. (2014). Delayed-response strategies in repeated games with observation lags. *Journal of Economic Theory*, 150:487–514.
- [Fudenberg and Levine, 2007] Fudenberg, D. and Levine, D. K. (2007). Continuous time limits of repeated games with imperfect public monitoring. *Review of Economic Dynamics*, 10(2):173–192.
- [Gerlach, 2009] Gerlach, H. (2009). Stochastic market sharing, partial communication and collusion. *International Journal of Industrial Organization*, 27(6):655 – 666.
- [Green and Porter, 1984] Green, E. J. and Porter, R. H. (1984). Noncooperative collusion under imperfect price information. *Econometrica*, 52(1):87–100.
- [Kandori and Matsushima, 1998] Kandori, M. and Matsushima, H. (1998). Private observation, communication and collusion. *Econometrica*, 66(3):627–652.
- [Keller et al., 2005] Keller, G., Rady, S., and Cripps, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68.
- [Kulkarni, 2009] Kulkarni, V. G. (2009). *Modeling and analysis of stochastic systems*. CRC Press.
- [Lagunoff and Matsui, 1997] Lagunoff, R. and Matsui, A. (1997). Asynchronous choice in repeated coordination games. *Econometrica*, 65(6):1467–1477.
- [Mailath and Samuelson, 2006] Mailath, G. J. and Samuelson, L. (2006). *Repeated games and reputations*, volume 2. Oxford university press Oxford.
- [Obara, 2009] Obara, I. (2009). Folk theorem with communication. *Journal of Economic Theory*, 144(1):120 – 134.
- [Rahman, 2014] Rahman, D. (2014). The power of communication. *The American Economic Review*, 104(11):3737–3751.
- [Rosenthal et al., 2013] Rosenthal, J. A., Lu, X., and Cram, P. (2013). Availability of consumer prices from us hospitals for a common surgical procedure. *JAMA internal medicine*, 173(6):427–432.

- [Sannikov and Skrzypacz, 2010] Sannikov, Y. and Skrzypacz, A. (2010). The role of information in repeated games with frequent actions. *Econometrica*, 78(3):847–882.
- [Stone, 1973] Stone, L. D. (1973). Necessary and sufficient conditions for optimal control of semi-markov jump processes. *SIAM Journal on Control*, 11(2):187–201.
- [Sugaya, 2015] Sugaya, T. (2015). Characterizing the limit set of perfect and public equilibrium payoffs with unequal discounting. *Theoretical Economics*, 10(3):691–717.
- [Takahashi and Wen, 2003] Takahashi, S. and Wen, Q. (2003). On asynchronously repeated games. *Economics Letters*, 79(2):239–245.
- [Wong et al., 2016] Wong, C., Faherty, L. J., Feingold, J., Town, R., and Werner, R. (2016). Price transparency for children and adolescent healthcare in the united states: Still opaque in 2015. *Journal of Adolescent Health*, 58(2):S48.