THE JOURNAL OF INDUSTRIAL ECONOMICS 0022-1821
Volume LXIX September 2021 No. 3

SUPPLY- VS. DEMAND-SIDE TRANSPARENCY: THE COLLUSIVE EFFECTS UNDER IMPERFECT PUBLIC MONITORING*

Luke Garrod[†]
Matthew Olczak[‡]

We analyse how market transparency affects collusion under imperfect monitoring where punishment phases occur on the equilibrium path. We show that increased transparency causes a 'pro-competitive' demand-side effect and an 'anti-competitive' supply-side effect on the optimal symmetric perfect public equilibrium (SPPE) profits. When transparency increases on both sides of the market, the optimal SPPE profits unambiguously increase at the perfect monitoring limit, because the pro-competitive demand-side effect vanishes. This result holds even when there is minimal structure on the competition game. The supply-side effect also dominates away from the limit under reasonable conditions. We draw conclusions for policy.

I. INTRODUCTION

There is some debate over when and how policymakers should intervene to affect the transparency of market information, such as firms' prices. The answer often depends upon whether the intervention will affect the supply side or the demand side. Greater demand-side transparency can make consumers more responsive to prices, which can make competition more effective or may undermine firms' incentives to sustain collusion.[1] Greater supply-side transparency can facilitate collusion by assisting coordination and/or by limiting firms' ability to implement secret price cuts.[2] Consequently, policies that aim to improve market outcomes by affecting transparency can cause unintended consequences, as was evident in the infamous Danish concrete case (see Albæk *et al*. [1997]). Thus, it is important to study the conditions under

*We are grateful for comments from Chris M. Wilson, two anonymous referees and the Editor. We have also benefitted from the comments of seminar participants at the Centre for Competition Policy (CCP), the Düsseldorf Institute for Competition Economics (DICE), and the European Association of Research in Industrial Economics (EARIE) conference 2019. The usual disclaimer applies.

[†]Authors' affiliations: School of Business and Economics, Loughborough University, Leicestershire, England.
e-mail: l.garrod@lboro.ac.uk
[‡]Aston Business School, Aston University, Birmingham, West Midlands, England.
e-mail: m.olczak@aston.ac.uk

[1] See Møllgaard and Overgaard [2001, 2002, 2008], OECD [2001] and Garrod *et al*. [2008].
[2] See Stigler [1964], Green and Porter [1984], Kühn and Vives [1995] and Kühn [2001].

which the anti-competitive effects of such policies are likely to dominate the pro-competitive effects.

The aim of this paper is to develop a theory of collusion in which a change in transparency causes *both* supply- and demand-side effects. Despite its importance for policy, this has not been adequately addressed by the previous literature which has, barring a couple of papers reviewed below, focused on either the supply or the demand side. Demand-side transparency is known to have an ambiguous effect on the sustainability of collusion when there is perfect monitoring of rivals' actions on the supply side (see, e.g., Møllgaard and Overgaard [2001, 2002, 2008]; Schultz [2005]). Imperfect monitoring on the supply side is known to undermine the profitability of collusion when transparency on the demand side is held constant (see, e.g., Green and Porter [1984]; Tirole [1988]; Harrington and Skrzypacz [2007]). However, this approach means there are some important policy questions currently unanswered. How does increased demand-side transparency affect collusion when price wars occur on the equilibrium path due to a lack of supply-side transparency? And under what conditions will any pro-competitive demand-side effects be dominated by anti-competitive supply-side effects?

To answer these questions, we analyse an infinitely repeated game where firms must imperfectly monitor each other's actions through a noisy public signal. Similar to Green and Porter [1984], punishment phases occur on the equilibrium path, and this allows us to capture the supply-side effects. To model the demand-side effects, we use the costly consumer search framework of Stahl [1989] as the underlying competition game. This follows Petrikaitė [2016] who analysed the demand-side effects under perfect monitoring and, under certain conditions, Stahl [1989] converges to Varian [1980] which has often appeared in the demand-side literature (see, e.g., Schultz [2005, 2017]; and Herre and Rasch [2013]). We initially focus on a general exogenous relationship between transparency on the supply and demand sides with a monitoring structure that encompasses a class of imperfect monitoring models, including Tirole [1988, p. 262-264] and Harrington and Skrzypacz [2007, p. 323-324]. After this, we endogenise this relationship using information available to the firms. While increased transparency has counteracting forces, we demonstrate that many of the effects can be signed unambiguously, with the main results also robust to other competition games.

Our primary contribution is to show the effects of increased transparency on the optimal symmetric perfect public equilibrium (SPPE) profits.[3] In our general analysis, increasing supply-side transparency strictly increases the

[3] In the main paper, we focus on a simple strategy profile, similar to Tirole [1988], where firms revert to the static Nash equilibrium for a number of periods when they receive a bad signal. However, in a supplemental appendix, we use the techniques of Abreu *et al*. [1986, 1990] to show that this simple strategy profile also generates the maximal SPPE profits. This analysis is closely related to Garrod and Olczak [2017], where we find a similar result in a Bertrand-Edgeworth oligopoly model.

optimal SPPE profits, other things equal. This is due to less frequent and shorter equilibrium punishment phases. In contrast, increasing demand-side transparency strictly decreases the optimal SPPE profits, other things equal. Despite counteracting effects, the dominant force here is that deviating is more attractive, which makes equilibrium punishment phases longer. Thus, when transparency increases on both sides of the market, the total effect consists of a 'pro-competitive' demand-side effect and an 'anti-competitive' supply-side effect. Both effects are smaller as monitoring becomes less imperfect, yet the pro-competitive demand-side effect vanishes at the perfect monitoring limit. Consequently, the anti-competitive supply-side effect dominates, so that increased market transparency unambiguously facilitates collusion. While this result is for the limiting case, it extends to a wider parameter space under reasonable conditions.

Following this, we extend Harrington and Skrzypacz [2007, p. 323-324] to endogenise monitoring. They analysed a setting where each firm's demand is subject to random shocks and firms monitor each other using information on their realised sales due to potential secret price cuts. However, in contrast to them, we assume that firms' prices are not always observable to all consumers. Thus, demand-side information about prices affects a deviant's sales, and this in turn affects monitoring. We consider two settings, the more complicated of which has firms' choosing the conditions that trigger a punishment phase, similar to Green and Porter [1984].[4] We show that in both examples the pro-competitive demand-side effect on the optimal SPPE profits is always dominated by the anti-competitive supply-side effect, even away from the perfect monitoring limit.

Our results have three main policy implications. First, interventions that increase demand-side transparency can successfully undermine collusion when the supply side is unaffected. This supports the view that: 'measures extending to consumers transparency which already exists among businesses should generally be pro-competitive' (OECD [2001, p. 9]). An example of such an intervention is the Australian Competition and Consumer Commission's (ACCC) investigation into the petrol market in 2014. The ACCC was concerned that retailers were using an online site to exchange information on prices, and their solution was to make the same information available to consumers.[5] According to our results, such interventions should be pro-competitive, but they would be less successful in markets where the supply-side is more transparent.

Second, interventions to promote demand-side transparency can cause anti-competitive effects when they also improve monitoring. This is consistent

[4] An interesting feature of this example is that it is optimal for firms to ensure that the chance of entering the punishment phase is at the lowest possible level.

[5] ACCC v Informed Sources, Federal Court of Australia, 2014. See *https://www.accc.gov.au/media-release/petrol-price-information-sharing-proceedings-resolved*

with the intervention by the Danish Competition Council (DCC) into the ready-mixed concrete industry in 1993 (see Albæk *et al*. [1997]). The DCC intended to enhance competition by publishing past transaction prices to make consumers more responsive to prices; instead, it facilitated collusion by limiting producers' ability to implement secret price cuts. Our general results imply that, conditional on firms' being able to coordinate their behaviour, such unintended consequences are more likely in markets where monitoring is almost perfect, because the pro-competitive demand-side effects will be small.

Third, the prohibition of facilitating practices that enhance supply-side transparency can undermine collusion, even when it has a negative impact on demand-side transparency. This is consistent with the approach taken by the U.K.'s Competition and Markets Authority (CMA) following a market investigation into the cement market in 2016.[6] The CMA concluded that coordination between the major firms had raised prices, and it implemented remedies designed to limit supply-side transparency, including prohibiting price announcements and restricting the publication of sales data. The CMA recognised that these practices might have previously had some pro-competitive demand-side effects but argued that their overriding effect was to facilitate collusion. Our general results suggest that such an intervention will be pro-competitive in markets where monitoring is close to perfect, because any demand-side effect will be small.

Finally, as mentioned above, a few papers have attempted to analyse both the supply- and demand-side effects in the same framework. In particular, in an extension of their demand-side analysis under perfect monitoring, Møllgaard and Overgaard [2001, 2008] also model supply-side transparency by allowing for imperfect monitoring off the equilibrium path (i.e., only deviations are detected imperfectly). However, this approach rules out equilibrium price wars that would occur if there were also imperfect monitoring on the equilibrium path. Furthermore, following a similar approach but in a more structured oligopoly model, Schultz [2017] argues that any collusive equilibrium with price wars will be less profitable and hence inferior to the collusive equilibria without price wars in his main analysis. However, he does not characterise any collusive equilibria with punishments on the equilibrium path, and his result relies on the assumption that there are histories in which firms' actions are observed by their rivals. In contrast, we analyse the optimal SPPE profits in a standard imperfect public monitoring framework, where each player's actions always remain unobservable to their rivals. Thus, the equilibria in the main analysis of Schultz [2017] do not exist in our standard imperfect monitoring framework.

[6] See 'Aggregates, Cement and Ready-Mix Concrete Market Investigation,' Competition and Markets Authority decision of 13th April, 2016.

The paper is organised as follows. In Section II, we present the assumptions of the model and then solve the game in Section III. In Section IV, we analyse the effects of market transparency on the optimal profits given a general exogenous relationship between supply- and demand-side transparency. We isolate the supply- and demand-side effects, before comparing their aggregate effect. In Section V, we analyse the effects when there is an endogenous relationship between supply- and demand-side transparency. In Section VI, we show that our main results are robust to the competition game employed. Section VII concludes. All proofs are relegated to the Appendix. Finally, in another supplemental appendix, we extend our analysis in two ways. First, we use the techniques of Abreu *et al.* [1986, 1990] to show that the simple strategy profile used in Section III generates the maximal equilibrium profits. Second, we analyse the effects of transparency on the critical discount factor under trigger strategies.[7]

II. THE MODEL

II(i). *Basic Assumptions*

Consider a market where there are two firms, denoted 1 and 2, that compete in prices to sell a homogeneous good in an infinitely repeated game. In period t, p_{it} is the price of firm $i \in \{1, 2\}$, p_{jt} is the price of its rival, $j \neq i$, and $\mathbf{p}_t \equiv \{p_{1t}, p_{2t}\}$ is the vector of both prices. There are $b \geq 1$ discrete buyers, each of whom will purchase $\frac{m}{b}$ units in each period if the price does not exceed their valuation of the product, ν. Thus, the expected market demand is m for any price less than or equal to ν. Firms' costs are normalised to zero and they have a common discount factor $\delta \in (0,1)$.

II(ii). *Information and Monitoring*

Regarding information on the demand side, we assume that in each period t a random process determines whether each buyer is a 'shopper' or a 'non-shopper'. Shoppers are fully informed of prices. They will purchase from the lowest-priced firm or select at random with iid decisions when prices are equal. Non-shoppers initially do not know any prices. However, they observe one firm's price freely and incur a search cost $s > 0$ to observe the second firm's price. They will purchase from the lowest-priced firm they know and select the firm they search first at random, again with iid decisions. We postpone assumptions relating to the random processes that determine the buyers' types until Section V. At this stage, we need only assume that firms expect when they set prices that market demand will consist of a fraction $\lambda \in (0,1)$ of

[7] The supplemental appendix may be found in the online version of this article at *http://wileyonlinelibrary.com/journal/joie*

shoppers and $1-\lambda$ of non-shoppers. Given more consumers are expected to observe both prices costlessly when λ is closer to 1, we say there is increased demand-side transparency when λ is higher.

On the supply side, we assume that firm i never observes rival j's price $p_{j\tau}$ for all $\tau \in \{0, \ldots, t\}$. This implies that firms must monitor a collusive scheme through some means other than prices. At this stage, we assume firms observe a noisy binary public signal of $\mathbf{p}_t$, denoted $y_t \in \left\{\underline{y}, \overline{y}\right\}$, where:

$$\Pr\left(y_t = \underline{y} | p_{1t}, p_{2t}\right) = 1 - \Pr\left(y_t = \overline{y} | p_{1t}, p_{2t}\right) = \begin{cases} \alpha \in (0, 1] & \text{if } p_{1t} = p_{2t}, \\ \beta \in (0, 1] & \text{otherwise.} \end{cases}$$

Without loss of generality, we henceforth refer to $\overline{y}$ and $\underline{y}$ as a positive and negative signal, respectively. Consequently, the above says α is the probability of a negative signal conditional on firms setting the same price, and β is the probability of a negative signal conditional on one firm undercutting. Thus, for any symmetric strategy profile, α represents the (conditional) probability of a 'false negative' and β represents the (conditional) probability of a 'true negative.' Given monitoring is closer to perfect when α is closer to 0 and when β is closer to 1, we say that there is increased supply-side transparency when α is lower and/or β is higher.

Having explained the information on each side of the market, we now discuss how they will interact. There are a number of ways in which the relationship between supply- and demand-side transparency could be modelled endogenously. For example, our monitoring structure is general enough to encompass as special cases both Tirole [1988, p. 262-264] and Harrington and Skrzypacz [2007, p. 323-324], so we could follow either of these approaches to place more structure on α and β. Nevertheless, to present the results as generally as possible, our initial approach in Section IV is to analyse the effects for a general exogenous relationship. This amounts to assuming there is a parameter that increases transparency on the supply side (via α and β) and on the demand side (via λ) at the same time. The results of this approach will then apply to any model of imperfect monitoring that has an endogenous relationship with the same properties. Following this, in Section V we construct such a model that endogenises the probabilities α and β in terms of λ by drawing on Harrington and Skrzypacz [2007]. This extends their work because it relaxes the assumption that firms' prices are *always* observable to consumers, despite *never* being observable to firms.[8]

In providing more detail of how we will later endogenise these probabilities, let us first briefly summarise the model of Harrington and Skrzypacz [2007, p. 323-324]. They analyse a setting where for a known market demand, m, all buyers are fully informed of prices (i.e., $\lambda = 1$) and they randomly choose between the firms with iid decisions when prices are equal. This implies that

[8] We show how our monitoring structure relates to Tirole [1988] in the supplemental appendix.

each firm's *realised* sales may not be evenly distributed when they have set the same price. Thus, colluding firms will face a non-trivial signal extraction problem when monitoring the scheme through their sales. Specifically, when all buyers purchase from one firm, its rival will be unsure whether its zero sales has resulted from chance or a deviation by its rival. Consequently, α and β are respectively given by the probabilities that all buyers purchase from one firm when they set the same price (a false negative) and when one undercuts (a true negative). In Section V, we follow this approach but with $\lambda < 1$. In this case, a higher fraction of shoppers who will observe a deviant's lower price, λ, will increase the probability that the non-deviant will receive low sales, and this in turn will increase the probability of a true negative, β. Hence, raising demand-side transparency will endogenously raise supply-side transparency. The specific nature of this endogenous relationship will depend upon the random processes that determine the buyers' types, and we return to this in Section V.

Before moving on, we should take a moment to defend our assumptions. While the monitoring structure allows us to encompass other papers in the literature, a limitation is that it forces us to restrict attention to symmetric strategies. The reason is that firms will only receive a positive signal if they set the same price. This is potentially problematic because asymmetric strategies may allow firms to monitor each other better than symmetric play, even in markets that have symmetric environments like ours, and this may allow them to earn higher profits for some discount factors.[9] However, symmetric strategies seem most appropriate for our application in our opinion. The reason is that asymmetric strategies often involve mechanisms, such as adjustments in market shares or side payments, that would be very difficult to implement or coordinate on without illegal express communication. In contrast, policies relating to market transparency are usually targeted at tacitly colluding firms, because there are more effective policy tools for dealing with illegal cartels that involve explicit communication. Thus, our analysis primarily applies to markets in which tacitly colluding firms are unable to coordinate on complex asymmetric strategies.

II(iii). *Static Nash Equilibrium Profits*

The competition game is equivalent to the unit-demand version of Stahl [1989] as first analysed by Janssen *et al*. [2005]. Thus, the statíc Nash equilibrium is

[9] For example, in Harrington and Skrzypacz [2007], the signals firms observe about each other's actions carry relevant information about which firm may have deviated. This enables the cartel to adjust their profits so that firms who sell too much compensate those firms who sell too little. Furthermore, Amelio and Biancini [2010] show, in a setting similar to Tirole [1988], that colluding firms can sometimes do better than symmetric play by sharing the market as monopolists in alternating periods.

in mixed strategies, where firms choose their prices from an interval $\left[\underline{p}, p^N\right]$ according to a distribution function $G(p)$ with density $g(p)$. An important part of the analysis is the upper bound $p^N \equiv \min\{v, p^*\}$, where p^* is the price that equates a non-shopper's expected marginal benefit and marginal cost of searching to discover a second firm's price,

$$\int_{\underline{p}}^{p^*} (p^* - p)\, g(p) dp = s. \tag{1}$$

Lemma 1 states the characteristics of the static Nash equilibrium and we sketch the proof below.[10]

Lemma 1. For any $\lambda \in (0,1)$ and $s > 0$, there exists a symmetric mixed strategy Nash equilibrium where each firm's profits are $\pi^N = p^N \frac{m}{2}(1-\lambda) \in \left(0, v\frac{m}{2}\right)$ where $p^N = \min\{v, p^*\}$. The equilibrium pricing distribution is $G(p) = 1 - \frac{(p^N - p)(1-\lambda)}{2\lambda p}$ on $\left[\underline{p}, p^N\right]$ where $\underline{p} = p^N\left(\frac{1-\lambda}{1+\lambda}\right) < p^N$ and $p^* = \frac{s}{1+\left(\frac{1-\lambda}{2\lambda}\right)\ln\left(\frac{1-\lambda}{1+\lambda}\right)}$.

Intuitively, when firm i sets $p \in \left[\underline{p}, p^N\right]$ it will receive expected profits of:

$$\pi_i = p\left(\lambda m(1 - G(p)) + \frac{m(1-\lambda)}{2}\right).$$

This says that firm i expects to supply all shoppers if it is the lowest-priced firm, which occurs with a probability $1-G(p)$, and also expects to supply an equal share of the non-shoppers, even when it is the highest-priced firm. Non-shoppers are always divided equally because no firm charges more than p^* in equilibrium. Thus, the marginal benefit of searching to discover a second firm's price is always less than the marginal cost, so non-shoppers never undertake costly search in equilibrium. Consequently, by charging $p_i = p^N$, firm i can obtain profits of π^N with certainty, because the firm will only supply its share of the non-shoppers, $\frac{m(1-\lambda)}{2}$. Furthermore, the equilibrium pricing distribution is derived from setting $\pi_i = \pi^N$, and this shows that firm i will never set a price below $\underline{p} = p^N\left(\frac{1-\lambda}{1+\lambda}\right)$ in an attempt to be the lowest-priced firm. The equilibrium pricing distribution can be used with (1) to find the upper bound p^N.

Let us end this section by discussing how the competition game relates to Schultz [2005] and Petrikaitė˙ [2016]. With this aim, first note that for any $s < v$ there exists a unique level of λ, denoted $\underline{\lambda} \in (0, 1)$, where the upper

[10] For more details, see Janssen *et al.* [2005].

bound price is $p^N = p^* < v$ if and only if $\lambda > \underline{\lambda}$.[11] Thus, if $\lambda > \underline{\lambda}$, then $p^N = p^* < v$ and the competition game is the same as Petrikaitė [2016] where π^N is an increasing function of the search cost, s. Yet, if $\lambda \leq \underline{\lambda}$, then $p^N = \nu$ and the profits are equivalent to the homogeneous goods analysis of Schultz [2005] where π^N is independent of s. Consequently, our analysis extends both of theirs to the case of imperfect monitoring, where $\alpha > 0$ and $\beta \leq 1$.

III. OPTIMAL SYMMETRIC EQUILIBRIUM PROFITS

We now solve the repeated game restricting attention to sequential equilibria in public strategies where firms condition their play on the public history (see Fudenberg and Tirole [1994, p. 187-191]). Such equilibria are known as perfect public equilibria (PPE). In the main text, we restrict attention to a particular class of PPE in which, similar to Green and Porter [1984] and Tirole [1988], firms punish each other by reverting to the static Nash equilibrium for a fixed number of periods upon observing a negative signal. The strategy profile for this approach is formally described below and we refer to it as trigger strategies. In the supplemental appendix, we use the techniques of Abreu *et al.* [1986, 1990] to solve for the set of (strongly) symmetric PPE.[12] This shows that trigger strategies are optimal equilibrium strategies in the sense that they support the maximal symmetric PPE (SPPE) payoffs.

Trigger strategies are formally defined as follows. There are 'collusive periods' and 'punishment phases'. In any collusive period $t \geq 0$, which includes period 0, firm i should set $p_{it} = \nu \forall i$. Then, if $y_t = \overline{y}$, such that the public signal is positive, period t+1 is a collusive period. Alternatively, if $y_t = \underline{y}$, such that the public signal is negative, then firms enter a punishment phase in period t+1. In a punishment phase, firm i should play the static Nash equilibrium for T periods, and then period t+T+1 is a collusive period. The sequence repeats.

To begin the analysis, let us describe the firms' per-period profits under three scenarios. First, in each period of a punishment phase, a firm expects to receive $\pi^N = p^N \frac{m(1-\lambda)}{2}$, from Lemma 1. Second, in a collusive period, each firm's expected profits are $\pi^c = v\frac{m}{2}$, because the shoppers and non-shoppers randomly select between the firms. This is due to shoppers being indifferent when prices are the same and non-shoppers only searching one firm, because the marginal benefit of searching the second firm is zero but the marginal cost is $s > 0$. Third, if a firm undercuts its rival in a collusive period it will not attract any additional non-shoppers, because its deviation will not be observed by them as they do not search a second firm in a symmetric

[11] This follows from $\lim_{\lambda \to 0} p^N = \nu$, $\lim_{\lambda \to 1} p^N = \min\{s, \nu\}$, and $\frac{\partial p^*}{\partial \lambda} < 0$.

[12] Such strategies are strongly symmetric in the sense that firms are required by the strategy profile to play the same actions as one another in future periods, even if they have not played the same actions in the past.

equilibrium. Thus, firm i's optimal deviation in a collusive period is to undercut ν marginally to supply the shoppers that would have otherwise purchased from its rival j. This amounts to an extra $\lambda\frac{m}{2}$ units, so the resultant expected deviation profits are $\pi^d = \nu\frac{m(1+\lambda)}{2}$.

Turning attention to the firms' incentives, let V^c and V^p denote each firm's expected (normalised) profits in a collusive period and at the start of a punishment phase, respectively, such that:

(2)
$$\begin{aligned} V^c &= (1-\delta)\,\pi^c + \delta\left[\alpha V^p + (1-\alpha)\,V^c\right] \\ V^p &= (1-\delta)\sum_{t=0}^{T-1}\delta^t\pi^N + \delta^T V^c \end{aligned}$$

Solving simultaneously yields:

(3)
$$V^c(\alpha,\lambda,s,T) = \pi^N + \frac{(1-\delta)}{1-\delta+\alpha\delta\left(1-\delta^T\right)}\left(\pi^c-\pi^N\right)$$

$$V^p(\alpha,\lambda,s,T) = \pi^N + \frac{(1-\delta)\,\delta^T}{1-\delta+\alpha\delta\left(1-\delta^T\right)}\left(\pi^c-\pi^N\right),$$

where $\pi^c > V^c(\alpha, \lambda, s, T) > V^p(\alpha, \lambda, s, T)$ for all $T > 0$ and $V^p(\alpha, \lambda, s, T) > \pi^N$ for all $T < \infty$.

The profile of trigger strategies is a PPE if, for each period t and any public history $h^t = (y_0, y_1, \ldots, y_{t-1})$, the strategies yield a Nash equilibrium from period t onwards. By definition each period of a punishment phase is a Nash equilibrium, so consider deviations in a collusive period, where each firm's incentive compatibility constraint (ICC) is:

(4)
$$(\beta-\alpha)\frac{\delta}{1-\delta}\left(V^c(\alpha,\lambda,s,T) - V^P(\alpha,\lambda,s,T)\right) \geq \nu\frac{m\lambda}{2}$$

This says that a firm will not deviate in any collusive period if the short-term gain from deviating, $\pi^d - \pi^c = \nu\frac{m\lambda}{2} > 0$, does not exceed the long-term loss from the punishment. The long-term loss is determined by the change in the probability of entering the punishment phase following a deviation $(\beta-\alpha)$ multiplied by the severity of the punishment when imposed, $\frac{\delta}{1-\delta}\left(V^c(.) - V^p(.)\right)$. Substituting $V^p(\alpha, \lambda, s, T)$ and $V^c(\alpha, \lambda, s, T)$ into (4), then rearranging yields:

$$\beta\left(\pi^c-\pi^N\right) - \alpha\left(\pi^d-\pi^N\right) - \frac{(1-\delta)\left(\pi^d-\pi^c\right)}{\delta} \geq \delta^T\left[\beta\left(\pi^c-\pi^N\right) - \alpha\left(\pi^d-\pi^N\right)\right]$$

(5)

Let α^* denote the critical probability of a false negative that sets the right hand-side of (5) to zero, where $\alpha^* \equiv \beta\left(\frac{\pi^c-\pi^N}{\pi^d-\pi^N}\right) \in (0,\beta)$. Furthermore, let δ^* denote the critical discount factor that sets the left hand-side to zero, where $\delta^* \equiv \frac{1}{(1-\alpha)+(\beta-\alpha)\left(\frac{\pi^c-\pi^N}{\pi^d-\pi^c}\right)}$, so that as $T\to\infty$, the ICC (5) holds for all $\delta \geq \delta^*$.

Proposition 1. For any $s > 0$, $0 < \lambda \leq 1$ and $0 < \beta \leq 1$, if $0 < \alpha < \alpha^*$ and $\delta \geq \delta^*$, then the optimal PPE profits under trigger strategies are:

$$V^* = v\frac{m}{2}\left(1-\frac{\alpha\lambda}{\beta-\alpha}\right) \in \left(\pi^N, \pi^c\right). \tag{6}$$

Proposition 1 shows that, despite firms earning $\pi^c = v\frac{m}{2}$ in every collusive period, the optimal PPE profits under trigger strategies are below this level, $V^* < \pi^c$ for all $\alpha > 0$. The reason is that punishment phases occur in equilibrium, because firms can only imperfectly monitor each other's actions. The optimal punishment phase duration that generates the optimal profits, denoted $T^*(\alpha,\beta,\lambda,s)$, ensures that the ICC (5) is binding with no slack.[13] Furthermore, note that while the optimal profits are a function of transparency on the supply side (via α and β) and the demand side (via λ), they are independent of the search costs, s. We delay discussion of this until Section VI and focus on the other parameters until then.

While in the main text our focus is on trigger strategies, in the supplemental appendix we use the techniques of Abreu *et al.* [1986, 1990] to show that trigger strategies generate the maximal SPPE profits for any $\delta \geq \delta^*$. This implies that there is no (strongly) symmetric strategy profile that generate higher profits than trigger strategies. Thus, henceforth we refer to such profits as the optimal SPPE profits. Nevertheless, as we show in the supplemental appendix, there may be alternative strategy profiles that support the maximal SPPE profits at some $\delta < \delta^*$. This occurs when the minimal SPPE profits are lower than π^N, so firms are able to implement a harsher punishment than under trigger strategies. This is not important for our objectives though, as we are primarily interested in the effects on the optimal SPPE profits, V^*, due to the presence of equilibrium punishment phases. Furthermore, focussing on trigger strategies allows us to describe the intuitions of the effects of transparency in the most intuitive strategy profile for our application.

IV. SUPPLY- VS. DEMAND-SIDE TRANSPARENCY

We now investigate how an increase in transparency affects the optimal SPPE profits, V^*. We proceed by first analysing the supply-side effects in isolation,

[13] Although $T^*(\alpha,\beta,\lambda,s)$ may not be an integer, the expected punishment phase duration could still equal this if firms can vary the length of punishment phases using a publicly observable randomisation device.

to show that our model has the usual 'anti-competitive' effects from improved monitoring.[14] Then we analyse the demand-side effects in isolation and show that these are 'pro-competitive'. Finally, we discuss the effects of *market* transparency when both sides of the market are affected at the same time. Throughout this section we impose $0 < \alpha < \alpha^* < \beta \leq 1$ and $\delta \geq \delta^*$ so V^* is as given in (6). The effects can often be represented in terms of how elastic a parameter is to transparency, so it is helpful to denote $\varepsilon_{g,h} \equiv \frac{\partial g}{\partial h}\frac{h}{g}$ as the elasticity of some parameter g with respect to another h.

Proposition 2. Increasing supply-side transparency strictly increases the optimal SPPE profits, $\frac{\partial V^*}{\partial \alpha} < 0$ and $\frac{\partial V^*}{\partial \beta} > 0$. The magnitude of the effects are smaller when the supply side is more transparent, $\frac{\partial^2 V^*}{\partial \alpha^2} < 0$, $\frac{\partial^2 V^*}{\partial \beta^2} < 0$ and $\frac{\partial^2 V^*}{\partial \alpha \partial \beta} > 0$.

Intuitively, increasing the probability of a true negative, β, raises the likelihood that a deviation will be detected. This implies that there is a larger change in the probability of punishment following a deviation, $\frac{\partial(\beta-\alpha)}{\partial \beta} > 0$, which introduces slack into the ICC. Consequently, punishment phases are shorter than before, $\frac{\partial T^*}{\partial \beta} < 0$. This indirectly increases the optimal SPPE profits, $\frac{\partial V^c}{\partial T}\frac{\partial T^*}{\partial \beta} > 0$, because there are expected to be more collusive periods in equilibrium. Similarly, decreasing the probability of a false negative, α, reduces the likelihood that firms will enter a punishment phase on-the-equilibrium path. This has two effects on the optimal SPPE profits. First, it directly increases the profits because, other things equal, there are expected to be fewer punishment phases in equilibrium, $\frac{\partial V^c}{\partial \alpha} < 0$. Second, it also introduces slack into the ICC, which results from i) a greater change in the probability of punishment following a deviation, $\frac{\partial(\beta-\alpha)}{\partial \alpha} < 0$, and ii) a rise in the severity of the punishment when imposed, $\frac{\partial(V^c - V^p)}{\partial \alpha} < 0$. Thus, punishment phases are shorter than before, $\frac{\partial T^*}{\partial \alpha} > 0$, and this indirectly increases the profits further, $\frac{\partial V^c}{\partial T}\frac{\partial T^*}{\partial \alpha} < 0$.

Proposition 3. Increasing demand-side transparency strictly decreases the optimal SPPE profits, $\frac{\partial V^*}{\partial \lambda} < 0$. The magnitude of the effect is smaller when the supply side is more transparent, $\frac{\partial^2 V^*}{\partial \lambda \partial \alpha} < 0$ and $\frac{\partial^2 V^*}{\partial \lambda \partial \beta} > 0$.

An increase in demand-side transparency reduces the static Nash equilibrium profits and raises the deviation profits, $\frac{\partial \pi^N}{\partial \lambda} < 0$ and $\frac{\partial \pi^d}{\partial \lambda} > 0$. This has three separate effects on the optimal SPPE profits. First, future equilibrium

[14] While we say 'usual', we note that a recent paper by Sugaya and Wolitzky [2018] shows that better information about rivals' prices can hinder collusion when firms agree to supply only their home market. The reason is that a deviant is better able to tailor its deviation price to the specific conditions of their rivals' markets when it has better information on its rival's prices.

punishment phases are expected to yield less profit than before, other things equal, so the profits *decrease*, $\frac{\partial V^c}{\partial \pi^N}\frac{\partial \pi^N}{\partial \lambda} < 0$. Second, there is an increase in the serevity of the punishment when imposed, $\frac{\partial (V^c - V^p)}{\partial \pi^N}\frac{\partial \pi^N}{\partial \lambda} > 0$, which introduces slack into the ICC, other things equal. This reduces the duration of the punishment phases, $\frac{\partial T^*}{\partial \pi^N} > 0$, so the profits *increase*, $\frac{\partial V^c}{\partial T}\frac{\partial T^*}{\partial \pi^N}\frac{\partial \pi^N}{\partial \lambda} > 0$. Third, there is a greater short-term gain from deviating, $\frac{\partial(\pi^d - \pi^c)}{\partial \lambda} > 0$, and this tightens the ICC, other things equal. As a result, the duration of the punishment phases lengthens, $\frac{\partial T^*}{\partial \pi^d} > 0$, so the profits *decrease*, $\frac{\partial V^c}{\partial T}\frac{\partial T^*}{\partial \pi^d}\frac{\partial \pi^d}{\partial \lambda} < 0$. Despite these counteracting forces, it is possible to sign the total effect unambiguously because the first and second effects perfectly offset each other. Consequently, the total effect of increasing demand-side transparency on the optimal SPPE is strictly negative and amounts to the indirect effect caused by a greater short-term gain from deviating.

We are now in a position to analyse the case of market transparency, where the demand and supply sides are affected at the same time. Formally, we index market transparency by $\theta \in (0, \bar{\theta})$, where $\lim_{\theta \to 0} \alpha = \alpha^*$ and $\lim_{\theta \to \bar{\theta}} \alpha = 0$. Thus, when $\theta = 0$, the market is so untransparent that collusion under imperfect monitoring is not sustainable for any $\delta < 1$. In contrast, when $\theta = \bar{\theta}$, the market is so transparent that firms can perfectly monitor each other on-the-equilibrium path so that punishment phases do not occur in equilibrium. We refer to $\theta \to \bar{\theta}$ as the perfect monitoring limit. Furthermore, we also assume:

Assumption 1a. $\frac{\partial \alpha}{\partial \theta} \le 0, \frac{\partial \beta}{\partial \theta} \ge 0, \frac{\partial \lambda}{\partial \theta} > 0, \forall\, \theta \in (0, \bar{\theta})$

Assumption 1b. If $\frac{\partial \alpha}{\partial \theta} = 0$, then $\frac{\partial \beta}{\partial \theta} > 0$, and if $\frac{\partial \beta}{\partial \theta} = 0$, then $\frac{\partial \alpha}{\partial \theta} < 0$

Assumption 1a implies that increasing market transparency, θ, raises the level of transparency on the demand side (by strictly increasing λ) and can also affect the supply side. Together Assumptions 1a and 1b ensure that there is a strict effect on the supply side that can work through improved monitoring on-the-equilibrium path only (via α) or off-the-equilibrium path only (via β), or both.

Given these assumptions, it is then clear from Proposition 2 and 3 that increasing market transparency can have counteracting pro- and anti-competitive effects. We now consider which effect dominates at the perfect monitoring limit.

Proposition 4. Increasing market transparency at the perfect monitoring limit increases the optimal SPPE profits, $\lim_{\theta \to \bar{\theta}} \frac{dV^*}{d\theta} \ge 0$, where the sign is strict if $\lim_{\theta \to \bar{\theta}} \frac{\partial \alpha}{\partial \theta} < 0$.

Proposition 4 shows that the supply-side effects (weakly) dominate the demand-side at the perfect monitoring limit. The intuition is that the effect of demand-side transparency is infinitesimally small at the limit, $\lim_{\theta\to\bar{\theta}}\frac{\partial V^*}{\partial\lambda}=0$, so increasing market transparency can only cause anti-competitive supply-side effects. In fact, the supply-side effect that works through improved monitoring off-the-equilibrium path (via β) is also infinitesimally small, $\lim_{\theta\to\bar{\theta}}\frac{\partial V^*}{\partial\beta}=0$. Thus, the result is only driven by the supply-side effect that works through improved monitoring on-the-equilibrium path (via α), $\lim_{\theta\to\bar{\theta}}\frac{\partial V^*}{\partial\alpha}<0$. The reason is that at the perfect monitoring limit the collusive profits in (3) are independent of the duration of the punishment phase, $\lim_{\theta\to\bar{\theta}}\frac{\partial V^c}{\partial T}=0$, because the probability of a false negative approaches zero, $\lim_{\theta\to\bar{\theta}}\alpha=0$. Consequently, all indirect effects are infinitesimally small. However, the supply-side effect that works through improved monitoring on-the-equilibrium path (via α) is comprised of a direct and an indirect effect, and the direct effect does not vanish if $\lim_{\theta\to\bar{\theta}}\frac{\partial\alpha}{\partial\theta}<0$. Therefore, this direct effect ensures that the anti-competitive supply-side effects dominate at the perfect monitoring limit.

The supply-side effects can dominate away from the perfect monitoring limit under reasonable conditions, as we demonstrate in Section V. As a precursor to that analysis, note that a sufficient condition for the supply-side effects to always dominate is V^* being concave in market transparency, $\frac{d^2V^*}{d\theta^2}\leq 0$. The reason is that an increase in market transparency has an even greater positive effect on V^* as θ moves away from the limit $\bar{\theta}$, such that $\frac{dV^*}{d\theta}\geq 0$ for all $\theta\in(0,\bar{\theta})$. Furthermore, in the opposite case where V^* is strictly convex in market transparency, $\frac{d^2V^*}{d\theta^2}>0$, the supply-side effects will still dominate when θ is sufficiently close to $\bar{\theta}$. More specifically, there will be a threshold $\underline{\theta}<\bar{\theta}$ for which $\frac{dV^*}{d\theta}>0$ for all $\theta\in(\underline{\theta},\bar{\theta})$. This lower threshold will be given by $\underline{\theta}=\max\{0,\theta^*\}$ where θ^* uniquely defines the point at which V^* is at a minimum, $\left.\frac{dV^*}{d\theta}\right|_{\theta=\theta^*}=0$.

V. ENDOGENOUS MONITORING

Having analysed the model with a general exogenous relationship between supply and demand-side transparency, we now endogenise α and β so they are determined by information available to firms. Following Harrington and Skrzypacz [2007, p. 323-324], we assume that firms monitor each other through their sales and model the way in which sales are subject to idiosyncratic shocks over time. To do so, we must explicitly model the random processes that determines whether buyers are shoppers or non-shoppers.

We analyse two examples. The first is the simplest possible example where the random processes that determine the buyers types are perfectly correlated across buyers. Thus, in each period, all buyers will be either shoppers or non-shoppers. This example is the simplest because, similar to Harrington and Skrzypacz [2007, p. 323-324], it ensures that firms will only ever receive a negative signal when one firm makes zero sales. The second is more complicated in that there are independently and identically distributed random processes that determine each buyer's type, so types can vary across buyers in the same period. This allows firms to choose the level of sales that constitute a negative signal, similar to Green and Porter [1984].

In both examples there is no supply-side effect that works through improved monitoring on-the-equilibrium path (via α), because $\frac{\partial \alpha}{\partial \lambda} = 0$ such that $\frac{\partial \alpha}{\partial \theta} = \frac{\partial \lambda}{\partial \theta}\frac{\partial \alpha}{\partial \lambda} = 0$. Consequently, there are only two effects of increased market transparency on V^*: the pro-competitive demand-side effect, via $\frac{\partial \lambda}{\partial \theta} > 0$, and the anti-competitive supply-side effect that improves monitoring off-the-equilibrium path, via $\frac{\partial \beta}{\partial \theta} = \frac{\partial \lambda}{\partial \theta}\frac{\partial \beta}{\partial \lambda} > 0$. Thus, from Proposition 4, these two effects are both infinitesimally small at the perfect monitoring limit, such that $\lim_{\theta \to \overline{\theta}} \frac{dV^*}{d\theta} = 0$. In what follows we show that the supply-side effects continue to (weakly) dominate the pro-competitive demand-side effects away from the limit, despite the fact that we have removed one of the avenues in which profits can be impacted by the anti-competitive supply-side effects.

V(i). *Example 1: Perfectly Correlated Random Processes*

Suppose that there is a random process that determines whether each buyer is a shopper or non-shopper, where draws are perfectly correlated across buyers. Thus, in any given period, λ represents the probability that all buyers are shoppers and $1-\lambda$ is the probability all buyers are non-shoppers. Clearly, this ensures that the expected fractions of shoppers and non-shoppers are λ and $(1-\lambda)$, respectively, as previously assumed.

If $p_1 \neq p_2$, then all buyers will purchase from the lowest-priced firm when they are shoppers, but they will each select at random the first and only firm they will search when they are non-shoppers. If $p_1 = p_2$, then each buyer will choose between the firms randomly, regardless of whether the buyers are shoppers or non-shoppers. This implies that firm i's realised sales will be subject to idiosyncratic shocks on-the equilibrium path. Furthermore, if all buyers purchase from one firm, then the other firm that makes zero sales will be uncertain as to whether this resulted from chance or a deviation. Thus, firms receive a negative public signal, $y_t = \underline{y}$, when all buyers purchase from one firm; otherwise, the public signal is positive, $y_t = \overline{y}$, because it is clear that the sales distribution is determined by chance when both firms make positive sales.

Consistent with Harrington and Skrzypacz [2007, p. 323-324], the probability of a false negative is $\alpha = 2\left(\frac{1}{2}\right)^b$. The reason is that, given all buyers choose randomly when $p_1 = p_2$, the probability that all b buyers purchase from firm i is $\left(\frac{1}{2}\right)^{b}$ and this event could occur for both firms. In contrast, the probability of a true negative is $\beta = \lambda + (1 - \lambda)\,2\left(\frac{1}{2}\right)^{b}$. The reason is that if $p_1 \neq p_2$, then all buyers will purchase from the lower-priced deviant when buyers are shoppers, but all buyers will randomly purchase from the same firm with a probability $2\left(\frac{1}{2}\right)^{b}$ when they are non-shoppers. Thus, increasing market transparency makes the demand side more transparent (via $\frac{\partial \lambda}{\partial \theta} > 0$) and this in turn endogenously raises supply-side transparency (via $\frac{\partial \beta}{\partial \theta} = \frac{\partial \lambda}{\partial \theta}\frac{\partial \beta}{\partial \lambda} > 0$).

Proposition 5. For any $b \in (2, \infty)$ and $s > 0$, if λ is sufficiently close to 1, such that $\alpha < \alpha^{*}$, and if $\delta \geq \delta^{*}$, then increasing market transparency has no effect on the optimal SPPE profits, $\frac{dV^{*}}{d\theta} = 0$.

Proposition 5 shows that the optimal SPPE profits, V^{*}, are independent of market transparency. The intuition is that the pro-competitive demand-side effect (via $\frac{\partial \lambda}{\partial \theta}$) and the anti-competitive supply-side effect that improves monitoring off-the-equilibrium path (via $\frac{\partial \beta}{\partial \theta}$) perfectly offset each other. Thus, given $\frac{\partial \alpha}{\partial \theta} = 0$, it follows that the result of Proposition 4 extends beyond the perfect monitoring limit in this setting.

V(ii). *Example 2: iid Random Processes*

Suppose the random processes that determine buyers' types have iid draws, so λ now represents the probability that an individual buyer is a shopper and $1-\lambda$ is the probability the buyer is a non-shopper. As before, the expected fractions of shoppers and non-shoppers are $\frac{1}{b}\sum_{k=0}^{b} k\binom{b}{k}(\lambda)^{k}(1-\lambda)^{b-k} = \lambda$ and $1-\lambda$, respectively.

In this setting, if $p_1 \neq p_2$, shoppers will purchase from the lowest-priced firm and non-shoppers will select at random the first and only firm they search. Thus, $\frac{1-\lambda}{2}$ is the probability that an individual buyer will unfortunately purchase from the highest-priced firm when they are a non-shopper, and $\lambda + \left(\frac{1-\lambda}{2}\right) = \frac{1+\lambda}{2}$ is the probability that the buyer purchases from the lowest-priced firm. The latter can arise because the buyer is a shopper or because they fortunately purchase from the lowest-priced firm when they are a non-shopper. In contrast, if $p_1 = p_2$, all buyers choose between the firms randomly, so the probability that an individual buyer will purchase from a given firm is $\frac{1}{2}$.

The probability that firm i will sell $\frac{m}{b}$ units each to $k \in \{0,..., b\}$ buyers is

$$(7)\qquad \Pr\left(s_i = \frac{m}{b}k \middle| p_i, p_j\right) = \begin{cases} \binom{b}{k}\left(\frac{1+\lambda}{2}\right)^k\left(\frac{1-\lambda}{2}\right)^{b-k} & \text{if } p_i < p_j \\ \binom{b}{k}\left(\frac{1}{2}\right)^k\left(\frac{1}{2}\right)^{b-k} & \text{if } p_i = p_j \\ \binom{b}{k}\left(\frac{1-\lambda}{2}\right)^k\left(\frac{1+\lambda}{2}\right)^{b-k} & \text{if } p_i > p_j, \end{cases}$$

where s_i denotes firm i's realised sales. Taking the case of $p_i<p_j$, the above says that k buyers will purchase from firm i with a probability $\left(\frac{1+\lambda}{2}\right)^k$ and $b-k$ buyers will purchase from firm j with a probability $\left(\frac{1-\lambda}{2}\right)^{b-k}$, where there are $\binom{b}{k} = \frac{b!}{k!(b-k)!}$ ways of distributing the k purchases among the b buyers. The other rows of (7) can be interpreted similarly.

In this example, the firms will always be unsure as to whether the sales distribution has resulted from pure chance or a deviation. So, to resolve this non-trivial signal extraction problem, firms can use a 'tail test' where a punishment phase is triggered whenever a sufficiently unlikely event occurs.[15] Consider a tail test, where the public signal is negative, $y_t = \underline{y}$, if at least K buyers purchase from one firm, $\frac{b}{2} + 1 \le K \le b$; otherwise, the public signal is positive, $y_t = \overline{y}$. Then, denoting the set of sales that generate a negative signal as Ω, where for tail tests $\Omega=\{K, K+1, \ldots, b-1, b\}$, the probabilities of a false negative and a true negative are, respectively:

$$(8)\qquad \begin{aligned} \alpha(K) &= \sum_{k\in\Omega} 2 \binom{b}{k}\left(\frac{1}{2}\right)^k\left(\frac{1}{2}\right)^{b-k} \\ \beta(K) &= \sum_{k\in\Omega} \binom{b}{k}\left[\left(\frac{1+\lambda}{2}\right)^k\left(\frac{1-\lambda}{2}\right)^{b-k} + \left(\frac{1+\lambda}{2}\right)^{b-k}\left(\frac{1-\lambda}{2}\right)^{k}\right] \end{aligned}$$

Thus, $\alpha(K)$ and $\beta(K)$ are determined by the tails of binomial distributions. Notice that $\alpha(K) < \beta(K) \le 1$ for all $\lambda > 0$ and that the public signal is less likely to be negative when the 'trigger level' K is closer to b, $\alpha(K) < \alpha(K-1)$ and $\beta(K) < \beta(K-1)$ for all $\frac{b}{2} + 1 \le K \le b$. There is also an endogenous relationship where increased market transparency raises demand-side transparency (via $\frac{\partial\lambda}{\partial\theta} > 0$) and this raises transparency on the supply side (via $\frac{\partial\beta}{\partial\theta} = \frac{\partial\lambda}{\partial\theta}\frac{\partial\beta}{\partial\lambda} > 0$).

In this setting, V^* represents the optimal SPPE profits for a given trigger level K. Hence, we denote it as $V^*(K)$ and refer to it as the optimal SPPE under tail tests. We do likewise for $\delta^*(K)$ and $\alpha^*(K)$ that are now also functions of K. Proposition 6 solves for the trigger level that maximises $V^*(K)$ and shows how such profits change with market transparency. As we discuss

[15] Applied work has often focussed on such tests in settings like ours. For example, the prime example of a tail test is in Green and Porter [1984], where quantity-setting firms enter the punishment phase if the market price falls below a given trigger level.

below, such profits will be the maximal SPPE profits when there is no other (non-tail) test that increases profits further.

Proposition 6. For any $b \in (2, \infty)$ and $s > 0$, if λ is sufficiently close to 1, such that $\alpha(b) < \alpha^*(b)$, and if $\delta \geq \delta^*(b) \in (0, 1)$, then the optimal trigger level is $K^* = b$ and the optimal SPPE profits under tail tests are $V^*(b) \in \left(\pi^N, \pi^c\right)$. Increasing market transparency strictly increases the optimal SPPE profits under tail tests, $\frac{dV^*(b)}{d\theta} > 0$.

Proposition 6 shows that if firms are sufficiently patient, then they will optimally set the trigger level to $K = b$. Intuitively, a stricter tail test with $K-1 < K$ will i) reduce the collusive profits due a higher probability of a false negative, $\alpha(K-1) > \alpha(K)$, and ii) raise the collusive profits due to a higher probability of a true negative, $\beta(K-1) > \beta(K)$. However, the former effect always dominates, implying that the optimal SPPE profits under tail tests are $V^*(b)$. Such profits are strictly increasing in market transparency, $\frac{dV^*(b)}{d\theta} > 0$. This contrasts with Example 1, where the optimal profits are independent of market transparency. The reason for the difference is that, while the pro-competitive demand-side effects (via $\frac{\partial\lambda}{\partial\theta}$) and the anti-competitive supply-side effects (via $\frac{\partial\beta}{\partial\theta}$) are larger than compared with Example 1, the anti-competitive supply-side effects now strictly dominate the pro-competitive demand-side effects.

Let us end by discussing when tail tests generate the maximal SPPE profits. This is the case when firms can do no better under a non-tail test, where there is some $\eta \in \{K+1, \ldots, b-1, b\}$ that does not generate a negative signal, $\eta \notin \Omega$. We next establish that a tail test with $K = b$ generates the maximal SPPE profits when λ is close to 1. To understand the intuition, consider when demand-side transparency is almost perfect, $\lambda \to 1$. Any test in equilibrium must then have $b \in \Omega$, so that the public signal is negative when all b buyers purchase from one firm. The reason is that if $b \in \Omega$, then $\lim_{\lambda\to1}\beta = 1$ because all buyers will observe a deviation and purchase from the deviant; yet if $b \notin \Omega$, then $\lim_{\lambda\to1}\beta = 0$. Of course, a tail test with $K = b$ satisfies this criterion. Now consider making the test harsher by including any $k < b$, so that $k \in \Omega$. This has no positive effect on β, since $\lim_{\lambda\to1}\beta = 1$ when $b \in \Omega$, and hence it also has no positive effect on V^*. Moreover, this harsher test also strictly increases α, which strictly lowers V^*. Thus, when the demand side is sufficiently transparent, such that λ is close to 1, firms cannot increase profits by making the test harsher using a non-tail test.

VI. ROBUSTNESS

Up to this point, we have modelled the competitive environment using Stahl [1989]. This raises the question whether our results are robust to other competition games. Consequently, the aim of this section is to show that the main results will extend to different environments. We proceed by rewriting Proposition 1 in terms of the underlying per-period profits, π^N, π^c and π^d, so the competition game has the structure of a noisy prisoners' dilemma. By using (5) and following the steps in the Proof of Proposition 1, we derive the following corollary, where the profits in (9) are the same as in (6), except that in the latter $\pi^c = v\frac{m}{2}$ and $\pi^d = (1+\lambda)\,v\frac{m}{2}$.[16]

Corollary 1. For any $\pi^d > \pi^c > \pi^N \geq 0$ and $0 < \beta \leq 1$, if $0 < \alpha < \beta\left(\frac{\pi^c - \pi^N}{\pi^d - \pi^N}\right)$ and $\delta \geq \frac{1}{(1-\alpha)+(\beta-\alpha)\left(\frac{\pi^c-\pi^N}{\pi^d-\pi^c}\right)}$, then the optimal PPE profits under trigger strategies are:

$$V^* = \pi^c - \frac{\alpha}{\beta - \alpha}\left(\pi^d - \pi^c\right) \in \left(\pi^N, \pi^c\right). \tag{9}$$

Corollary 1 implies that $\frac{\partial \pi^d}{\partial \lambda} > 0$ and $\frac{\partial \pi^c}{\partial \lambda} = 0$ suffice for the collusive profits to decrease as demand-side transparency rises, $\frac{\partial V^*}{\partial \lambda} = -\frac{\alpha}{(\beta-\alpha)}\frac{\partial \pi^d}{\partial \lambda} < 0$. These effects were assumed by Møllgaard and Overgaard [2002] in their general analysis of demand-side transparency, so the results of Proposition 3 will apply to other competitive environments. The sign of $\frac{\partial \pi^N}{\partial \lambda}$ is unimportant because, as noted in Section IV, the direct effect of a change in π^N, $\frac{\partial V^c}{\partial \pi^N}$, is perfectly offset by the change in the optimal punishment duration, $\frac{\partial V^c}{\partial T}\frac{\partial T^*}{\partial \pi^N}$. In relation to market transparency, all indirect effects remain infintesimally small at the perfect monitoring limit, because (3) is independent of the duration of the punishment phase, $\lim_{\alpha \to 0} \frac{\partial V^c}{\partial T} = 0$. Thus, consistent with Proposition 4, only the direct effect through α can be positive, so that the anti-competitive supply-side effects will dominate at the perfect monitoring limit regardless of the signs of $\frac{\partial \pi^d}{\partial \lambda}$ and $\frac{\partial \pi^N}{\partial \lambda}$.

We can now address why search costs have a different effect compared to the fraction of shoppers in our model. The reason is that only the static Nash equilibrium profit is affected by search costs in Stahl [1989], so that $\frac{\partial \pi^d}{\partial s} = 0$, $\frac{\partial \pi^c}{\partial s} = 0$ and $\frac{\partial \pi^N}{\partial s} \geq 0$. Thus, the optimal SPPE profits are independent of search costs, because the two counteracting effects from increasing π^N perfectly offset each other and there is no effect that works through π^d. However, this result is specific to Stahl [1989] and arises due to non-shoppers not undertaking costly search in equilibrium. In an alternative search framework developed by Wolinsky [1986], which is also analysed by Petrikaitė [2016], consumers sequentially search to discover their match values for differentiated products. In this framework, buyers do undertake costly search in equilibrium and so a deviation attracts more consumers when search costs are lower, $\frac{\partial \pi^d}{\partial s} < 0$; but search costs do not affect per-period collusive profits when

[16] The profits in (9) are also the maximal SPPE profits in the noisy prisoners' dilemma (see, e.g., Abreu *et al.* 1991).

search costs are low, $\frac{\partial \pi^c}{\partial s} = 0$. Thus, if our monitoring structure was applied to this alternative framework, then it would follow from above that increasing demand-side transparency, by reducing search costs, would be pro-competitive, $\frac{\partial V^*}{\partial s} = -\frac{\alpha}{(\beta-\alpha)}\frac{\partial \pi^d}{\partial s} > 0$.

VII. CONCLUDING REMARKS

In this paper, we have analysed both the supply- and demand-side effects of transparency in a model of collusion under imperfect monitoring, where punishment phases occur on-the-equilibrium path. We showed that in general the total effect of increasing market transparency on the optimal SPPE profits consists of a 'pro-competitive' demand-side effect and an 'anti-competitive' supply-side effect. Both effects are smaller as monitoring becomes less imperfect, but the pro-competitive demand-side effect vanishes at the perfect monitoring limit. Consequently, the anti-competitive supply-side effects dominate at the limit, implying that an increase in market transparency raises collusive profits. After endogenising monitoring, we also showed that the anti-competitive supply-side effects on the optimal SPPE profits dominate away from the limit.

Our results have three main policy implications. First, while it is generally presumed that policy interventions that increase demand-side transparency will hinder collusion when the supply side is unaffected, our model suggests that the pro-competitive demand-side effects will be smaller in markets where monitoring is easier. Second, our model reinforces the view that interventions that increase demand-side transparency should be avoided when they are also likely to facilitate collusion through increased supply-side transparency. The supply-side effects tend to dominate the demand-side effects and this result holds quite generally when monitoring is close to perfect. Finally, prohibiting facilitating practices that firms use to enhance supply-side transparency can undermine collusion even when this reduces transparency on the demand side.

APPENDIX

Proof of Proposition 1. The left-hand side of (5) is strictly less than the expression in square brackets on the right-hand side for all $\delta < 1$, so (5) can only hold if both sides are non-negative, since $\delta^T \in (0, 1]$ for all $T \in [0,\infty)$. Thus, a PPE in trigger strategies with $V^c(\alpha, \lambda, s, T) > \pi^N$ must satisfy three conditions: i) $\alpha \leq \alpha^*$, so that the right-hand side of (5) is non-negative; ii) $\delta \geq \delta^*$, so that the left-hand side of (5) is non-negative, where $\delta^* \in \left(\frac{1}{2}, 1\right)$ if $\alpha < \alpha^*$; iii) $T \geq T^*(\alpha, \beta, \lambda, s)$, where $T^*(\alpha, \beta, \lambda, s)$ denotes the level of T that ensures (5) holds with equality. To solve for the optimal PPE profits under trigger strategies, note that $\frac{\partial V^c}{\partial T} < 0$ such that $V^* = V^c\left(\alpha, \lambda, s, T^*(\alpha, \beta, \lambda, s)\right)$. Thus, it follows from (5) that:

$$(10) \qquad 1-\delta^{T^*(\alpha,\beta,\lambda,s)}=\frac{(1-\delta)\left(\pi^d-\pi^c\right)}{\delta\left[\beta\left(\pi^c-\pi^N\right)-\alpha\left(\pi^d-\pi^N\right)\right]}.$$

Substituting (10) into (3) shows that V^* is as claimed, where $V^*\in\left(\pi^N,\pi^c\right)\ \forall 0<\alpha<\alpha^*$. ■

Proof of Proposition 2. Differentiating V^* with respect to α and β yields $\frac{\partial V^*}{\partial\alpha}=-\frac{\beta\lambda v\frac{m}{2}}{(\beta-\alpha)^2}<0$ and $\frac{\partial V^*}{\partial\beta}=\frac{\alpha\lambda v\frac{m}{2}}{(\beta-\alpha)^2}>0$, respectively. The second-order derivatives are $\frac{\partial^2 V^*}{\partial\alpha^2}=-\frac{2\beta\lambda v\frac{m}{2}}{(\beta-\alpha)^3}<0$, $\frac{\partial^2 V^*}{\partial\beta^2}=-\frac{2\alpha\lambda v\frac{m}{2}}{(\beta-\alpha)^3}<0$ and $\frac{\partial^2 V^*}{\partial\alpha\partial\beta}=\frac{(\alpha+\beta)\lambda v\frac{m}{2}}{(\beta-\alpha)^3}>0$. ■

Proof of Proposition 3. Differentiating V^* with respect to λ gives $\frac{\partial V^*}{\partial\lambda}=-\frac{\alpha v\frac{m}{2}}{\beta-\alpha}<0$. The second-order derivatives are $\frac{\partial^2 V^*}{\partial\lambda^2}=0$, $\frac{\partial^2 V^*}{\partial\lambda\partial\alpha}=-\frac{\beta v\frac{m}{2}}{(\beta-\alpha)^2}<0$ and $\frac{\partial^2 V^*}{\partial\lambda\partial\beta}=\frac{\alpha v\frac{m}{2}}{(\beta-\alpha)^2}>0$. ■

Proof of Proposition 4. The total effect of increased market transparency on V^* is $\frac{dV^*}{d\theta}=\frac{\partial\alpha}{\partial\theta}\frac{\partial V^*}{\partial\alpha}+\frac{\partial\beta}{\partial\theta}\frac{\partial V^*}{\partial\beta}+\frac{\partial\lambda}{\partial\theta}\frac{\partial V^*}{\partial\lambda}$. From Proposition 2 and 3, $\frac{\partial V^*}{\partial\alpha}=-\frac{\beta\lambda v\frac{m}{2}}{(\beta-\alpha)^2}$, $\frac{\partial V^*}{\partial\beta}=\frac{\alpha\lambda v\frac{m}{2}}{(\beta-\alpha)^2}$ and $\frac{\partial V^*}{\partial\lambda}=-\frac{\alpha v\frac{m}{2}}{\beta-\alpha}$, such that:

$$(11) \qquad \frac{dV^*}{d\theta}=\frac{\lambda v\frac{m}{2}}{\theta(\beta-\alpha)}\left[-\theta\frac{\partial\alpha}{\partial\theta}\frac{\beta}{\beta-\alpha}+\alpha\left(\frac{\theta}{\beta-\alpha}\frac{\partial\beta}{\partial\theta}-\varepsilon_{\lambda,\theta}\cdot\varepsilon_{(\pi^d-\pi^c),\lambda}\right)\right]$$

Thus, it follows from $\lim_{\theta\to\bar\theta}\alpha=0$, $\lim_{\theta\to\bar\theta}\frac{\partial\alpha}{\partial\theta}\le 0$, and $\beta>0$ that $\lim_{\theta\to\bar\theta}\frac{dV^*}{d\theta}=-\lim_{\theta\to\bar\theta}\frac{\partial\alpha}{\partial\theta}\frac{\lambda v\frac{m}{2}}{\beta}\ge 0$, where the sign is strict if $\lim_{\theta\to\bar\theta}\frac{\partial\alpha}{\partial\theta}<0$. ■

Proof of Proposition 5. Substituting in $\alpha=\left(\frac{1}{2}\right)^{b-1}$ and $\beta=\lambda+(1-\lambda)\left(\frac{1}{2}\right)^{b-1}$ yields $V^*=v\frac{m}{2}\left(\frac{1-\left(\frac{1}{2}\right)^{b-2}}{1-\left(\frac{1}{2}\right)^{b-1}}\right)$ where $\frac{\partial V^*}{\partial\lambda}=0$ such that $\frac{dV^*}{d\theta}=\frac{\partial\lambda}{\partial\theta}\frac{\partial V^*}{\partial\lambda}=0$. Next recall $\alpha<\alpha^*$ implies $\pi^N<V^*$. Thus, given i) $\lim_{\lambda\to 1}\pi^N=0<V^*$ for all $b>2$, ii) $lim_{\lambda\to 0}\pi^N=v\frac{m}{2}>V^*$ for all $b<\infty$, and iii) $\frac{\partial\pi^N}{\partial\lambda}<0$ and $\frac{\partial V^*}{\partial\lambda}=0$, it follows that there exists a unique $\lambda^*\in(0,1)$ that sets $\alpha=\alpha^*$ for any $b\in(2,\infty)$, where $\alpha<\alpha^*$ if $\lambda>\lambda^*$. ■

Proof of Proposition 6. First, we establish $V^*(K)>V^*(K-1)$ for all $\frac{b}{2}+1\le K\le b$. From Proposition 1, $V^*(K)=v\frac{m}{2}\left(1-\frac{\alpha(K)\lambda}{\beta(K)-\alpha(K)}\right)$ where $V^*(K)>V^*(K-1)$ if:

$$(12) \qquad \beta(K)>\alpha(K)\left[\frac{\beta(K-1)-\beta(K)}{\alpha(K-1)-\alpha(K)}\right]$$

Note from (8) that:

$$\beta(K-1)-\beta(K)=\left[\frac{\alpha(K-1)-\alpha(K)}{2}\right]\left((1+\lambda)^{K-1}(1-\lambda)^{b-K+1}+(1+\lambda)^{b-K+1}(1-\lambda)^{K-1}\right).$$

Substituting into (12) then yields:

$$\beta(K)>\frac{\alpha(K)}{2}\left[(1+\lambda)^{K-1}(1-\lambda)^{b-K+1}+(1+\lambda)^{b-K+1}(1-\lambda)^{K-1}\right],$$

where if $\lambda = 0$, the expression in square brackets equals 2 and $\beta(K) = \alpha(K)$. Then from this and (8) it suffices to show that, for all $\sigma \in \{0, \dots, b-K\}$ and all $K \in \left\{\frac{b}{2}+1, \dots, b\right\}$:

(13)
$$(1+\lambda)^{b-\sigma}(1-\lambda)^{\sigma}+(1+\lambda)^{\sigma}(1-\lambda)^{b-\sigma} > (1+\lambda)^{K-1}(1-\lambda)^{b-K+1}+(1+\lambda)^{b-K+1}(1-\lambda)^{K-1}$$

Rearranging (13) yields:

$$(1+\lambda)^{\sigma}(1-\lambda)^{\sigma}\left[(1+\lambda)^{b-K+1-\sigma}-(1-\lambda)^{b-K+1-\sigma}\right]\left[(1+\lambda)^{K-1-\sigma}-(1-\lambda)^{K-1-\sigma}\right]>0,$$

for all $\lambda > 0$, $K \in \left\{\frac{b}{2}+1, \dots, b\right\}$ and $\sigma \in \{0, \dots, b-K\}$. Thus, $V^{*}(K) > V^{*}(K-1)$ for all $\frac{b}{2}+1 \leq K \leq b$ and $\lambda > 0$, such that $V^{*}(K)$ is at its highest value when $K = b$.

We next establish that $\frac{dV^{*}(b)}{d\theta} > 0$ for any $0<\lambda<1$. Given $\frac{\partial\alpha}{\partial\theta} = 0$, it follows from (11) that a necessary and sufficient condition for $\frac{dV^{*}(b)}{d\theta} > 0$ is

$$\frac{\theta}{\beta-\alpha}\frac{\partial\beta}{\partial\theta} > \varepsilon_{\lambda,\theta}\cdot\varepsilon_{(\pi^{d}-\pi^{c}),\lambda}$$

Under Example 2, $\varepsilon_{(\pi^{d}-\pi^{c}),\lambda} = 1$ and $\frac{\theta}{\beta-\alpha}\frac{\partial\beta}{\partial\theta} = \varepsilon_{\lambda,\theta}\cdot\left(\frac{\lambda\frac{\partial\beta}{\partial\lambda}}{\beta-\alpha}\right)$, such that $\frac{dV^{*}(b)}{d\theta} > 0$ if $\lambda\frac{\partial\beta}{\partial\lambda} > \beta-\alpha$. Substituting in for α, β and $\frac{\partial\beta}{\partial\lambda} = \frac{b}{2}\left[\left(\frac{1+\lambda}{2}\right)^{b-1}-\left(\frac{1-\lambda}{2}\right)^{b-1}\right] > 0$ shows that $\lambda\frac{\partial\beta}{\partial\lambda} > \beta-\alpha$ if:

(14)
$$b\lambda\left((1+\lambda)^{b-1}-(1-\lambda)^{b-1}\right) > (1+\lambda)^{b}+(1-\lambda)^{b}-2$$

To prove (14) holds for all $\lambda > 0$, first note that if $\lambda = 0$, then the left-hand side (LHS) equals the right-hand side (RHS). Then note $\frac{\partial \text{LHS}}{\partial\lambda} > \frac{\partial \text{RHS}}{\partial\lambda}$ for all $\lambda > 0$, from:

$$\frac{\partial \text{LHS}}{\partial\lambda}=b\left((1+\lambda)^{b-2}-(1-\lambda)^{b-2}\right)+b^{2}\lambda\left((1+\lambda)^{b-2}+(1-\lambda)^{b-2}\right)$$

$$\frac{\partial \text{RHS}}{\partial\lambda}=b\left((1+\lambda)^{b-2}-(1-\lambda)^{b-2}\right)+b\lambda\left((1+\lambda)^{b-2}+(1-\lambda)^{b-2}\right).$$

Consequently, it follows from the above that $\frac{\theta}{\beta-\alpha}\frac{\partial\beta}{\partial\theta} > \varepsilon_{\lambda,\theta}\cdot\varepsilon_{(\pi^{d}-\pi^{c}),\lambda}$ such that $\frac{dV^{*}(b)}{d\theta} > 0$ for all $0<\lambda<1$.

Finally, given $\alpha<\alpha^{*}$ implies $\pi^{N}<V^{*}$, then from i) $\lim_{\lambda\to 1}\pi^{N} = 0<V^{*}$ for all $b > 2$, ii) $\lim_{\lambda\to 0}\pi^{N} = v\frac{m}{2} > V^{*}$ for all $b<\infty$, and iii) $\frac{\partial\pi^{N}}{\partial\lambda} < 0$ and $\frac{\partial V^{*}}{\partial\lambda} > 0$, it follows that there exists a unique $\lambda^{*}(b) \in (0,1)$ for any $b \in (2,\infty)$ that sets $\alpha = \alpha^{*}(b)$, where $\alpha<\alpha^{*}(b)$ if $\lambda > \lambda^{*}(b)$. ■

REFERENCES

Abreu, D.; Pearce, D. and Stacchetti, E., 1986, 'Optimal Cartel Equilibria with Imperfect Monitoring,' *Journal of Economic Theory*, 39(1), pp. 251–269.

Abreu, D.; Pearce, D. and Stacchetti, E., 1990, 'Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,' *Econometrica*, 58(5), pp. 1041–1063.

Abreu, D.; Milgrom, P. and Pearce, D., 1991, 'Information and Timing in Repeated Partnerships,' *Econometrica*, 59(6), pp. 1713–1733.

Albæk, S.; Møllgaard, H. P. and Overgaard, P., 1997, 'Government-Assisted Oligopoly Coordination? A Concrete Case,' *Journal of Industrial Economics*, 45(4), pp. 429–443.

Amelio, A. and Biancini, S., 2010, 'Alternating Monopoly and Tacit Collusion,' *Journal of Industrial Economics*, 58(2), pp. 402–423.

Fudenberg, D. and Tirole, J., 1994, *Game Theory* (MIT Press, Cambridge, Massachusetts, U.S.A.).

Garrod, L.; Hviid, M.; Loomes, G. and Waddams Price, C., 2008, 'Assessing the Effectiveness of Potential Remedies in Final Consumer Markets,' report OFT994. (Office of Fair Trading).

Garrod, L. and Olczak, M., 2017, 'Collusion under Imperfect Monitoring with Asymmetric Firms,' *Journal of Industrial Economics*, 65(3), pp. 654–682.

Green, E. and Porter, R., 1984, 'Non-Cooperative Collusion under Imperfect Price Information,' *Econometrica*, 52(1), pp. 87–100.

Harrington, J. and Skrzypacz, A., 2007, 'Collusion under Monitoring of Sales,' *RAND Journal of Economics*, 38(2), pp. 314–331.

Herre, J. and Rasch, A., 2013, 'Customer-Side Transparency, Elastic Demand, and Tacit Collusion under Differentiation,' *Information Economics and Policy*, 25(1), pp. 51–59.

Janssen, M.; Moraga-González, J. and Wildenbeest, M., 2005, 'Truly Costly Sequential Search and Oligopolistic Pricing,' *International Journal of Industrial Organization*, 23, pp. 451–466.

Kühn, K.-U., 2001, 'Fighting Collusion by Regulating Communication between Firms,' *Economic Policy*, 16(32), pp. 167–197.

Kühn, K.-U. and Vives, X., 1995, *Information Exchanges among Firms and Their Impact on Competition* (European Commission, Office for Official Publications of the European Communities, Luxembourg).

Møllgaard, H. P. and Overgaard, P., 2001, 'Market Transparency and Competition Policy,' *Rivista di Politica Economica*, 91, pp. 11–58, reprinted in Baldassarri, M. and Lambertini, L. (eds.), *Antitrust, Regulation and Competition* (Palgrave Macmillan, 2003).

Møllgaard, H. P. and Overgaard, P., 2002, 'Market Transparency: A Mixed Blessing?' mimeo (Departments of Economics, Copenhagen Business School and University of Aarhus, Denmark) See *https://www.researchgate.net/publication/5160970_Market_Transparency_A_Mixed_Blessing*.

Møllgaard, H. P. and Overgaard, P., 2008, 'Information Exchange, Market Transparency and Dynamic Oligopoly,' in Collins, W. (ed.), *Issues in Competition Law and Policy* (American Bar Association, Section of Antitrust Law, Chicago, Illinois, U.S.A.).

OECD, 2001, 'Price Transparency,' document DAFFE/CLP(2001)22 (Organisation for Economic of Co-operation and Development, 2, rue André Pascal, Paris, France).

Petrikaitė, V., 2016, 'Collusion with Costly Consumer Search,' *International Journal of Industrial Organization*, 44, pp. 1–10.

Schultz, C., 2005, 'Transparency on the Consumer Side and Tacit Collusion,' *European Economic Review*, 49(2), pp. 279–297.

Schultz, C., 2017, 'Collusion in Markets with Imperfect Price Information on Both Sides,' *Review of Industrial Organization*, 50(3), pp. 287–301.

Stahl, D. O., 1989, 'Oligopolistic Pricing with Sequential Consumer Search,' *American Economic Review*, 79(4), pp. 700–712.

Stigler, G., 1964, 'A Theory of Oligopoly,' *Journal of Political Economy*, 72, pp. 44–61.

Sugaya, T. and Wolitzky, A., 2018, 'Maintaining Privacy in Cartels,' *Journal of Political Economy*, 126(6), pp. 2569–2607.

Tirole, J., 1988, *The Theory of Industrial Organisation* (MIT Press, Cambridge, Massachusetts, U.S.A.).

Varian, H., 1980, 'A Model of Sales,' *American Economic Review*, 70(4), pp. 651–659.

Wolinsky, A., 1986, 'True Monopolistic Competition as a Result of Imperfect Information,' *Quarterly Journal of Economics*, 101(3), pp. 493–512.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of this article at *http://wileyonlinelibrary.com/journal/joie* or via the Journal's website, *http://www.jindec.org*.

A.1 Relation to Tirole [1988], p. 262-264

A.2 Critical Discount Factor

A.3 Optimal Symmetric Equilibrium Profits