The Role of Information and Monitoring on Collusion

Christian Rojas[1]

Abstract

This paper reports the results of experiments constructed to test the effects of demand information (knowledge of the demand function or schedule) and monitoring (knowledge of rivals' actions) on collusion in infinitely repeated games. In addition, the treatments encompass the assumptions of two highly influential theories (Green and Porter [1984] and Rotemberg and Saloner [1986]), thereby allowing for a test of the predictions of each theory. Given the numerous theoretically possible equilibria, this last exercise is important as it improves our understanding of which equilibria appear more *empirically plausible*. Results indicate that monitoring is a key factor in facilitating collusion, but, contrary to conventional wisdom, demand information does not improve collusion and in some cases it may even decrease cooperation. Both theories tested receive empirical support as possible explanations of behavior; however, in both cases data appears to be best described by permanent price wars (i.e. grim-trigger strategies) rather than by the temporary price wars for which both theories are known. We discuss these findings within the broad set of possible strategies and equilibria.

[1] Department of Resource Economics, University of Massachusetts at Amherst, rojas@resecon.umass.edu.
I am extremely thankful to Catherine Eckel for her encouragement, support and constructive comments. I also thank Vince Crawford, Chetan Dave, Alex Yuskavage, Ernan Haruvy, John Spraggon and Bart Wilson for helpful suggestions. Sheryl Ball provided grant administration support. Funding was provided by the National Science Foundation, grant SES-0526229 and by the Department of Resource Economics at the University of Massachusetts Amherst.

1. Introduction

Firms' coordination to obtain high profits has been a continuous concern for researchers and antitrust authorities. As a consequence, there exists a large body of theoretical work on the factors that determine the likelihood of collusion. But analyzing collusion empirically is difficult because the illegal status of cartels makes field data scarcely available. Importantly, with many exogenous and unobservable factors in field data, the task of identifying and estimating the effect of different market conditions on collusion becomes problematic. One objective of this paper is to improve the understanding of the role of two factors that have been prominent in models of repeated interaction with demand uncertainty: demand information (knowledge of the demand function or schedule) and monitoring (knowledge of rivals' actions). The general strategy is to analyze the effects of these factors on collusion by generating data from controlled experiments that resemble various demand information and monitoring conditions.

The motivation comes from two models that have been highly influential in the development of the theoretical and empirical literature on cartel stability: Green and Porter (1984) and Rotemberg and Saloner (1986) [GP and RS henceforth]. Both models assume an uncertain (stochastic) demand structure, but differ on their assumption about firms' information regarding the actual demand realization (e.g. high, low). RS assume that firms have perfect foresight about demand next period (i.e. the demand realization can be anticipated) whereas GP assume that firms are always uncertain about (future and past) demand realizations. In addition, GP assume that monitoring among cartel members is imperfect (i.e. comes in the form of a noisy public signal),[2] whereas RS assume that monitoring is perfect. Our experimental design is guided by these differences in assumptions: in two of our treatments monitoring and demand information differ in the same way as GP and RS differ. In a third, 'intermediate', treatment there is uncertainty about next period's demand realization (as in GP) but monitoring is perfect (as in RS); this treatment allows us to separate the imperfect monitoring effect from the imperfect demand information effect.

[2] This assumption is needed so that uncertainty about past demand realizations persists into the future.

Interestingly, the differential treatment of demand information and monitoring assumptions by GP and RS generate theoretical predictions that appear to be "at odds" (Ellison, 1994: 38). GP show that price wars in a cartel may be triggered by unusually small demand shocks, but RS show that a cartel may experience price wars during periods of unusually large demand shocks.[3] A usual interpretation of these theories is that GP predicts more collusion during booms, while RS predicts more collusion during recessions.[4] Given the divergent predictions of these theories and our encompassing experimental design, we test each theory's internal validity; this is the second objective of this research.

It is important to point out, however, that GP and RS are not mutually exclusive theories, and results from our second objective may well indicate that each model is valid in its own domain. Instead, our effort is to investigate the empirical plausibility of each theory. Studying each theory's internal validity is important for at least two reasons. First, with field data there is no guarantee that firm behavior that appears to correspond to the predictions of a given theory (even if the theory's assumptions appear to hold) is a consequence of the theory at work (e.g. Frechette, Kagel and Morelli, 2005). Second, from a practical perspective, the multiplicity of equilibria in infinitely repeated games and the large number of theories on collusion makes it important for empirical economists to identify the empirically *plausible* equilibria from the set of theoretically *possible* equilibria; experiments can be a particularly useful tool in this effort. For example, if the predictions of a collusive theory hardly emerge in a simple controlled environment, then there should be more skepticism in labeling an industry as collusive when its behavior resembles the predictions of such theory.

[3] RS is also interpreted as a model of 'countercyclical' pricing. Here, the alternative interpretation of price wars during booms is adopted since it fits better with the experimental design of two strategies (collusion and non-collusion). Because of our specific design, the terms "price wars", "non-collusive outcome" and "competitive outcome" will be used interchangeably throughout the paper.

[4] A subtle difference is that in GP the duration of a price war is endogenously determined (i.e. the price war is triggered by low demand and remains in place regardless of demand realizations during the 'punishment' period), whereas in RS a price war is state dependent: it only emerges when the demand shock is unusually large.

Specifically, we test whether the collusive equilibrium path for which GP is known (finite price wars triggered by low demand) is supported by the data. For the RS theory, we test the equilibrium prediction that price wars should occur during high demand periods whereas collusion should occur otherwise. We also study strategies at the individual level to determine whether the strategies implied by each theory constitute a reasonable explanation of behavior when compared with other plausible strategies. Finally, we study how the RS and GP equilibria explain behavior with respect to other possible equilibria.

Results indicate that monitoring appears to be a critical factor in facilitating collusion. Conversely, contrary to conventional wisdom, demand information does not appear to have the expected effect on collusion: removing demand information does not decrease (and in some cases increases) collusion. The results provide some support for both the RS and the GP predictions; however, evidence appears to be stronger for permanent price wars (i.e. grim-trigger strategies) rather than the temporary reversions for which both theories are known for. This is important as one of the several GP equilibria allows for permanent price war, while the RS equilibria do not permit this possibility.

Section 2 reviews the literature while section 3 describes the model. Section 4 provides details of the experimental design and section 5 describes its implementation. Section 6 presents the results and section 7 discusses our main findings.

2. Literature Review

Friedman (1971) showed that if firms are patient enough in a non-cooperative infinitely repeated game, a trigger strategy (reversion to Cournot production levels when market price dropped below a threshold) would produce an equilibrium in which no firm has an incentive to deviate. According to this early view, the existence of price wars in oligopoly markets was interpreted as a sign of cartel breakdown. However, GP show that instead of a symptom of unsuccessful collusion, finite price wars may be part of a 'collusive' equilibrium path. GP modify Friedman's model by allowing for a stochastic demand structure and imperfect monitoring of rivals' actions.[5] As in Friedman, collusion can be sustained through the use of trigger strategies, but now switching from the collusive outcome to the competitive outcome (after an unusually low demand state) is only temporary. More importantly, the seminal result of this model is that price wars are part of the collusive equilibrium path as they constitute a self-enforcing mechanism used by successful colluders.

RS propose a model with a stochastic environment similar to that of GP. The main differences between RS and GP is that firms know next period's demand shock realization prior to setting their quantity (or price) and that firms can perfectly monitor rivals' choices. In this environment, firms' incentive to deviate from the collusive outcome is positively correlated with next period's demand shock and for unusually large (and positive) demand shocks this incentive more than offsets the expected future losses of a reversion to the competitive outcome. A cartel is thereby predicted to be less stable during "booms". To avoid the competitive outcome during large demand shock periods, firms limit the incentives to deviation by reducing (increasing) their "collusive" price (quantity) below (above) the monopoly level. The resulting collusive equilibrium path has firms pricing in a countercyclical fashion.[6]

Empirical work assessing the validity of the GP theory has been restricted to data from the 19th century Joint Executive Committee (JEC). However, limited data and an uncontrolled field environment do not allow a direct test of the GP theory. As a result most of the work with these data has been concerned with finding evidence for the existence of regime switching between high and low prices (see Ellison, 1994, and references cited therein).[7] Empirical work on the RS theory has focused on its countercyclical pricing prediction; this has been a puzzling issue as it runs counterintuitive to conventional wisdom (i.e. a rightward shift in demand should increase

---

[5] Note that with imperfect monitoring, uncertainty about demand in period t is extended to periods $t+i$, $i=1,\ldots\infty$.

[6] RS consider IID demand shocks. Haltiwanger and Harrington (1991) and Kandori (1991) show that, under certain conditions, countercyclical pricing also holds with cyclical demand, but Bagwell and Staiger (1997) show that in other stochastic conditions the RS results are reversed.

[7] Ellison also uses the JEC data to search for evidence supporting the RS theory; however, he finds better support for the GP theory.

equilibrium price) but is nevertheless frequently observed in many markets (e.g. soft drinks during summer, turkey during Thanksgiving). Prior research has tried to explain this pattern against other competing models and has found little support for RS as an explanatory theory (Chevalier, Kashyap and Rossi, 2003).

On the experimental front, there has been work studying how demand information and monitoring affect collusion/cooperation in repeated games. This literature has addressed either monitoring or demand information, but not both, and in rather specific ways. The role of imperfect monitoring on collusion has been studied Holcomb and Nelson (1991, 1997), Bereby-Meyer and Roth (2006), and Aoyagi and Frechette (2008). Holcomb and Nelson study repeated duopoly games in which opponent's quantity choices are randomly changed by the experimenter 50% of the time; they find that collusion increases in treatments where this manipulation of the opponent's quantity choice is not present. Imperfect monitoring in this experiment, however, comes from a private signal administered by the experimenter to either subject rather than from noisy demand (as assumed by theory).

Bereby-Meyer and Roth study the speed of learning in the *finitely* repeated prisoner's dilemma game and find that cooperation decreases when payoffs are noisy, even when players can monitor opponents' actions. Aoyagi and Frechette study experiments on the infinitely repeated prisoner's dilemma game when players are given a noisy public signal about the pair of strategies (out of the four possible in the prisoner's dilemma game) chosen by two players.[8] The authors find that payoffs decrease with noise. The authors also test the existence of a threshold strategy, similar to that of GP, that switches between cooperation and punishment as a function of the noise and find that subjects follow a simple threshold strategy that considers only the most recent public signal.

Cason and Mason (1999) and Feinberg and Snyder (2002) investigate the role of demand information on collusion. Cason and Mason analyze how firms' decisions of whether to share information can facilitate collusion in an environment of demand

[8] A downside of this design is that to make imperfect monitoring possible, subjects are never informed of their period earnings but only of their accumulated earnings at the end of the rounds.

uncertainty. In different treatments, subjects accessed information about the state of demand in different ways. The main result is that subjects generally decided to share information to reduce uncertainty, which led to output reductions. However, in treatments where subjects did not have the choice of sharing information, information itself did not increase tacit collusion. Feinberg and Snyder claim that uncertain demand shocks do interfere with collusion, although few data points and an apparent 'group effect' do not allow a clear interpretation of the results.

3. The Model

The model is based on the prisoner's dilemma game. There are at least two reasons for studying collusion in an environment that is a highly simplified version of the models that motivate this research: a) we want to give collusion its best possible chance of occurrence - subjects' coordination on the collusive outcome is less likely if a game has multiple (or continuous) strategies, and b) we want to give theories their best possible chance of occurrence - if theories fail to produce the predicted results in simple environments, it is less likely that such results will be observed in more complex situations. In addition, simplicity of the experiment reduces subjects' confusion thereby producing more reliable results.

This simplified model can be thought of as a 2-firm Cournot game with homogeneous products, constant marginal cost, symmetric firms, and discrete choices.[9] Three demand states (High (h), Medium (m) and Low (l)) are assumed. There are two quantity choices: Low (L) and High (H). The game is infinitely repeated and three demand states (i.e. three prisoner's dilemma games) occur with probabilities 0.60 for Medium, and 0.20 for both High and Low. The payoff table for demand state "s" has the following structure:[10]

[9] Both RS and GP can be cast as either price or quantity games. With homogeneous products, quantity competition seems a more natural wording as the matrices used have non-zero payoffs at the Nash equilibrium. Without loss of generality, the prior language of "price war" should be interpreted as the Cournot outcome (H).

[10] Although 20% seems like reasonable probability for a positive (negative) demand shock, this choice of probabilities may seem somewhat arbitrary (i.e. it is not immediately clear how or whether results would be sensitive to the equilibria implied by different probabilities). However, the equilibria implied by a different distribution of demand states can also be obtained by changing the values of the payoff tables; we choose

Table 1: Typical Payoff Table

		Player 2	
		L	H
Player 1	L	Π_s^C, Π_s^C	Π_s^{ND}, Π_s^D
	H	Π_s^D, Π_s^{ND}	Π_s^{NE}, Π_s^{NE}

where, Π_s^C denotes the Collusive payoff, Π_s^D the payoff to a 'deviating' or cheating firm, Π_s^{ND} the payoff to a 'Non-deviating' firm that has been cheated upon, and Π_s^{NE} the Nash-Equilibrium payoff, with $\Pi_s^D > \Pi_s^C > \Pi_s^{NE} > \Pi_s^{ND}$. The subscript *s* represents the demand state (*h, m* or *l*). For a given demand state *s*, and if firms use a grim-trigger strategy (reverting to *H* forever after a firm deviates from *L*), it is well known that the collusive outcome (*L,L*) can be supported in equilibrium if (*s* temporarily omitted):

$$\begin{aligned} \Pi^D - \Pi^C &< \sum_{t=1}^{\infty} \delta^t (\Pi^C - \Pi^{NE}) \\ \Pi^D - \Pi^C &< \frac{\delta}{1-\delta} (\Pi^C - \Pi^{NE}) \end{aligned} \tag{1}$$

where δ is the discount factor. For a given set of payoffs, collusion will be sustainable if firms' discount factor is larger than the critical value, $\delta^* = \left[\Pi^D - \Pi^C\right] / \left[\Pi^D - \Pi^{NE}\right]$. A main feature in the RS and GP models, however, is that demand is stochastic (i.e. *s* is determined probabilistically); this assumption is critical for obtaining the models' predictions: 'price wars' (i.e. temporary deviations to *H*) are observed for some periods as a consequence of 'extreme' demand shocks and not as a consequence of a 'breakdown' in collusion.

The RS Model: Perfect Demand Foresight and Perfect Monitoring

Although demand is stochastic, RS assume that firms know next period's demand shock (and can monitor rivals' strategies) but are uncertain about future demand (i.e. firms know the distribution of the demand shock for periods *t+2* onwards, but not their

---

this latter route and consider three different payoff structures to check the robustness of our results (see sections 6 and 7).

realizations). RS assume that players use a grim-trigger strategy and show that for a given punishment (RHS of (1)) the incentive to deviate (LHS of (1)) is increasing in the demand shock. As a consequence, for a sufficiently high demand shock and a given discount factor, the incentive to deviate is greater than the future punishment and thus collusion breaks down. In the literature this breakdown is interpreted in two different ways: *a)* as countercyclical pricing and *b)* as price wars during booms. RS indicate that depending on how the strategy space is constructed the model can yield either result: if the strategy space is continuous, (smooth) countercyclical pricing will be observed in equilibrium; on the other hand, if the strategy space is constrained to either compete or collude (as is the case here), then price wars will be observed (RS: 396).

Accommodating the RS model to our design gives the modified version of equation (1):

$$\begin{aligned}&\Pi_s^D - \Pi_s^C < \frac{\delta}{1-\delta}[E(\Pi_s^C - \Pi_s^{NE})] \\ &\Pi_s^D - \Pi_s^C < \frac{\delta}{1-\delta}\left[0.2(\Pi_h^C - \Pi_h^{NE}) + 0.6(\Pi_m^C - \Pi_m^{NE}) + 0.2(\Pi_l^C - \Pi_l^{NE})\right]\end{aligned} \tag{2}$$

To test the RS theory, a set of payoff tables (parameterization 1, see section 4 below) are constructed so that equation (2) holds only for the medium (*m*) and low (*l*) demand states but not for the high (*h*) demand state. As is the case with infinitely repeated games, however, the RS prediction is one of several equilibria. We adapt the results in Stahl II (1991) to our stochastic setting and compute the set of equilibria that can be derived from a grim-trigger strategy in the parameterizations considered. In addition, we consider whether equilibria can be supported by finite punishment strategies and compute the corresponding minimum price war length. Appendix A contains computational details and the results of this exercise. Understanding the set of possible equilibria and the set of possible optimal punishment strategies allows us to compare the evidence for the RS equilibrium with respect to alternative explanations.

The GP Model: No Demand Foresight and Imperfect Monitoring

At any *t*, GP assume that firms only know the distribution of the demand schedule. In addition, firms can only imperfectly monitor their rivals' quantity choices through a noisy signal: the market price. These two sources of uncertainty impede firms from inferring their opponents' strategies even *after* the realization of demand. GP show that finite punishment strategies (reversion to *H*) can be sustained in an equilibrium path in which no firm deviates from the collusive agreement. These finite punishment periods are triggered by a low market price: since there is imperfect monitoring of rivals' actions, a low price can either denote a negative demand shock or a rival's deviation. As with other cartel models, GP entertain a collusive equilibrium in which no firm has an incentive to deviate from the agreement. Hence, in this sort of equilibrium, the low price can only be a consequence of a negative demand shock. Nonetheless, rational firms will want to punish themselves and their rivals during periods of low demand, otherwise the threat of reverting to the Nash-Equilibrium is not credible.

The GP model involves a more complex optimization problem in which the length of the punishment, the trigger price and the level of collusion have to be adjusted. An advantage of our experimental design is that it allows calibration of the payoff tables to a GP equilibrium with any desired punishment length. The GP equilibrium is based on the assumption that firms use a trigger strategy which establishes that all players reverse to the Nash equilibrium (H,H) for N periods when the noisy signal (price) falls below threshold level *k*.[11]

Denote firm i's choice as $y_i \in \{L, H\}$ and the expected discounted future payoffs of choosing y_i as $V_i(y_i)$. Also, denote $f(p \mid y)$ as the density function that determines the probability of observing price level *p* given the outcome $y = (y_i, y_j)$, and $F(p \mid y)$ its corresponding cumulative distribution function. The expected payoff in each reversionary period is given by: $\lambda_i = 0.2\Pi_h^{NE} + 0.6\Pi_m^{NE} + 0.2\Pi_l^{NE}$, and the (next-period) expected profit when the opponent is sticking to the collusive quantity (*L*) is:

$$\gamma_i(y_i) = \begin{cases} 0.2\Pi_h^C + 0.6\Pi_m^C + 0.2\Pi_l^C & \text{if } \; y_i = L \\ 0.2\Pi_h^D + 0.6\Pi_m^D + 0.2\Pi_l^D & \text{if } \; y_i = H \end{cases}$$

[11] Our payoff matrices are constructed to accommodate imperfect monitoring (see section 4). Appendix C describes how the implied price of our payoff matrices represents a noisy public signal (as required by the GP theory).

and the resulting Bellman's equation is:

$$V_i(y_i)=\begin{cases}\gamma_i(L)+(1-F_k^L)\delta V_i(L)+F_k^L\left[\sum_{t=1}^{N}\delta^t\lambda_i+\delta^{N+1}V_i(L)\right] & \text{if } y_i=L\\ \gamma_i(H)+(1-F_k^H)\delta V_i(H)+F_k^H\left[\sum_{t=1}^{N}\delta^t\lambda_i+\delta^{N+1}V_i(H)\right] & \text{if } y_i=H\end{cases} \quad (3)$$

where, $F_k^L=F(k\mid y_i=y_j=L)$ and $F_k^H=F(k\mid y_i=H;y_j=L)$. That is, for a given threshold level *k*, firm *i*'s decision shifts the probability with which the observed price level may fall below *k*. The GP equilibrium path specifies that a price war is triggered when the noisy signal falls below a given threshold level *k* and (since firms are colluding) the only way price may fall is because demand contracts. In our design, finding the *k* suggested by GP is straightforward: the price that corresponds to the collusive profit in the low demand state (Π_l^C), which we denote as $\tilde{p}$. Thus, for $\tilde{p}$ and a given punishment period N, the GP equilibrium exists if:

$$\gamma_i(L)+(1-F_{k=\tilde{p}}^L)\delta V_i(L)+F_{k=\tilde{p}}^L\left[\sum_{t=1}^{N}\delta^t\lambda_i+\delta^{N+1}V_i(L)\right]>$$

$$\gamma_i(H)+(1-F_{k=\tilde{p}}^H)\delta V_i(H)+F_{k=\tilde{p}}^H\left[\sum_{t=1}^{N}\delta^t\lambda_i+\delta^{N+1}V_i(H)\right]$$

Solving for $V_i(y_i)$ and after some manipulation, the equilibrium condition becomes:

$$V_i(L)=\frac{\gamma_i(L)+F_{k=\tilde{p}}^L\dfrac{\delta(1-\delta^N)}{1-\delta}\lambda_i}{1-(1-F_{k=\tilde{p}}^L)\delta-F_{k=\tilde{p}}^H\delta^{N+1}}>\frac{\gamma_i(H)+F_{k=\tilde{p}}^H\dfrac{\delta(1-\delta^N)}{1-\delta}\lambda_i}{1-(1-F_{k=\tilde{p}}^H)\delta-F_{k=\tilde{p}}^H\delta^{N+1}}=V_i(H) \quad (4)$$

Of course, there could be many values of *N* for which this condition holds (or none at all), including the *N*=∞. Similarly, there can be several thresholds (prices) for which one can find a finite *N* such that (4) holds. Computation of all feasible equilibria (including all feasible punishment lengths) can be found in Appendix A (table A.2). As with the case of the RS design, we employ the set of feasible equilibria to guide our econometric analysis of the likelihood of the GP equilibrium compared with other alternatives.

4. Experimental Design

Subjects play the prisoner's dilemma game for 30 rounds with certainty and the continuation probability is set to 0.75 thereafter; this simulates an infinitely repeated game with a discount factor (δ) of 0.75 (Fudenberg and Tirole, 1989). We do not implement a continuation probability in round 1 (as suggested by Dal Bó, 2005) because the focus of the paper is to create an environment with a horizon long enough where finite periods of reversion to competition (as predicted by theory) can emerge.[12]

In each round subjects simultaneously choose a quantity (High (*H*) or Low (*L*)) and payoffs are determined by one of three tables: High Demand (*h*), Medium Demand (*m*) and Low Demand (*h*), see table 2. One of the three tables is chosen each round with the probabilities indicated above; this probability is known by all subjects in all treatments, but the way in which demand information is presented varies across treatments.[13]

There are two sets of payoff matrices. Note that the main difference between these two parameterizations is that deviation is less attractive in parameterization 2: the additional profits from deviation are smaller, and reversion to NE play is more costly; the reason for two parameterizations is to investigate the robustness of our results and to create variation in our data. To illustrate the latter point, parameterization 1 implies that the GP equilibrium is not feasible (not even with the use of a grim strategy), while in

[12] Our design is a combination of the two types of treatments that Dal Bó proposes for disentangling the "shadow of the future" effect from that given by the increased number of expected rounds. In Dal Bó's analysis, one type of treatment has the random stopping rule implemented from round 1 (e.g. $\delta = 0.75$), whereas the other type of treatment has a finitely repeated game of equivalent expected length (e.g. $1/(1-\delta)= 1/(1- 0.75) = 4$ rounds). Dal Bó finds that under the first treatment cooperation is higher, but it is not clear if there exists an important difference in the strategies used that would compromise the intended incentives of our design (Dal Bó only indicates that strategies in the first type of treatment appear to be more consistent with the grim strategy, p. 1601). For our design, Dal Bó's main result means that the *level* of cooperation may be different than that observed under a design that implements a random stopping rule from period 1 (i.e. our design is perhaps more likely to have a lower level of cooperation as the number of fixed periods is large relative to the number of expected periods under the random stopping rule). Hence, it is not clear how using a mixed design (33 rounds for sure and then a random stopping rule) would result in strategies (or equilibrium outcomes) that are very different from those one would obtain if the random stopping rule were to be implemented from round 1.

[13] To keep treatments comparable, subjects in all sessions are not informed about their opponent's choice. This means that when monitoring is allowed the opponent's choice can be inferred by looking at the payoff tables.

parameterization 2 the GP equilibrium is feasible with a punishment length of at least 3 periods.

Table 2: Payoff Tables for Three Demand States

Parameterization 1

High Demand (probability=0.2)

		Player 2	
		L	H
Player 1	L	26.00 , 26.00	7.50 , 43.00
	H	43.00 , 7.50	12.50, 12.50

Medium Demand (probability=0.6)

		Player 2	
		L	H
Player 1	L	7.50 , 7.50	2.10 , 12.50
	H	12.50 , 2.10	3.50 , 3.50

Low Demand (probability=0.2)

		Player 2	
		L	H
Player 1	L	2.10 , 2.10	0.60 , 3.50
	H	3.50 , 0.60	1.00 , 1.00

Parameterization 2

High Demand (probability=0.2)

		Player 2	
		L	H
Player 1	L	31.00 , 31.00	9.00 , 43.00
	H	43.00 , 9.00	12.50, 12.50

Medium Demand (probability=0.6)

		Player 2	
		L	H
Player 1	L	9.00 , 9.00	2.50 , 12.50
	H	12.50 , 2.50	3.50 , 3.50

Low Demand (probability=0.2)

		Player 2	
		L	H
Player 1	L	2.50 , 2.50	0.70 , 3.50
	H	3.50 , 0.70	1.00 , 1.00

Thus, if the GP predictions are likely to occur, a natural test would be to check whether the predictions of the GP theory are more likely to occur in parameterization 2.

Likewise, parameterization 1 is calibrated in accordance with the predictions of the RS theory (collusion is an equilibrium only in the medium and low demand states), while parameterization 2 is not (i.e. collusion is an equilibrium in all demand states). Parameterization 2 also serves to check the robustness of our other main result, namely the more prominent role of monitoring (rather than demand information) on collusion.

Instead of specifying a demand function, the payoff tables are constructed so that the percentage difference between payoffs across entries remains invariant across demand states. For example, in parameterization 1 the payoff in the collusive outcome is about 100% higher (with some rounding error) than the payoff in the Nash-Equilibrium. The reason for constructing payoff matrices in this fashion is that individuals seem to care about relative variation in payoffs rather than the absolute variation (Weber, Shafir and Blais, 2004); thus, the potential confounding effect of significant variation in relative payoffs across demand states is reduced.

Another behavioral aspect considered is the possibility of risk aversion by subjects, which contrasts with the risk neutrality assumption of the theories we entertain.[14] Because risk aversion may compromise the external validity of our results, we estimate the level of risk aversion in our sample by including a risk measurement task in the experimental protocol.[15] To ensure that our results are robust to the presence of risk aversion, we carry out additional sessions with an alternative parameterization (parameterization 3, see appendix D); section 7 reports these results.

The first objective of the paper is to test the role of demand information and monitoring on collusion, as motivated by the difference in assumptions between the RS and GP theories. To this end, we design two treatments that resemble each theory's assumptions. We also consider a third treatment that contemplates an intermediate case.

i) Full Information Treatment (FI) [RS theory]: Before a round starts, subjects are told which payoff table they will play. Hence, uncertainty about next period's demand state is removed; demand states for future rounds remain unknown, however. Similarly, monitoring is possible because in our design it possible to infer the rival's

[14] See Harrison and Ruström (2008) for an extensive review of the evidence of risk aversion in the lab.

[15] Appendix B contains a brief description of the task and the estimation of risk aversion in our sample.

strategy after a round's profit realization. Payoff tables in parameterization 1 are constructed such that the incentive to deviate in the high demand state (LHS of (3)) is smaller than the expected value of a future infinite punishment (RHS of (3)). Parameterization 2, on the other hand, does not provide incentives to deviate in any of the three demand states. Comparisons of results between parameterizations 1 and 2 will indicate whether the incentives devised by RS work as originally intended.

ii) Imperfect Monitoring Treatment (IM) [GP theory]: Subjects only know the distribution about next period's demand state: demand next period will most likely be normal (medium, probability 0.60) but there is also a chance of experiencing a demand shock (high with probability 0.20 and low with probability 0.20).

With the constructed payoff tables, there is imperfect monitoring about the opponent's choice of output even after profits have been realized for that round. To see this, suppose that player 1 chooses to collude (*L*) and his profit turns out to be 2.10 (parameterization 1); if demand was low, it means that the opponent also chose the collusive outcome (*L*), but if demand was medium it means that the opponent deviated from it (*H*). This imperfect monitoring is possible whether a firm follows the collusive outcome (*L*) or the deviates from it (*H*).[16] Importantly, as assumed by GP, our design implies that imperfect monitoring is *public*. To see this, note that because subjects know the profit and the quantities chosen, they can also perceive the corresponding price; this implied price is the noisy public signal as defined by GP. Appendix C explains how both players perceive the same noisy price.[17]

iii) The Monitoring Treatment (M) [Intermediate Treatment]: Here, imperfect monitoring is removed from the IM treatment, thus only allowing for imperfect demand foresight. This treatment similar to the IM treatment, except that subjects are informed of the demand state after the round takes place, which is equivalent to informing them of

[16] A system of equations allows us to maintain the imperfect monitoring structure (available upon request).
[17] Alternatively, subjects could be directly informed of the noisy price. Such design, while more realistic for a test of the GP theory, would significantly increase the level of difficulty of the instructions. See Aoyagi and Frechette for an example of the degree of complexity in instructions when this approach is used (see also footnote 8 above for another possible caveat of such design).

their opponent's choice.[18] The motivation for adding this treatment is twofold. First, while demand uncertainty and imperfect monitoring (as assumed by GP) may be realistic sometimes, it is plausible that firms accrue information to infer rivals' past actions. Secondly, this treatment isolates one of the two factors that differentiate the FI treatment from the IM treatment.

The treatments are organized in a 2x2 matrix (Table 3). The perfect demand foresight/no monitoring treatment is unfeasible because subjects can infer the opponent's strategy. This yields a 3 (treatments) x 2 (parameterizations) experimental design.

5. Implementation

Twelve sessions with a total of 288 subjects were run. Six sessions were run with each parameterization (two sessions for each treatment). Subjects were recruited from Economics, Statistics and Management courses at the University of Massachusetts-Amherst. Demographic composition was not unusual for laboratory experiments with college students: 40% were females, 72% were white, and the combined number of freshmen and sophomores was 51% (with the remaining 49% distributed relatively evenly among juniors, seniors and graduate students). Subjects received a $5 show-up fee and earned additional money from their decisions; earnings from decisions were in experimental dollars ($1=10 experimental $). Average earnings in dollars ($) per session, as well as the corresponding dates and number of subjects are presented in table 4.

Table 3: Experimental Design

	Monitoring	**No Monitoring**
Perfect Demand Foresight	Full Information (FI)	-
Imperfect Demand Foresight	Monitoring (M)	Imperfect Monitoring (IM)

All experiments were computerized and programmed in Z-tree (Fischbacher, 1999).[19] Students were assigned a computer terminal and advised that they would be randomly paired with someone else in the room for the duration of the experiment and that communication with other participants was forbidden. Special efforts were made to achieve subjects' comprehension and familiarity with the experiment before the start of

[18] In order to keep the experimental design consistent across treatments, subjects in the IM treatment are informed of the demand state instead of their opponent's choice.

[19] Instructions and decision screens are available at: http://www.umass.edu/resec/faculty/rojas/z-tree.html

the game. Extensive instructions were coupled with 3 practice rounds, each with one demand state, and a quiz. If a subject did not respond correctly to a question, the participant was approached by the experimenter for explanation.[20]

Table 4: Number of Participants and Average Earnings per Session*

Parameterization 1						
Treatment	Full Information		Monitoring		Imperfect Monitoring	
Session Number	I	II	III	IV	V	VI
Date	04/25/08	04/25/08	04/28/08	04/30/08	04/28/08	04/28/08
# of Participants	24	24	24	24	24	24
Avg. Earnings $	25.76	27.13	31.42	30.40	24.74	26.86
Parameterization 2						
Treatment	Full Information		Monitoring		Imperfect Monitoring	
Session	VII	VIII	IX	X	XI	XII
Date	04/30/08	04/30/08	05/05/08	05/05/08	05/02/08	05/02/08
# of Participants	24	24	24	24	24	24
Avg. Earnings $	34.40	33.03	33.22	35.78	29.86	28.04

*Excludes show-up fee ($5) and earnings in risk task (see Appendix B for details of risk task)

Three colors were used for the different demand states and to distinguish own payoffs from the counterpart's payoffs; also, the three possible payoff tables were permanently displayed on the left hand side of the screen (see figures 1-4). After the quiz and the practice rounds, but before the actual rounds started, subjects were allowed to send their opponent a message from a pre-specified menu.[21] The objective of the message is twofold: a) to increase the possibility of a collusive environment (which is the focus of this paper), and b) to further enhance subjects' understanding of the experiment.[22]

[20] Only 2.7% of the questions (23 out of 864) were answered incorrectly.

[21] The menu of messages is: "Let's both play A every round"; "Let's both play B every round"; "I will always play A"; "I will always play B"; "I will play A only if you play A"; "If you play B once, I will never play A again"; "If you play B, I will not play A for some time"; "I don't want to tell you anything".

[22] In several pilot sessions we noticed that subjects spent a considerable amount of time thinking about what message to send. Also, this was the time when most questions were asked to the experimenter. These two facts were an indication that subjects were thinking intensely about how the rules of the game worked.

Figure 1: Decision Screen in the Full Information (FI) Treatment (Parameterization 1)

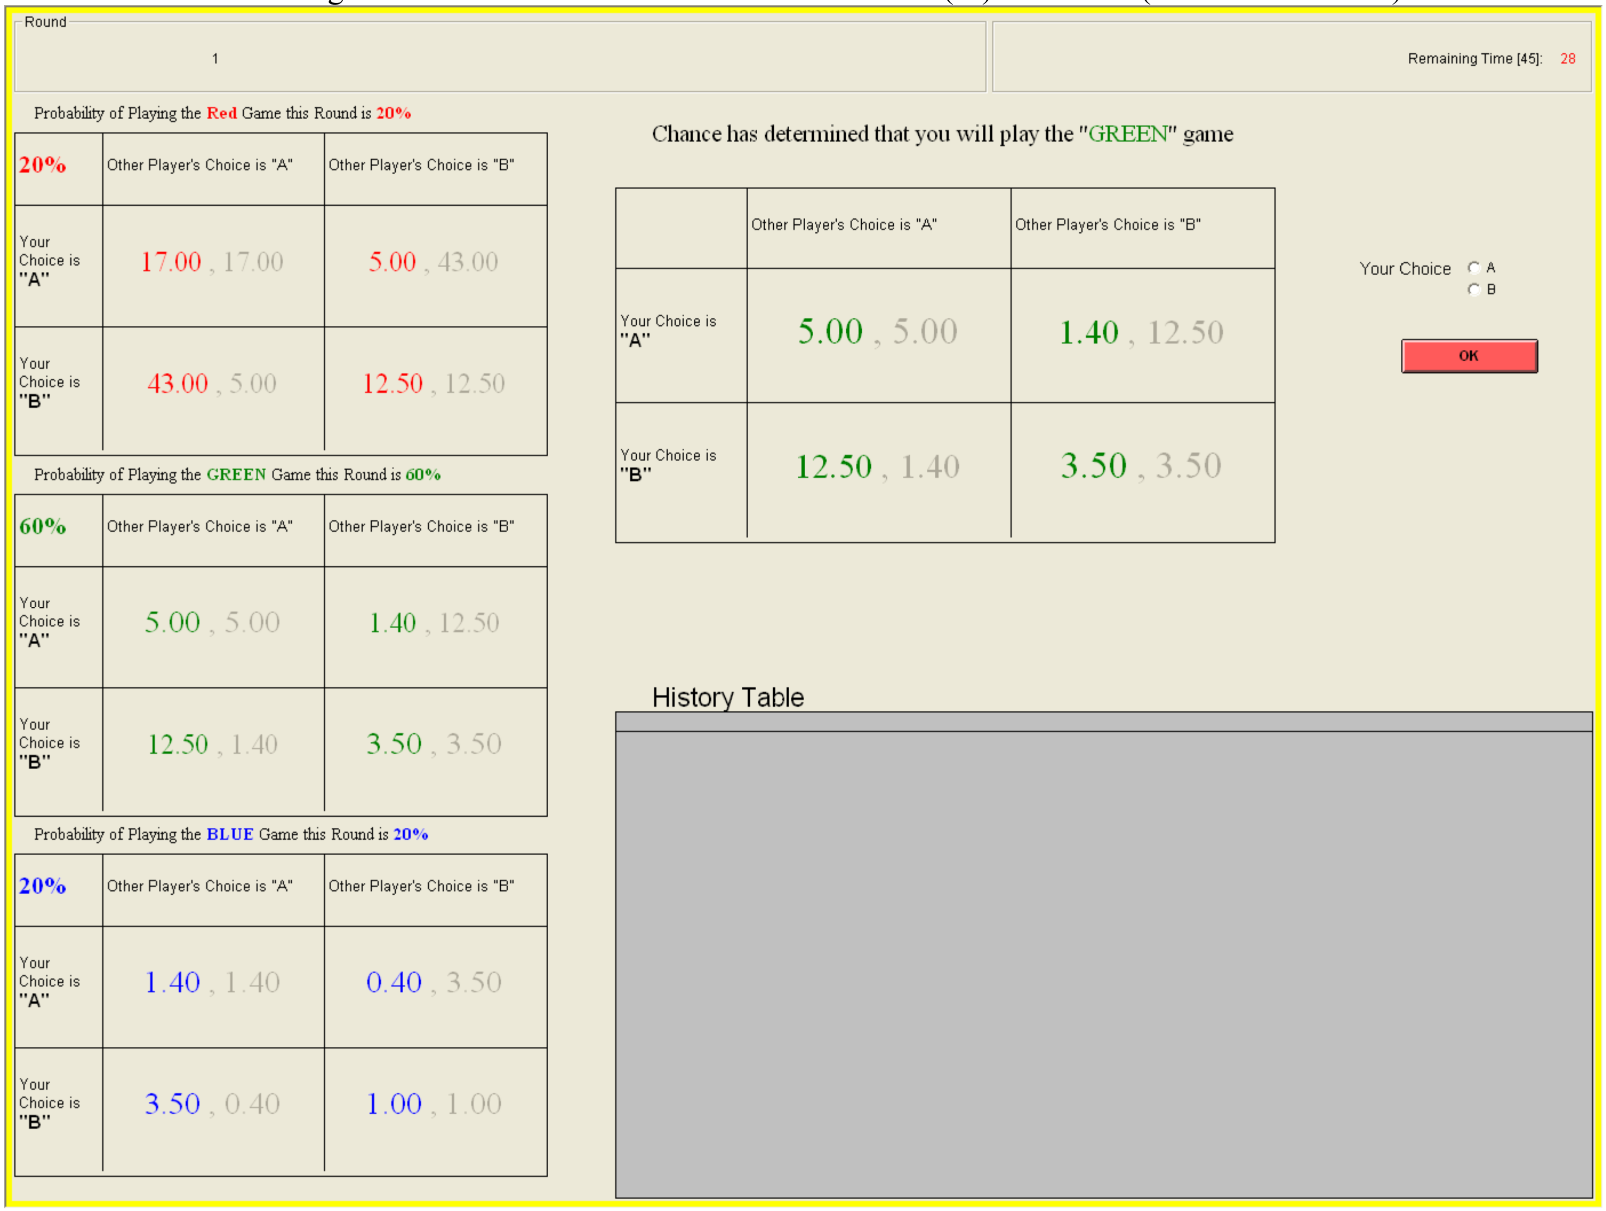

Figure 2: Decision Screen in the Monitoring (M) and Imperfect Monitoring (IM) Treatments (Parameterization 1)

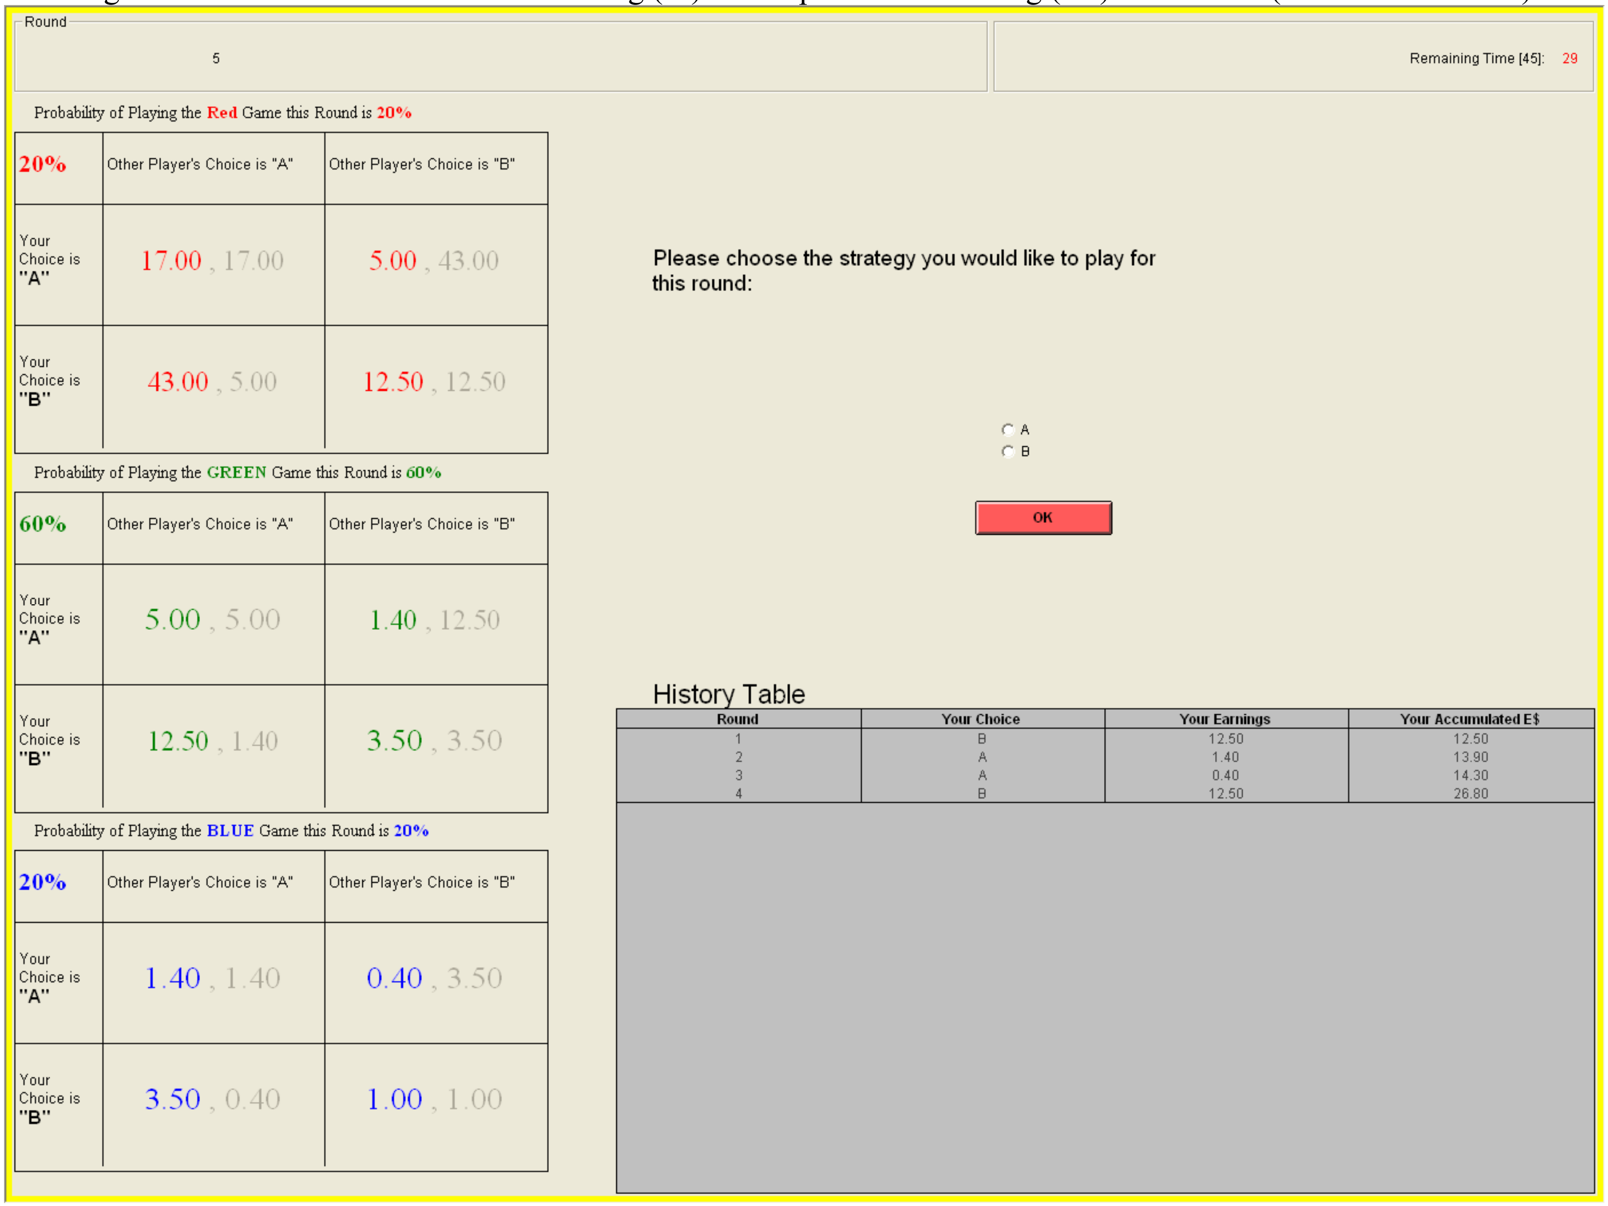

Figure 3: Profit Screen in the Full Information (FI) and Monitoring (M) Treatments (Parameterization 1)

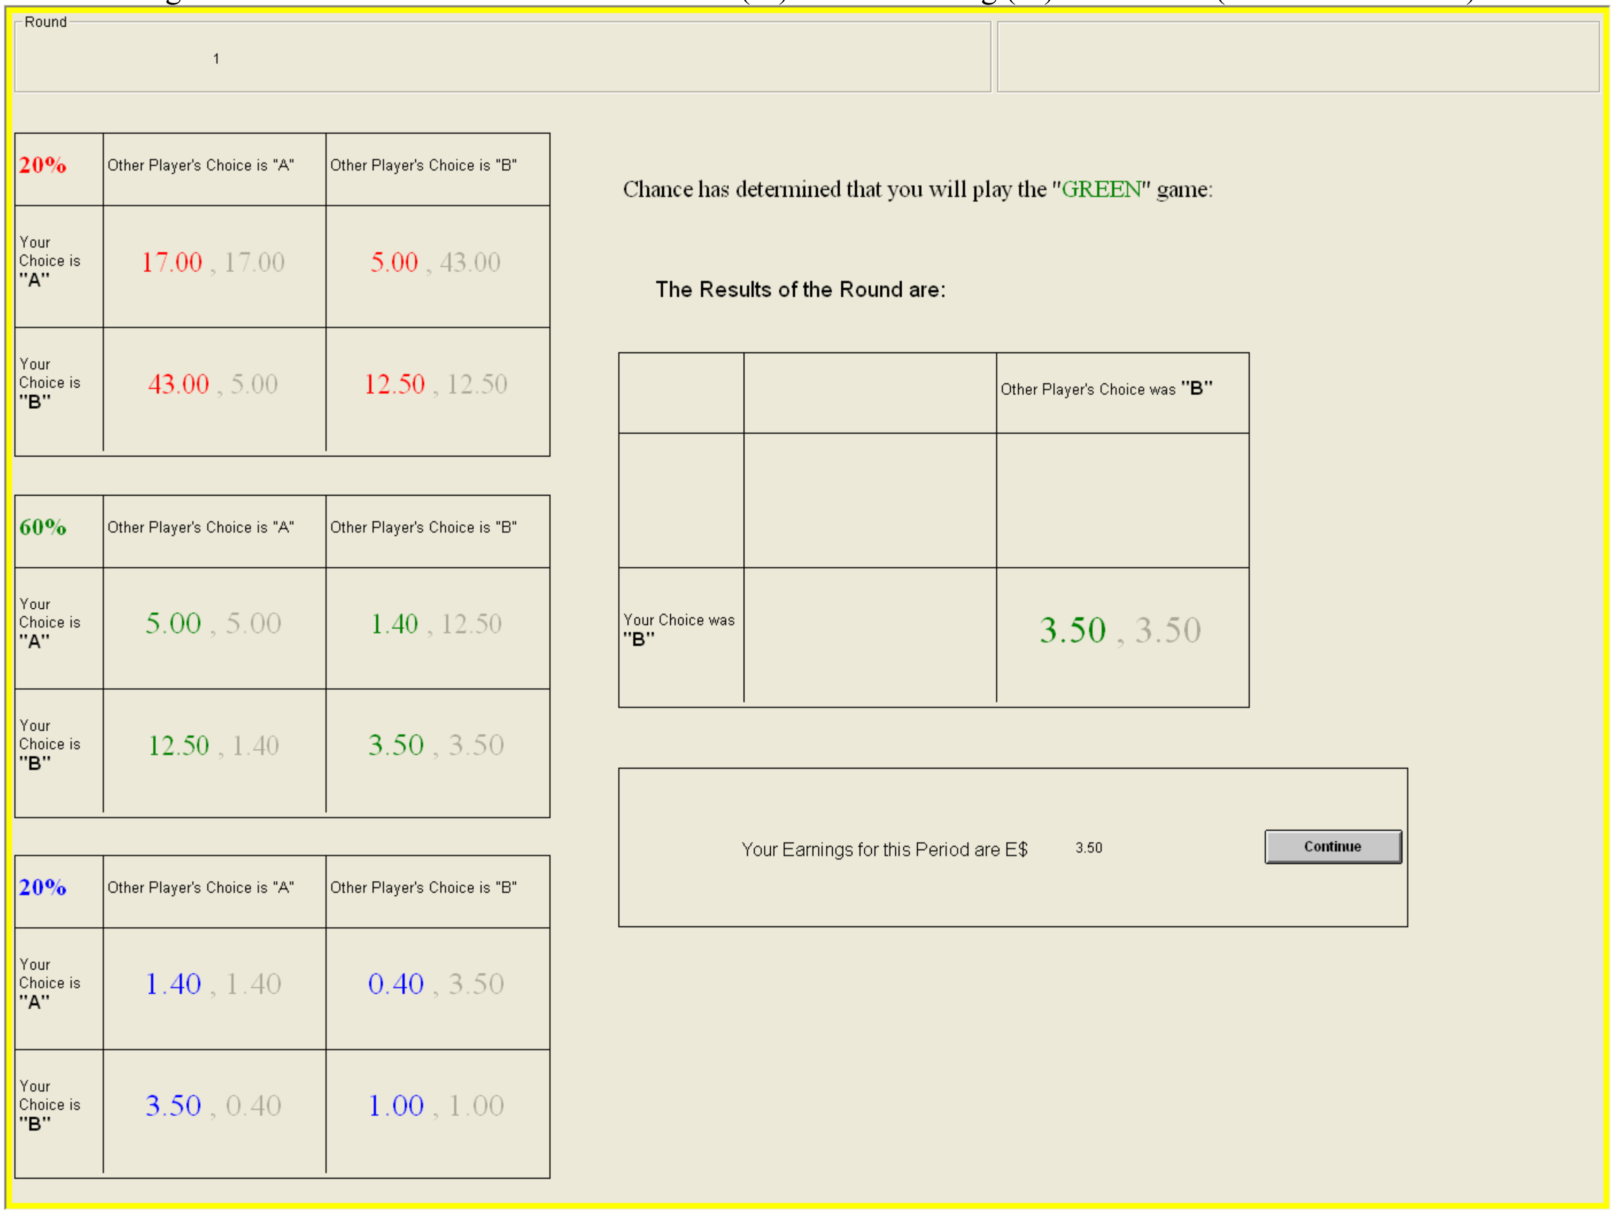

Figure 4: Profit Screen in the Imperfect Monitoring (IM) Treatment (Parameterization 1)

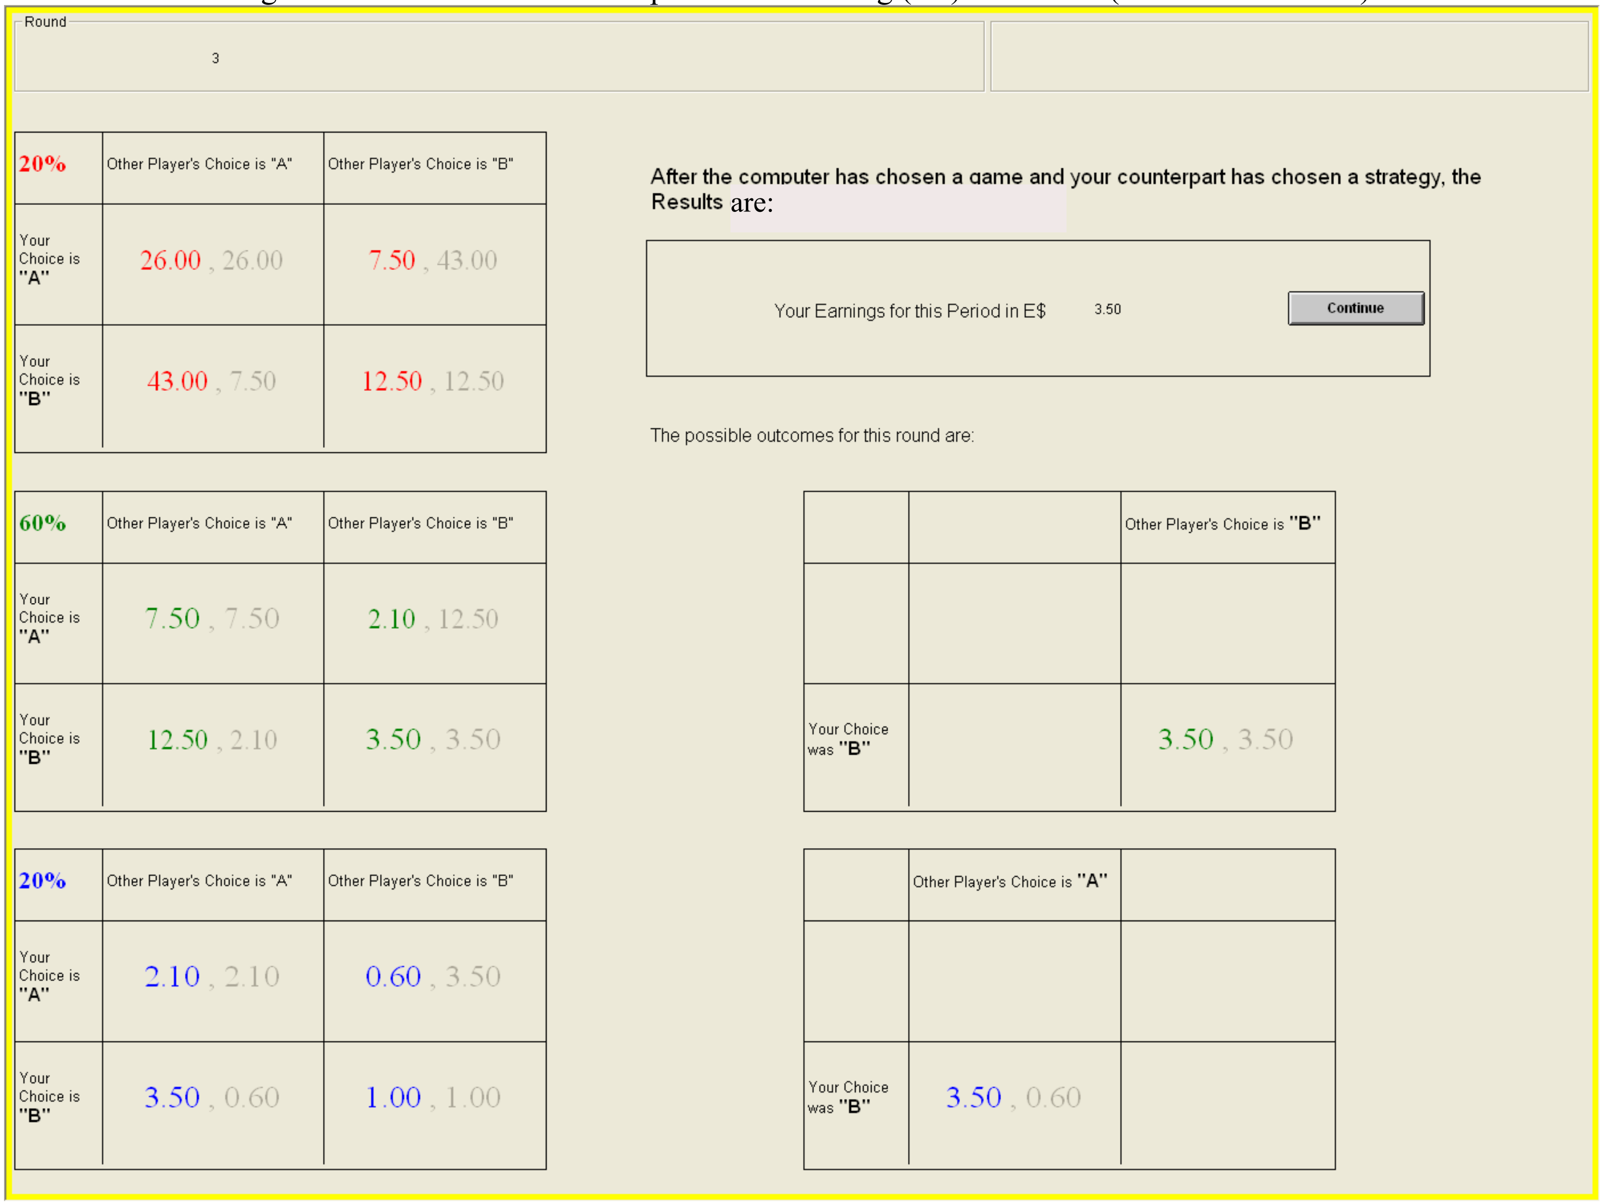

All subjects in all treatments were informed about the probability of appearance of each payoff table. Subjects played 30 rounds with certainty; after round 30, the computer terminated the game with 20% probability. To keep treatments comparable, the same draw was used to terminate the game in all treatments; the total number of periods turned out to be 33. To determine the demand state (high, medium or low), 33 random draws from a uniform distribution were taken once, and the same set of demand states implied by these draws was used in all treatments to preserve comparability.

A round consisted of subjects making a simultaneous decision between low output (*L*) and high output (*H*) (decision screen, figures 1 and 2); after a decision, subjects were informed of profits and the round ended (profit screen, figures 3 and 4). In the FI treatment, the decision screen presented subjects with the payoff matrix that they would play (figure 1). Conversely, in the IM and M treatments the decision screen only reminded subjects of the probability with which each payoff table will be chosen for play (figure 2).

In the M treatment, the profit screen reveals the chosen demand state. Also, this screen highlighted the cell in the chosen payoff table that determined the subject's profit. Because the FI and M treatments imply perfect monitoring and demand information *after* the round is played (ex-post), the profit screen for both of these treatments was the same (figure 3). In the IM treatment, the profit screen presented subjects with the *possible* outcomes that might have occurred (figure 4), effectively implementing the desired imperfect monitoring.

As depicted in the figures above, the program also contains a history table where subjects can see their cumulative earnings. In addition to the experiment on collusion, subjects completed a risk task (see Appendix B) and a small survey that contained questions on demographics, and on the subjects' assessment of the clarity of the experiment (97% of the subjects believed that the instructions were clear).[23] All sessions lasted approximately one and a half hours, including instructions. A total of 9,504

[23] Answers to the statement "The instructions for the experiment were clear and easy to follow" were: Strongly Agree (62%), Agree (35%), Neither Agree nor Disagree (0%), Disagree (1%) and Strongly Disagree (2%).

observations were collected (33 rounds x 288 participants). At the end of the experiment, subjects were individually called in private and were paid their cumulative earnings from the task in cash.

6. Results

6.1 Effect of Demand Information and Monitoring

For each of the three treatments, table 5 presents the frequency of individual cooperation (at least one player chooses the collusive outcome - *L*)[24] as well as the frequency of collusion (both players choose *L*) for the two parameterizations considered. Contrary to a stylized fact in industrial organization, when demand information is removed (from Full Information to Monitoring), cooperation and collusion increase in parameterization 1 and appear unchanged in parameterization 2. Conversely, cooperation and collusion diminish in both parameterizations when imperfect monitoring is introduced (from Monitoring to Imperfect Monitoring).

Table 5: Frequencies of Cooperation and Collusion (standard deviation)

Treatment	Parameterization	Frequency of Cooperation*	Frequency of Collusion**
Full Information	1	0.72 (0.45)	0.51 (0.50)
	2	*0.83 (0.38)*	*0.71 (0.46)*
Monitoring	1	0.76 (0.42)	0.59 (0.49)
	2	*0.84 (0.37)*	*0.71 (0.46)*
Imperfect Monitoring	1	0.63 (0.48)	0.31 (0.46)
	2	*0.66 (0.47)*	*0.41 (0.49)*

* At least one player chooses *L*. ** Both players choose *L*. # of observations in all treatments is 1,584

We test and confirm these observations using several non-parametric tests (Wilcoxon, Kolmogorov-Smirnov, Pearson's Chi-square and Epps-Singleton) as well as the parametric *t*-test: frequencies (for both cooperation and collusion) in parameterization 1 are statistically larger in the M treatment than in the FI treatment (all p-values<0.01), but frequencies from these two treatments are not statistically different from each other in parameterization 2 (p-values>0.39); frequencies (for both collusion and cooperation) in both parameterizations are statistically larger in the M and FI treatments when (individually) compared with the IM treatment (all p-values<0.01). Finally, using the

[24] Alternatively, one can define cooperation as the number of "*L*" choices (a smaller number). The results in the paper are invariant to either definition.

same battery of parametric and non-parametric tests, the level of cooperation and collusion in a given treatment is statistically larger in parameterization 2 (p-values<0.01), except for cooperation in the IM treatment (p-values tests range from 0.05 to 0.39). It important to note that the described results support theoretical predictions:

a) As noted earlier, the incentives to collude are stronger in parameterization 1, regardless of the treatment. Further, in the FI treatment, parameterization 2 implies that the left hand side of (2) is smaller than its right hand side *for all* three demand states, whereas in parameterization 1 this is true only for the medium and low demand states (deliberately, to test the RS theory); this reinforces the fact that a larger amount of collusion should be observed in parameterization 2 in the FI treatment.

b) In the M treatment, collusion is an equilibrium if the following condition (a modified version of equation (1)) is met:

$$E(\Pi_s^D - \Pi_s^{NE}) < \frac{\delta}{1-\delta}[E(\Pi_s^C - \Pi_s^{NE})]$$

$$0.2(\Pi_h^D - \Pi_h^{NE}) + 0.6(\Pi_m^D - \Pi_m^{NE}) + 0.2(\Pi_l^D - \Pi_l^{NE}) < \frac{\delta}{1-\delta} 0.2(\Pi_h^C - \Pi_h^{NE}) + 0.6(\Pi_m^C - \Pi_m^{NE}) + 0.2(\Pi_l^C - \Pi_l^{NE}) \quad (5)$$

where $E(\Pi_s^C - \Pi_s^{NE}) = 0.2(\Pi_h^C - \Pi_h^{NE}) + 0.6(\Pi_m^C - \Pi_m^{NE}) + 0.2(\Pi_l^C - \Pi_l^{NE})$. In parameterization 1, the left hand size of equation (5) is equal to 6.68 and the right hand side is equal to 15.96; this inequality is even more pronounced in parameterization 2: the left hand size of equation (5) is equal to 4.70 and the right hand side is equal to 21.90.[25] Assuming no mistakes by subjects, we should rarely observe deviations from cooperation/collusion in this treatment (especially in parameterization 2).

c) Theoretically, if the GP equilibrium were supported by the data, larger levels of collusion and cooperation should be expected in parameterization 2 as the collusive scheme predicted by GP can not be an equilibrium in parameterization 1. This only holds for collusion, however.

[25] These numbers correspond to the case of risk neutrality. After adjusting for the level of risk aversion observed in our sample (see Appendix B), the inequalities remain unchanged: 1.89 (left hand side) and 5.99 (right hand side) for parameterization 1 and 1.28 and 7.82 for parameterization 2.

Figures 5A and 5B show the frequency of cooperation throughout the 33 periods of the experiment, in both parameterizations. The figures confirm the higher level of cooperation in the FI and M treatments (with respect to the IM treatment) in both parameterizations. Also, the figures confirm the higher cooperation rate in the M treatment than in the FI treatment in parameterization 1 (5A), and the similar cooperation rates in these two treatments in parameterization 2 (5B). Figures 6A, 6B and 6C compare the frequency of cooperation across parameterizations for each of the treatments and confirm the larger cooperation in parameterization 2. The level of cooperation appears to decrease with time in all three treatments and parameterizations.[26]

A stylized fact in industrial organization is that both market factors considered should facilitate collusion. Overall, the results of this experiment suggest that this relation appears to apply only to monitoring. Lack of demand information (as modeled here) does not decrease collusion, and, conversely, may even increase it.

6.2 The Evidence for the RS theory

Descriptive Evidence

The second question we analyze is whether the predictions of the theories of interest are supported by the data. Figure 6A provides "visual" support for the RS model prediction that price wars, or breakdowns in collusion, should be observed during periods of high demand shocks: the six periods of high demand ("h", the dotted vertical lines) coincide with important drops in the frequency of cooperation in the FI treatment. Interestingly, the level of cooperation appears to return to prior levels immediately after the positive demand shock, suggesting that (some) participants did not regard defection as a trigger that caused a temporary or permanent reversion to competitive levels. On the other hand, these drops in cooperation in the FI treatment are either absent or less pronounced in parameterization 2 (see FI(2) line in figure 6B); recall that parameterization 2 implies that collusion is an equilibrium in all three demand states and hence the frequency of collusion should be approximately the same regardless of the demand state.

---

[26] Similar patterns to the ones displayed by figures 5 and 6 are present in figures that plot either the frequency of collusion or the frequency of "*L*" choices (not shown).

Figure 5: Frequency of Cooperation over 33 Periods of Stochastic Demand: h=high [---], m=medium or l=low [—]; by Parameterization (in parenthesis)

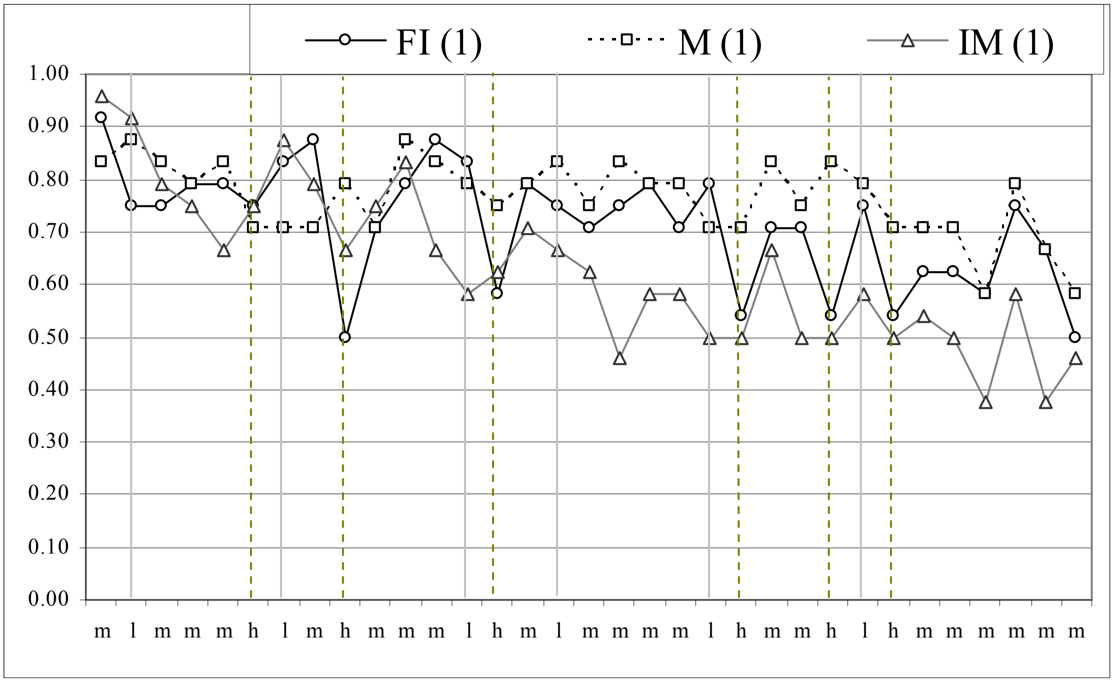

Figure 5A: Parameterization 1

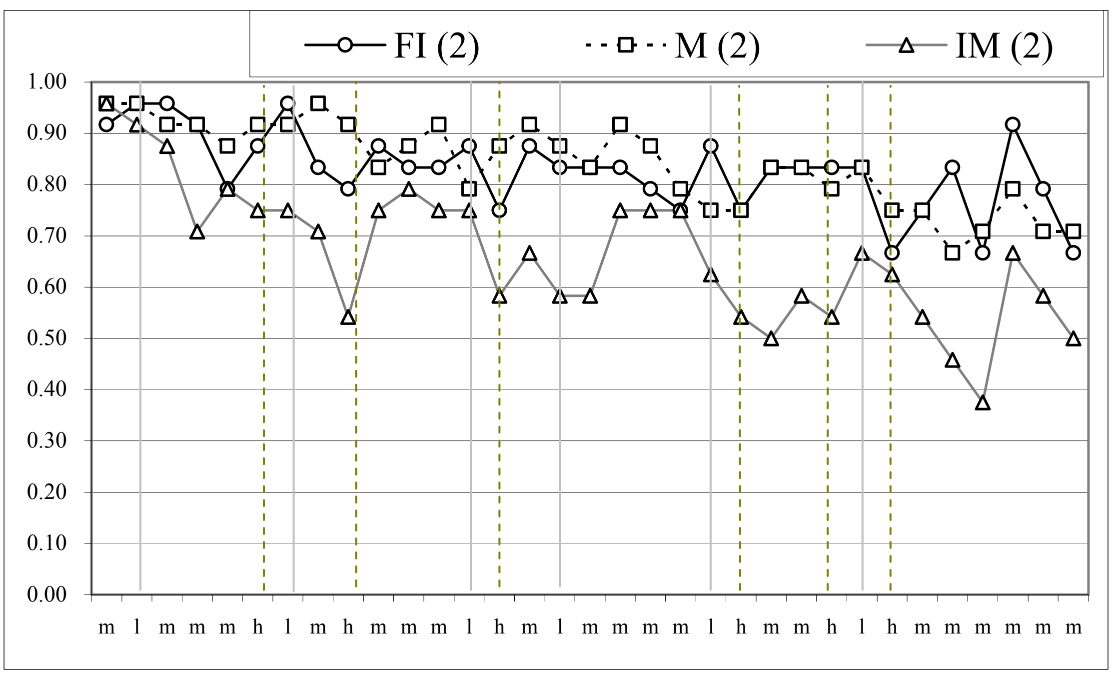

Figure 5B: Parameterization 2

Figure 6: Frequency of Cooperation over 33 Periods of Stochastic Demand: h=high [---], m=medium or l=low [—]; by Treatment (parameterization)

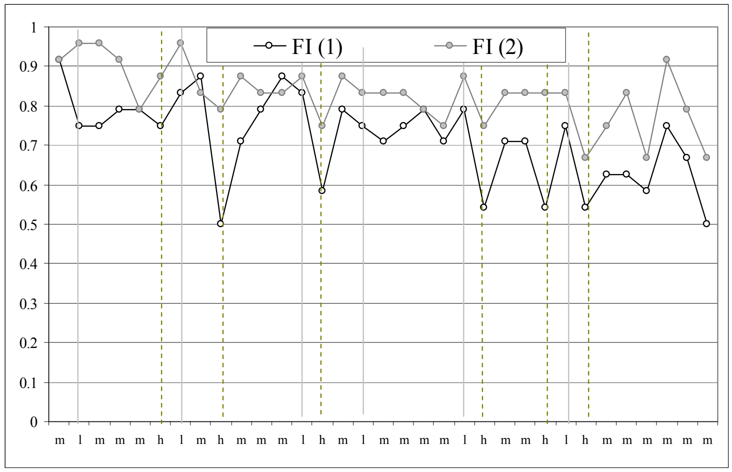

Figure 6A: Full Information Treatment

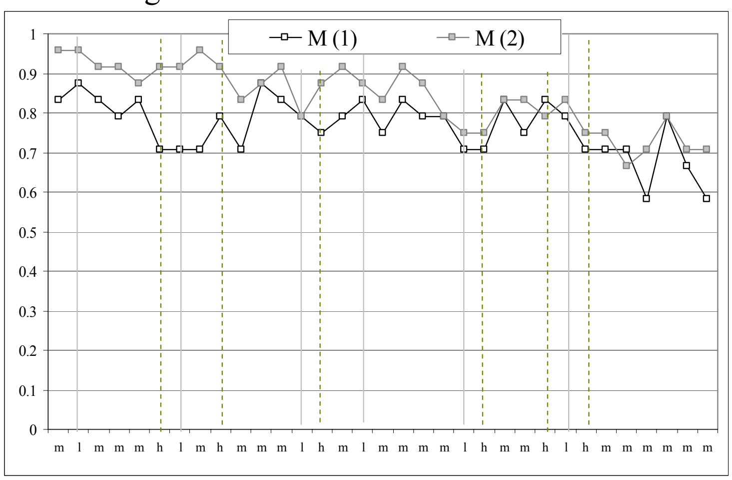

Figure 6B: Monitoring Treatment

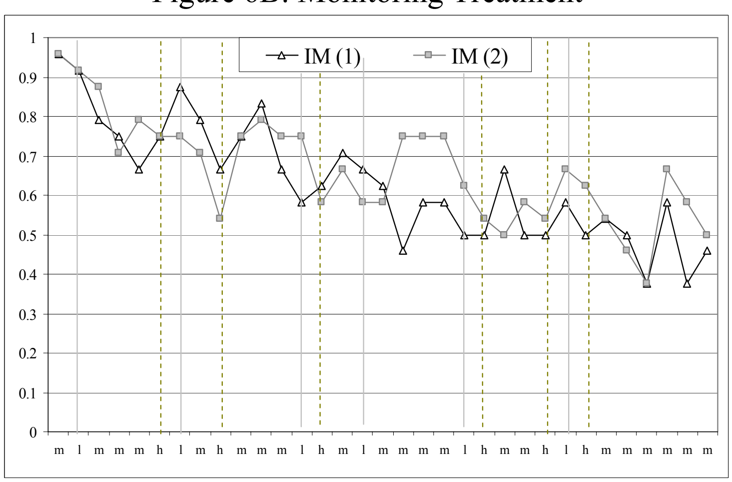

Figure 6C: Imperfect Monitoring Treatment

Table 6: Frequencies of Cooperation and Collusion in FI Treatment (St. Dev.)

Demand State	Parameterization	All Observations (Periods 1-33)			Periods 1-25		
		# Obs.	**Freq. Coop.***	**Freq. Collusion****	**# Obs.**	**Freq. Coop.***	**Freq. Collusion****
High	1	288	0.58 (0.49)	0.42 (0.49)	240	0.58 (0.49)	0.43 (0.50)
(*h*)	*2*	*288*	*0.78 (0.42)*	*0.65 (0.48)*	*240*	*0.80 (0.40)*	*0.67 (0.47)*
Medium	1	1,008	0.73 (0.44)	0.52 (0.50)	720	0.78 (0.42)	0.56 (0.50)
(*m*)	*2*	*1,008*	*0.83 (0.38)*	*0.70 (0.46)*	*720*	*0.85 (0.36)*	*0.73 (0.44)*
Low	1	288	0.78 (0.41)	0.57 (0.49)	240	0.79 (0.41)	0.59 (0.49)
(*l*)	*2*	*288*	*0.89 (0.31)*	*0.76 (0.43)*	*240*	*0.90 (0.30)*	*0.77 (0.42)*

* Frequency of either player choosing *L*. ** Frequency of *both* players choosing *L*.

Table 6 summarizes the frequency of cooperation and collusion in the FI treatment under both parameterizations. To analyze any potential end-of-game effect, table 6 also presents the frequencies in all 33 rounds as well as frequencies in rounds 1-25. There appears to be a small end-of-game effect, as lower frequencies of cooperation and collusion appear to be slightly larger (in both parameterizations) when only the first 25 rounds are considered; however, the only statistically significant difference appears in the medium demand state (parameterization 1), where the difference in frequencies (for both collusion and cooperation) between the two period lengths (33 vs. 25) is somewhat significant (p-values range from 0.04 to 0.47).[27]

Focusing on rounds 1-25[28], our battery of tests reveal that the frequency of cooperation and collusion in parameterization 1 is significantly smaller in the high demand state when compared with either the medium or low demand states (all p-values < 0.01), whereas cooperation and collusion are not statistically different between medium and low demand states (p-values>0.37). Turning to parameterization 2, it is clear that the difference in collusion and cooperation frequencies between the high demand state and the other two states is much smaller than in parameterization 1; this is reflected by a weaker statistical difference (p-values range from 0.06 to 0.79). A similarly small and

[27] A similar conclusion is reached if the end-of-game effect is analyzed with other cut-off periods (24, 26).
[28] We report the results of rounds 1-25 because of the evidence found of an end-of-game effect for the medium demand period. The conclusion in this paragraph, however, is similar if all rounds are used.

weakly significant difference in cooperation and collusion frequencies exists between the medium and low demand states (p-values range from 0.04 to 0.97). Taken together, these results provide some informal evidence for the predictions of the RS theory.

Subjects' Strategies

More formally, we test how closely subjects' strategies are in accordance to the predictions of the RS theory by analyzing individual choices in a limited dependent variable model. Theoretically, only parameterization 1 yields the predictions of the RS theory (price wars during high demand); hence the analysis that follows uses data from this parameterization. Specifically, we consider various possible *individual* strategies and use standard likelihood ratio tests to analyze how well the RS theory can explain the observed choices. The decision to cooperate of individual i in period t is given by:

$$y_{it} = 1\left[\gamma z_{it} + c_i + e_{it} > 0\right]$$

where y_{it} denotes cooperation (1) or defection (0), $1[\cdot]$ is an indicator function, z_{it} is a vector of explanatory variables, $c_i \mid z_{it} \sim N(\alpha + \psi\zeta_i, \sigma_c^2)$ is an unobserved and random heterogeneity term, and $e_{it} \sim N(0,1)$ is an idiosyncratic error term. We assume a non-zero mean for c_i to deal with the initial conditions problem (Wooldridge, 2002, p. 495) and set ζ_i equal to y_{i0}.[29] The vector of parameters γ as well as α, ψ and σ_c^2 are estimated from the data. The vector z_{it} contains different variables depending on which strategy is being tested. Table 7 displays the different strategies considered and presents the recursive definition (third column) of the corresponding variable (s_{it}) that enters z_{it}. Note that this specification accommodates subjects' mistakes (through e_{it}) as well as heterogeneity (through c_i); for example if $z_{it} = s_t^{RS}$ and $e_{it} = c_i = 0$, one would have a deterministic rule: whenever demand is high one would observe defection, whereas collusion would be observed otherwise (assuming γ_1 is positive).

[29] Other specifications that set ζ_i equal to the mean or the median of y_{it} over the 33 periods generated qualitatively similar results.

Note that all these strategies can be defined as trigger strategies, each with three characteristics: a) the type of trigger, b) the duration of reversion to the non-cooperative outcome, and c) the rule for returning to the collusive outcome. It is important to note, however, that while s_t^{RS} is a trigger-like strategy, it tests whether subjects are on the *equilibrium path* predicted by RS; that is, the trigger strategy upon which the RS prediction is based (i.e. reversion to the NE forever if deviation occurs in the medium and low demand states) is not observed. Put differently, s_t^{RS} tests whether the equilibrium outcome predicted by RS occurs, whereas s_{it}^{∞}, for example, tests whether subjects use a strategy that is *consistent* with the conditions needed to obtain the RS equilibrium (i.e. the grim-trigger strategy). In a sense, then, s_t^{RS} entails a demanding test of the RS theory.

Table 7: Strategies Considered and Corresponding Variables (FI treatment)

Strategy	γz_{it}	**Definition**
Random	N/A	N/A
RS	$\gamma_1 s_t^{RS}$	$s_t^{RS} = 1[demand\ state = medium\ or\ low]$
Tit-for-Tat (TT)	$\gamma_2 s_{it}^{tt}$	$s_{it}^{tt} = y_{jt-1}$
Punishment Strategy (P-*N*)	$\gamma_3 s_{it}^{N}$	For $t=1$: $s_{i1}^{N}=1$ For $t>1$: $s_{it}^{N} = \begin{cases} 1 & \text{if } s_{it-1}^{N}=1 \text{ and } (y_{it-1}=1 \text{ and } y_{jt-1}=1) \\ & \text{or if } s_{it-1}^{N}=0 \text{ and } s_{it-(N+1)}^{N}=1 \\ 0 & \text{otherwise} \end{cases}$

Notes: N = punishment length; y_i = subject i's choice

Note that s_t^{RS} is determined exclusively by the *current* demand state; s_{it}^{tt} depends exclusively on the opponent's choice *last* period, whereas punishment strategies depend on both own (*i*) and opponent's (*j*) choices. Given the wide range of punishment lengths that can sustain collusion (Appendix A, table A.1), we consider seven punishment strategies (P-*N*), six finite ($N = 1,...,6$) and the grim strategy $N = \infty$. The tit-for-tat strategy, is not related to theoretical predictions we have in mind, but it is considered for its reported predictive power (see, for example, Engle-Warnick and Slonim, 2006).

Table 8 reports the results of probit regressions of the strategies considered excluding the last 8 rounds of the game (results are robust if all rounds are included).

While searching for patterns that may explain subjects' strategies better than the RS theory alone, informal inspection of the data revealed that subjects may appear to be basing their strategies on both the demand state (high or not) as well as on their opponent's choice; the last column presents the estimates of this "combined" strategy.

In all specifications, the statistical significance of $\rho = \sigma_c^2 / (\sigma_c^2 + 1)$ does not reject the random effects specification, indicating that heterogeneity is important. The positive and significant estimate of ψ indicates that a subject that cooperates in period 1 is more likely to do so later on. The log-likelihood reveals that (besides the "combined" RS + tt strategy) the best fit is given by the infinite punishment strategy (P-∞), while the second best fit is given by the RS strategy, followed by the tit-for-tat strategy. The finite punishment strategies tend to increase in predictive power with the punishment length; importantly, the fit of the model appears to converge to that of P-∞ as N grows beyond 6 periods (results not reported).

The table also presents pair-wise likelihood ratio tests between each of the strategies and the random strategy. The random strategy can not be rejected in favor of P-1 and can be rejected only at the 10% level in favor of P-3. Consistent with the fit of each strategy, the random strategy can be rejected in favor of RS, TT, P-∞ and the P-2, P-4, P-5 and P-6 strategies.[30] The single most important variable explaining behavior is P-∞. The combined RS+tt strategy has a better fit than the P-∞ strategy and is also preferred over the RS strategy alone (p-value of LR-test < 0.01) and the tt strategy alone (p-value < 0.01). Considering the fact that s_t^{RS} is a demanding test of the RS theory, and that most of the explanatory power in the RS+tt strategy appears to be attributable to the s_t^{RS} variable, the results are interpreted as supportive of the RS theory.

[30] Because the feasible equilibria are state dependent (see Appendix A), it may make more sense to search for three different punishment lengths; that is, the N with the highest explanatory power may depend on which demand state a punishment phase starts. We carried out such an estimation and, for all three demand states, we found a similar pattern to the one observed in table 8: coefficients of strategies with short punishment length have weak (or no) statistical significance and increase in size and significance with N. This indicated that there is a unique N that minimizes subject's mistakes in all three demand states: $N = \infty$ (i.e. the P-∞ specification in table 8).

Table 8: Probit Estimates of Different Strategies in the FI treatment, Parameterization 1, Rounds 1-25

Parameter	**Random**	**RS**	**tt**	**P-1**	**P-2**	**P-3**	**P-4**	**P-5**	**P-6**	**P-** ∞	**RS + tt**
α	-0.80***	-0.66	-0.97*	-0.88**	-0.93**	-0.85**	-0.90**	-0.88**	-0.86**	-0.69*	-0.86**
	(0.43)	(0.46)	(0.36)	(0.42)	(0.39)	(0.39)	(0.36)	(0.36)	(0.35)	(0.16)	(0.38)
γ_1		0.92*									0.99*
		(0.14)									(0.14)
γ_2			0.56*								0.68*
			(0.12)								(0.12)
γ_3				0.17	0.37*	0.23**	0.51*	0.49*	0.53*	2.39*	
				(0.11)	(0.11)	(0.12)	(0.12)	(0.13)	(0.14)	(0.24)	
ψ	2.40*	2.56*	2.03*	2.31*	2.20*	2.20*	2.02*	2.03*	1.98*	0.69*	2.14*
	(0.52)	(0.53)	(0.43)	(0.48)	(0.46)	(0.45)	(0.41)	(0.42)	(0.41)	(0.23)	(0.44)
ρ	0.69*	0.72*	0.60*	0.67*	0.65*	0.65*	0.60*	0.61*	0.59*	0.69*	0.63*
LL	-450.84	-427.84	-440.41	-449.80	-445.30	-449.17	-443.37	-444.33	-444.16	-422.07	-413.77
LR Test	N/A	46.00	20.85	2.07	11.08	3.34	14.94	13.01	13.35	57.53	74.13
(p-value)†		(<0.01)	(<0.01)	(0.15)	(<0.01)	(0.07)	(<0.01)	(<0.01)	(<0.01)	(<0.01)	(<0.01)

* Significant at 1%. ** Significant at 5%. *** Significant at 10%. † Likelihood ratio test with respect to the random strategy

Notes: # of Observations = 1,152 in all models to keep number of observations comparable across strategies (first period is lost in TT strategy). RS=Rotemberg and Saloner equilibrium strategy, tt=Tit-for-Tat, P-*N*=punishment for *N* periods, P-∞= infinite punishment. LL=Log-likelihood. Standard errors in parentheses.

Equilibrium Outcomes

As opposed to the econometric model presented above, the focus of the analysis here is on *equilibrium outcomes* rather than on individual strategies. Specifically, we analyze how the data lends support to the different feasible equilibria presented in Appendix A (table A.1). Table 9 displays the frequencies of the outcomes observed in each of the three demand states; the bold numbers indicate that the cell is a feasible equilibria. There are several patterns worth noting. First, collusion (*L*,*L*) is the most frequently observed outcome, except when theory predicts it is not an equilibrium (high demand, parameterization 1). Second, (*H*,*L*)/ (*L*,*H*) is the least frequently observed outcome, except in one case (low demand, parameterization 2). Third, within a parameterization, collusion appears a more likely outcome during "bad times" (i.e. its frequency decreases as demand becomes larger), whereas the one-shot NE becomes more likely during "good times"; the frequency of the (*H*,*L*)/(*L*,*H*) equilbria, on the other hand, is relatively stable within a parameterization.

Table 9: Frequencies of Observed Outcomes

Demand State (outcomes)		Parameterization 1	Parameterization 2
High (*h*)	(*L*,*L*)	41.67%	**65.28%**
	(*H*,*H*)	**42.36%**	**22.22%**
	(*H*,*L*)/(*L*,*H*)	**15.97%**	**12.50%**
Medium (*m*)	(*L*,*L*)	**51.79%**	**70.44%**
	(*H*,*H*)	**26.59%**	**17.46%**
	(*H*,*L*)/(*L*,*H*)	**21.63%**	**12.10%**
Low (*l*)	(*L*,*L*)	**57.64%**	**75.69%**
	(*H*,*H*)	**21.53%**	**11.11%**
	(*H*,*L*)/(*L*,*H*)	**20.83%**	**13.19%**

Notes: Bold numbers indicate that entry is a feasible equilibrium (see Appendix A for details)

To contrast the predictive power of the RS equilibrium with that of other feasible equilibria, we conduct a simple test. First, we create an indicator variable that takes a value of 1 if the observed outcome coincides with that predicted by a given equilibria. Then, we compute the fraction of "correctly" predicted outcomes for each of the equilibria considered and rank the equilibria according to this fitness measure. The equilibria that we consider are: a) the one shot NE (*H*,*H*) in all periods, b) the collusive outcome (*L*,*L*) in all periods, c) the RS prediction (*H*,*H* during high demand and *L*,*L* otherwise), d) and the (*H*,*L*)/(*L*,*H*) outcome in all periods.[31,32] The best fit in parameterization 1 is given by the RS equilibrium (54%) followed by the "always collude" outcome (51%), the "always defect" outcome (28.54%), and the (*H*,*L*)/(*L*,*H*) outcome (20.45%). The best fit in parameterization 2 is given by the "always collude" outcome (70.45%), followed by the RS equilibrium (64.90%), the "always defect" outcome (17.17%), and the (*H*,*L*)/(*L*,*H*) outcome (12.37%). Thus, the RS equilibrium has stronger support when expected (parameterization 1) thereby providing additional support for the RS predictions; again, however, this support is not overwhelming.

6.3 The Evidence for the GP theory

Figure 6C (IM treatment) shows a decline in cooperation after the occurrence of a low demand period in four out of six cases in parameterization 1, and in all but one case in parameterization 2. This evidence may suggest that low demand triggers defection (i.e. a price war). While there is no discernable regime switching pattern in either parameterization, as predicted by GP, cooperation tends to have a downward trend after a low demand realization and an upward trend after a high demand realization, especially in earlier rounds. As shown by table 5 and figure 6C, the frequency of cooperation and collusion in the IM treatment is the smallest of all treatments which may suggest that the GP environment may hinder collusive behavior.

In search of further evidence for the GP theory, we construct an indicator variable that takes a value of 0 for price war and 1 otherwise. To determine whether a price war has started, we use the GP trigger price p_2 (see Appendix A for details) as the relevant threshold in this exercise: a price war is triggered if a price of p_2 or lower is observed. We construct several indicator variables, one for each possible price war duration (from one period to ∞). Finally, we compute the frequencies of cooperation and collusion when the

[31] Note that (*L*,*L*) in all periods is only an equilibrium in parameterization 2. Similarly, the RS outcome is only an equilibrium in parameterization 1; for completeness of the tests, however, we consider these two outcomes as possible explanations of behavior even if they are not a feasible theoretical equilibrium.
[32] Our indicator variable of the (*H*,*L*)/(*H*,*L*) equilibria takes a value of 1 if either outcome is observed, and zero otherwise; alternatively, one can construct two different indicators, one for each equilibria, but in our case this is not relevant as our conclusions remain unchanged.

indicator variable predicts collusion (C) and compare them with the frequencies of cooperation and collusion when the indicator variable predicts a price war (R).

Table 10 reports the cooperation and collusion frequencies during the two regimes, as predicted by the different lengths of punishment after price drops to p_2 (or below). Cooperation and collusion frequencies are almost always statistically different (with the frequencies in the collusive regime (C) always larger); this difference tends to increase in size with the punishment length (to conserve space we only display a few punishment lengths here). As shown in appendix A, in parameterization 1 the GP equilibrium is not feasible (even with a permanent price war $N=\infty$); conversely, parameterization 2 can support price wars that range from 3 to ∞. Interestingly, frequencies appear similar across parameterizations, even though the parameterizations imply different incentives. Overall, the evidence in table 10 in support of the GP equilibrium appears to be mixed; we would expect no difference between "C" and "R" frequencies for parameterization 1, but there are several cases when this does not occur. Also, it is not clear why finite price wars (as suggested by GP) appear less likely than price wars of infinite duration. We next turn to a more exhaustive analysis that compares the strategies and outcomes predicted by GP with those predicted by other alternatives.

Table 10: Frequencies of Cooperation and Collusion in Collusive (C) and Reversionary (R) Regimes in IM treatment, Various Punishment Lengths, Rounds 1-25.

Punishment Length (N)	**Cooperation**						**Collusion**					
	Parameterization 1			**Parameterization 2**			**Parameterization 1**			**Parameterization 2**		
	R	C	*p-value**	R	C	*p-value**	R	C	*p-value**	R	C	*p-value**
2	0.66	0.69	*0.27*	0.65	0.74	*<0.01*	0.30	0.41	*<0.01*	0.42	0.48	*0.02*
3	0.67	0.68	*0.83*	**0.65**	**0.74**	***<0.01***	0.33	0.39	*0.02*	**0.43**	**0.48**	***0.05***
4	0.67	0.68	*0.83*	**0.66**	**0.78**	***<0.01***	0.33	0.44	*<0.01*	**0.42**	**0.51**	***<0.01***
16	0.67	0.71	*0.24*	**0.68**	**0.80**	***<0.01***	0.34	0.46	*<0.01*	**0.44**	**0.54**	***<0.01***
17	0.66	0.74	*0.04*	**0.68**	**0.81**	***<0.01***	0.34	0.49	*<0.01*	**0.43**	**0.59**	***<0.01***
18	0.66	0.79	*<0.01*	**0.68**	**0.83**	***<0.01***	0.34	0.57	*<0.01*	**0.43**	**0.63**	***<0.01***
∞	0.65	0.94	*<0.01*	**0.68**	**0.94**	***<0.01***	0.33	0.73	*<0.01*	**0.43**	**0.73**	***<0.01***

Note: Bold numbers indicate that the entry entails a feasible punishment length in the GP equilibrium. The results are qualitatively similar if all rounds (1-33) are considered; p_2 is assumed to be the price trigger.
* Pearson's Chi-Square statistic; p-values of other non-parametric tests (Wilcoxon, Kolmogorov-Smirnov, and Epps-Singleton) and the parametric t-test produce similar p-values.

Subjects' Strategies

Recall that the basis for the GP equilibrium is the presence of trigger strategies; thus, we investigate how the predictive power of individual trigger strategies as predicted by the GP equilibrium compares to the predictive power of other plausible trigger strategies. As in the previous section, we recursively define a variety of plausible trigger strategies (s_{it}) and use a probit model that allows for errors in subjects' decisions. The strategies considered, summarized in table 11, are more complex than in the FI treatment; this is because the opponent's choice here is observed with noise. To be consistent with the public nature of monitoring assumed by GP, the relevant noisy signal in these strategies is the price "implied" by our design (see Appendix C).[33]

Two strategies can be considered as tests of the GP theory. The GP strategy, which corresponds to the regimes used in table 10 (1=collusive, 0=competitive), can be considered as a more demanding test of the GP theory as it tests whether individual strategies are consistent with the GP *equilibrium path*. The trigger strategy that depends on the public signal only (*T1*) can be considered as a less demanding test of the GP theory as it only tests whether defection is triggered by the threshold level predicted by theory (p_2 in parameterization 2, see Appendix A).

Except for the GP strategy, all strategies are subject-specific; these strategies differ in whether their transition rules (from the collusive regime to the competitive regime -"down"- and from the competitive regime to the collusive regime –"up") depend on the public signal (p_{t-1}), the own action (y_{it-1} in *TT1* and *TT2*) and whether there is a singe threshold (*k*) (i.e. return to the collusive regime is determined by a given number of punishment periods -*T1* and *TT1*) or two thresholds (*T2* and *TT2*). The reason for defining strategies to be a function of the own action (*TT1* and *TT2*) is because the distribution of prices is a function of the own action.

[33] In a separate analysis of trigger strategies that consider observed profit (not price) to be the relevant noisy signal we obtained qualitatively similar conclusions; we omit these results to conserve space.

Table 11: Strategies Considered and Corresponding Variables, IM treatment

Strategy	γz_{it}	**Definition**
Random	N/A	N/A
GP, with: $N = 1,\ldots,\infty$	$\gamma_1 s_t^{GP_N}$	For $t=1$: $s_1^{GP_N}=1$ $s_t^{GP_L} = \begin{cases} 1 & \text{if } s_{t-1}^{GP_N}=1 \text{ and demand=high or medium} \\ & \text{or } s_{t-1}^{GP_N}=0 \text{ and } s_{t-(N+1)}^{GP_N}=1 \\ 0 & \text{otherwise} \end{cases}$
One Threshold Strategy (T1), with: k =f(p) $N = 1,\ldots,\infty$	$\gamma_2 s_{it}^{T1_N}$	For $t=1$: $s_{i1}^{T1_N}=1$ For $t>1$ $s_{it}^{T1_N} = \begin{cases} 1 & \text{if } s_{it-1}^{T1_N}=1 \text{ and } p_{t-1}>k \\ & \text{or } s_{it-1}^{T1_N}=0 \text{ and } s_{t-(N+1)}^{T1_N}=1 \\ 0 & \text{otherwise} \end{cases}$
One Threshold Strategy (TT1), with: k=f(p , y_i) N= 1,…, ∞	$\gamma_3 s_{it}^{TT1_N}$	For $t=1$: $s_{i1}^{TT1_N}=1$ For $t>1$ $s_{it}^{TT1_N} = \begin{cases} 1 & \text{if } s_{it-1}^{TT1_N}=1 \text{ and if:} \begin{cases} y_{it\text{-}1}=1\text{: and } p_{t-1}>k(L) \\ \text{or } y_{it\text{-}1}=0\text{: and } p_{t-1}>k(H) \end{cases} \\ & \text{or if } s_{it-1}^{TT1_N}=0 \text{ and } s_{it-(N+1)}^{TT1_N}=1 \\ 0 & \text{otherwise} \end{cases}$
Two-Threshold Strategy (T2), with: k =f(p)	$\gamma_4 s_{it}^{T2}$	For $t=1$: $s_{i1}^{T2}=1$ For $t>1$: $s_{it}^{T2} = \begin{cases} 1 & \text{if } s_{it-1}^{T2}=1 \text{ and if } p_{t-1}>k^{down} \\ & \text{or if } s_{it-1}^{T2}=0 \text{ and if } p_{t-1}>k^{up} \\ 0 & \text{otherwise} \end{cases}$
Two-Threshold Strategy (TT2) with: k=f(p , y_i)	$\gamma_5 s_{it}^{TT2}$	For $t=1$: $s_{i1}^{TT2}=1$ For $t>1$: $s_{it}^{TT2} = \begin{cases} 1 & \text{if } s_{it-1}^{TT2}=1 \text{ and if:} \begin{cases} y_{it\text{-}1}=1\text{: and } p_{t-1}>k^{down}(L) \\ \text{or } y_{it\text{-}1}=0\text{: and } p_{t-1}>k^{down}(H) \end{cases} \\ & \text{or if } s_{it-1}^{TT2}=0 \text{ and if:} \begin{cases} y_{it\text{-}1}=1\text{: and } p_{t-1}>k^{up}(L) \\ \text{or } y_{it\text{-}1}=0\text{: and } p_{t-1}>k^{up}(H) \end{cases} \\ 0 & \text{otherwise} \end{cases}$

N = punishment length; p = price $\in (p_0, p_1, p_2, p_3, p_4)$; k = threshold p; y = choice; L=low quantity; H=high quantity

These strategies allow several possibilities. For example, if a subject cooperated in *t-1*, the threshold price for reverting to the competitive regime may be different than if the subject did not cooperate ($k_{TT1}(H) \neq k_{TT1}(L)$ or $k_{TT2}^{down}(H) \neq k_{TT2}^{down}(L)$). Also, reversion to collusion may depend on a threshold ($k_{T2}^{up}, k_{TT2}^{up}(y_i)$, $y_i \in \{H, L\}$) rather than occurring after a given period of time. Finally, the threshold level for the transition to competition may be different than the threshold level for the transition to collusion (i.e. $k_{T2}^{up} \overset{\neq}{=} k_{T2}^{down}$; $k_{TT2}^{down}(y_i) \overset{\neq}{=} k_{TT2}^{up}(y_i)$).

We investigated different punishment lengths (*N=1,…,*∞) and all possible threshold values. Table 12 presents estimates (for both parameterizations) of the strategies that had the highest explanatory power. In both parameterizations, the random strategy is rejected in favor of a variety of strategies (including GP strategies).[34] There are several patterns worth noting. First, the grim-trigger strategy ($N=\infty$) has the largest explanatory power in both parameterizations as indicated by its substantially larger LL value; this is true whether defection is triggered by a low demand realization (GP strategy) or by a low threshold level (*T1*, *TT1*). Second, the most likely trigger for reverting to defection seems to be p_1, regardless of the strategy or parameterization considered; this threshold is lower than that predicted by GP for parameterization 2 (p_2).

Third, the random strategy is rejected in favor of several GP strategies of various punishment lengths (in both parameterizations); recall that the GP equilibrium is not feasible in parameterization 1 and is feasible for punishment lengths that range from 3 to ∞ in parameterization 2. Lastly, behavior appears to be very similar in both parameterizations, as the strategies with the highest explanatory power are almost identical.

[34] Periods 1-25 are included in the estimation. Conclusions are qualitatively similar if all periods are included.

Table 12: Probit Estimates of Different Strategies in the IM treatment, Rounds 1-25

Parameterization 1

	Random	GP_N			$T1_N$ $\{N\}$			$TT1_N$ $\{N\}$	$T2$		$TT2$	
		N=12[1]	N=13	N=∞	$k=p_1$ [2] $\{7\}$	$k=p_1$ [2] $\{\infty\}$	$k=p_2$ [2] $\{\infty\}$	$k(L)=p_1$, $k(H)=p_0$ [3] $\{\infty\}$	$k^{down}=p_1$, $k^{up}=p_3$	$k^{down}=p_1$, $k^{up}=p_4$	$k^{up}(L)=p_4$, $k^{up}(H)=p_3$; $k^{down}(L)=p_1$, $k^{down}(H)=p_0$	$k^{up}(L)=p_4$, $k^{up}(H)=p_3$; $k^{down}(L)=p_2$, $k^{down}(H)=p_1$
γ	N/A	0.81*	1.04*	1.61*	0.86*	1.21*	1.86*	1.02*	1.14*	1.21*	0.92*	1.74*
LL	-592.2	-570.9	-565.9	-555.3	-555.7	-537.1	-554.54	-553.4	-539.7	-537.1	-551.9	-555.6
LR†	N/A	42.56	52.45	73.77	73.05	110.22	75.34	77.61	105.03	110.22	80.57	73.15
p-value		<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01

Parameterization 2

	Random	GP_N			$T1_N$ $\{N\}$			$TT1_N$ $\{N\}$	$T2$		$TT2$	
		N=3[1]	N=13	N=∞	$k=p_1$ [2] $\{8\}$	$k=p_1$ [2] $\{\infty\}$	$k=p_2$ [2] $\{\infty\}$	$k(L)=p_1$, $k(H)=p_0$ [3] $\{\infty\}$	$k^{down}=p_1$, $k^{up}=p_3$	$k^{down}=p_1$, $k^{up}=p_4$	$k^{up}(L)=p_4$, $k^{up}(H)=p_3$; $k^{down}(L)=p_1$, $k^{down}(H)=p_0$	$k^{up}(L)=p_4$, $k^{up}(H)=p_3$; $k^{down}(L)=p_2$, $k^{down}(H)=p_1$
γ	N/A	0.33*	0.79*	1.30*	0.74*	1.33*	1.45*	0.91*	1.25*	1.33*	0.84*	1.43*
LL	-549.8	-543.6	-534.4	-523.7	-526.2	-502.3	-523.6	-521.3	-503.9	-502.3	-521.6	-522.6
LR†	N/A	12.50	30.85	52.25	47.22	95.02	52.42	56.95	91.78	95.02	54.41	56.35
p-value		<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01

Notes: Estimates of α, ρ and ψ are significant at the 1% level in all specifications (not shown). Number of observations: 1,200 in both parameterizations. N=number of punishment periods. LL= Log-likelihood. GP=Green and Porter theoretical prediction; strategies T1, TT1, T2 and TT2 are defined in table 11. k are threshold levels, also defined in table 11.

[1] Parameterization 1: the GP equilibrium is not feasible; Parameterization 2: the GP equilibrium is feasible for punishment lengths 3-∞. Shown are only the three punishment lengths with the highest explanatory power (except for N=3 in parameterization 2, shown for reference only).

[2] Other finite punishment lengths also had statistical significance, but were less significant than the one reported.

[3] Strategies with threshold levels $k(L)=p_2$ and $k(H)=p_1$ also had statistical significance, but were less significant than the one reported.

† Likelihood ratio test with respect to the random strategy

* Significant at 1% level

The evidence presented in table 12 is interpreted as providing support for the existence of trigger strategies in general, and in particular for the grim-trigger strategy. Also, results provide support for the strategies predicted by the GP theory, with some caveats. First, we do not observe a substantial difference in behavior across parameterizations, even when the incentives favor a stronger evidence for GP in parameterization 2. Also, collusion is not significantly larger in parameterization 2, even though the incentives for defection are smaller than in parameterization 1; this means that some subjects are cooperating less than what theory predicts. Finally, punishment strategies seem to be infinite rather than finite (which GP is known for).

Equilibrium Outcomes

We adopt a similar strategy as in the analysis of equilibrium outcomes of the FI treatment. Recall that the equilibrium set is given by the "always defect" outcome (*H*,*H*), in both parameterizations, and by the GP equilibria in parameterization 2 (with threshold level p_2 and punishment lengths from 3 to ∞). The predictive power of different GP "equilibrium paths" is presented in table 13; note that this test is different than that presented in the previous section: here we analyze whether *both* players defect when the GP path predicts so. The results are consistent with what was observed in the preceding section: a) subjects appear to apply the grim-trigger strategy, and b) there are no substantial differences in behavior across parameterizations. Moreover, it appears as if the GP equilibrium paths have a higher predictive power for parameterization 1, which is contrary to what one would expect: GP equilibrium paths considered can only be sustained as such in parameterization 2.

7. Robustness Checks

With the level of risk aversion observed in our sample, the intended theoretical incentive for the RS theory in parameterization 1 (i.e. collusion is not an equilibrium when demand is high) no longer holds. We conducted additional sessions with parameters that restore the intended incentives under the observed level of risk aversion (parameterization 3, see Appendices A and D). In addition, these sessions allow us to further check whether demand information remains as the key factor affecting collusion in our design. Appendix D reports the results of the estimations, which strongly confirm the latter finding. Conversely, the evidence for the RS theory still exists but is not as strong as in parameterization 1. In particular, the tit-for-tat and the grim strategies now appear to describe data better than the RS strategy.

Table 13: Fraction of Times the Equilibrium Path Correctly Predicts Outcomes

Equilibrium Path	**Parameterization 1**	**Parameterization 2**
(*H*,*H*) every period	**36.87%**	**33.59%**
GP_3	50.00%	**50.63%**
GP_4	56.82%	**53.91%**
GP_5	48.48%	**48.36%**
GP_6	52.27%	**51.14%**
GP_{15}	66.16%	**60.48%**
GP_{16}	67.68%	**60.73%**
GP_{17}	68.69%	**61.74%**
GP_{18}	70.20%	**61.74%**
GP_{∞}	71.72%	**62.25%**

Notes: Bold numbers indicate a theoretically feasible equilibrium (see Appendix A for details). The GP_N path takes a value of 1 when collusion is predicted and 0 when a price war is predicted; a price war is assumed to be triggered by a low signal (price$\leq p_2$) which lasts *N* periods.

With the observed level of risk aversion, the set of GP equilibria gets larger (see Appendix B). Specifically, the GP equilibrium becomes feasible in parameterization 1 for punishment lengths that range from 6 periods to ∞. Another possible equilibrium emerges for parameterization 2: with threshold level p_1, the feasible range of punishment lengths for the GP equilibrium is [6,..,∞]; in addition, the punishment length for a threshold of p_2 increases its range to [2,..,∞]. This attenuates our interpretation of the results regarding the similar behavior across parameterizations being construed as lack of evidence for the incentives implied by the GP equilibrium. In addition, since p_1 is now a threshold level that yields feasible GP equilibria, the estimation results for strategy *T1* are no longer inconsistent with the GP predictions.

We still note, however, that the finite punishment behavior for which GP is known does not describe the data as well as the infinite punishment strategy. To be sure, we carried out an additional check: we varied the random draws that determine the demand states and conducted additional sessions with parameterization 2. Our main results are robust (Appendix D).

8. Discussion

In this paper we focus on two factors (demand information and monitoring) that have played a key role in the theory of infinitely repeated games with stochastic demand. Guided by theory, we construct experiments to study the effect of these two factors on collusion. Results indicate that monitoring appears to always increase collusion, whereas the effect of knowing the demand schedule (to be faced next period) is either negligible or may even reduce collusion (if the parameters of the game are appropriately calibrated). We show that, while counterintuitive, this result is consistent with theoretical predictions of our design. Thus a central conclusion of this work is that theory plays a crucial role determining the effect of each factor.

The large number of equilibria that are theoretically possible in infinitely repeated games has been frequently criticized by many empirical economists. As a second objective, we attempt to bridge this gap between theory and empirics by studying whether the predictions of two influential theories are supported by experimental data from treatments that resemble the assumptions of each theory. We carry out two types of analysis. First we study individual strategies and find that data can be explained relatively well by several alternatives. Specifically, we reject the random strategy (i.e. flipping a coin every period to decide whether to cooperate) in favor of the strategies predicted by the RS and GP theories. But behavior can also be explained relatively well by other strategies, especially the grim-trigger strategy which explains the data best in both cases. This result suggests that reversions to competition are more likely to be *permanent* rather than temporary. What triggers permanent defection is different in each model, however. In the RS model, the grim strategy is triggered by defection from either player, whereas in the GP model it is triggered by a low price (or a negative demand shock).

Second, we analyze how observed outcomes (not individual strategies) lend support to the various equilibria that are theoretically possible. Results indicate that within the set of possible equilibria, the RS and GP equilibria, in their respective treatment, tend to have the highest explanatory power.

The traditional (or simple) interpretation of the RS model is that it is a theory of countercyclical pricing (*temporary* low price during high demand), whereas the GP theory is usually attributed with somewhat an opposite prediction (*temporary* low price after low demand). Taking the traditional interpretation of the theories at face value, the strong support for the grim strategy would suggest that the theories considered, while plausible, are not the best explanation of observed data. It is important to note, however, that the theoretical feasibility of a GP equilibrium path that allows for a finite punishment of length N, also allows for equilibrium paths with punishment lengths $N+1,\ldots,\infty$. In this broader interpretation of the GP equilibrium, our results are supportive of the GP theory, just not the finite price war equilibria for which it is well known. In this sense, our results cast some doubt on the likelihood of whether observed finite price wars (e.g. Porter, 1983) are evidence for the GP theory at work.

Conversely, the finite price war considered by RS is different in nature: it is not a punishment mechanism to deter collusion; instead it is observed because demand is sometimes too high to prevent firms from deviating. As such, it does not predict price wars of extensive length. While we do find evidence for a drop in collusion when demand is high, the evidence still tends to more strongly favor permanent price wars (triggered in *any* demand state). However, while observing the grim-trigger strategy is not consistent with the RS equilibrium path, this is the behavior RS *assume* players should have in order obtain their result. In this sense, observing the grim strategy is not entirely inconsistent with the RS theory.

Given this preceding discussion, our overall assessment of whether the RS and GP theories explain data well is positive, even though the evidence is not overwhelming. A reason for this assessment is that the theories are being tested against several alternatives. For example, while the tit-for-tat strategy has been reported to be perhaps the most successful strategy in the repeated prisoner's dilemma game, the explanatory power of the strategy implied by the RS theory is superior in one of the two parameterizations considered.

In general, we find that the RS environment is more "collusive friendly" than the GP environment. We conjecture that the uncertain environment implied by GP may be responsible for this as well as for why finite punishments do not emerge as a major explanation of behavior.

Finally, we observe a large level of heterogeneity in our sample as evidenced by the acceptance of the random effects specification in all regressions, and by the existence of a variety of strategies and equilibria that explain data relatively well. Classifying subjects into "types" is another possibility of analysis that we leave for further research.

REFERENCES

Abreu, D., D. Pearce, and E. Stacchetti. 1990. "Toward a theory of discounted repeated games with imperfect monitoring," *Econometrica,* 58, 1041-1063.

Aoyagi, M. and G. Frechette. 2008. "Collusion as Public Monitoring Becomes Noisy: Experimental Evidence," Working Paper, NYU.

Bagwell, K. and R. W. Staiger. 1997. "Collusion Over the Business Cycle," *Rand Journal of Economics*, 28, 1, 82-106.

Bereby-Meyer, Y. and A. Roth. 2006. "The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation," *American Economic Review*, 96, 1029-42.

Cason, T. and C. Mason. 1999. "Information Sharing and Tacit Collusion in Laboratory Duopoly Markets," *Economic Inquiry*, 37, 2, 258-79.

Chevalier J., A. Kashyap and P. Rossi. 2003. "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," *American Economic Review*, 93, 1, 15-37.

Dal Bó, P. 2005. "Cooperation Under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games," *American Economic Review*, 95, 5, 1591-604.

Eckel, C. and P. Grossman. 2007. "Forecasting Risk Attitudes: An Experimental Study Using Actual and Forecast Gamble Choices." Forthcoming, *Journal of Economic Behavior and Organization*.

Ellison, G. 1994. "Theories of Cartel Stability and the Joint Executive Committee," *Rand Journal of Economics*, 25, 1, 37-57.

Engle-Warnick, J. and R. Slonim. 2006. "Inferring Repeated-Game Strategies from Actions:
Evidence from Trust Game Experiments," *Economic Theory*, 28, 603-32.

Feinberg, R. and C. Snyder. 2002. "Collusion with Secret Price Cuts: An Experimental Investigation," *Economics Bulletin*, 3, 6, 1-11.

Fischbacher, U. 1999. "z-Tree - Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual," *Working Paper*, Nr. 21, Institute for Empirical Research in Economics, University of Zurich.

Frechette, G., J. Kagel and M. Morelli, 2005. “Behavioral Identification in Coalitional Bargaining: An Experimental Analysis of Demand Bargaining and Alternating Offers,” *Econometrica,* 73, 6, 1893-1938.

Friedman, J. 1971. “A Noncooperative Equilibrium for Supergames,” *Review of Economic Studies*. 38, 113, 1-12.

Fudenberg, D. and J. Tirole, 1989 “Non-Cooperative Game Theory for Industrial Organization: An Introduction and Overview,” in R. Schmalensee and R. Willig, eds., *Handbook of Industrial Organization*, Vol. 1, Amsterdam: Elsevier Science.

Green, E. and R. Porter. 1984. “Noncooperative Collusion under Imperfect Price Competition.” *Econometrica*, 52, 87-100.

Haltiwanger, J. and J. Harrington. 1991. “The Impact of Cyclical Demand Movements on Collusive Behavior,” *RAND Journal of Economics*, 22, 1, 89-106.

Harrison, G. and E. Ruström. 2008. “Risk Aversion in the Laboratory,” in J.C. Cox and G.W. Harrison (eds.), *Risk Aversion in Experiments*, Bingley, UK: Emerald, Research in Experimental Economics, Volume 12.

Holcomb, J. and P. Nelson. 1991. “Cartel Failure: A Mistake of Do They Do It to Each Other on Purpose.” *Journal of Socio-Economics*, 20, 3, 235-249.

Holcomb, J. and P. Nelson. 1997. “The Role of Monitoring in Duopoly Market Outcomes.” *The Journal of Socio-Economics*, 26, 1, 79-93.

Kandori, M. 1991. “Correlated Demand Shocks and Price Wars during Booms,” *Review of Economic Studies*, 58, 1, 171-80.

Porter, R. 1983. 1983. “A Study of Cartel Stability: The Joint Executive Committee, 1880-1886,” *Bell Journal of Economics*, 301-14

Rotemberg, J.J. and G. Saloner. 1986. “A Supergame-Theoretic Model of Price Wars During Booms.” *American Economic Review*, 76, 390-407.

Weber, E. U., Shafir, S., & Blais, A.-R. 2004. “Predicting risk-sensitivity in humans and lower animals: Risk as variance or coefficient of variation.” *Psychological Review*, 111, 2, 430-445.

Wooldridge, J. 2003. *Econometric Analysis of Cross Section and Panel Data*. Cambridge: MIT Press.

APPENDIX A: Equilibrium Set

A.1. The Full Information Treatment (RS theory)

In order for collusion (*L*,*L*) to be a feasible equilibrium under a grim-trigger strategy, we know that condition (2) must hold. Further, it is easy to show that playing the one-shot NE equilibrium (*H*,*H*) in every period is also an equilibrium. In addition, as shown by Stahl II (1991), the strategy pair (*L*,*H*) can also be sustained as an equilibrium outcome. Adapting Stahl's approach to our stochastic demand structure, the strategy pair (*L*,*H*) *can* be an equilibrium outcome in demand state *s* if there exists at least one payoff Π^x for player 1, such that the following condition is met:[i]

$$(1-\delta)\Pi_s^{ND} + \delta E(\Pi^x) \geq (1-\delta)\Pi_s^{NE} + \delta E(\Pi^y)$$

where $E(\Pi^z) = 0.2\Pi_h^z + 0.6\Pi_m^z + 0.2\Pi_l^z$. The left hand side consists of the payoff of playing *L* today (when the opponent is playing *H*) plus the expected discounted profits of obtaining profits Π^x for each of the remaining rounds of the game; the right hand side consists of the payoff of switching to *H* today (when the opponent is playing *H*) plus the expected discounted payoffs of playing a strategy that yields a payoff of Π^y. By setting Π^y equal to Π^{NE} (i.e. the grim strategy) one makes the right hand side as small as possible, thus making the following condition a sufficient one for (*L*,*H*) to be an equilibrium:

$$E(\Pi^x) \geq \frac{(1-\delta)(\Pi_s^{NE} - \Pi_s^{ND}) + \delta E(\Pi^{NE})}{\delta}$$

In our case, it suffices to check if, for each demand state, the above condition holds when we set $\Pi^x = \Pi^C$. If it does, then, by symmetry, the equilibrium (*H*,*L*) is also feasible. It is important to note, again, that these equilibria are *state dependent*.

The above equilibria are derived using a grim-trigger strategy. However, less severe punishments (i.e. finite reversions to the NE outcome) can also support the same equilibria. For the (*L,L*) outcome to be an equilibrium in state *s* under a finite punishment strategy of length *N*, the following inequality (using Bellman's equation) should hold:

$$(1-\delta)\Pi_s^C + \delta E(\Pi^C) \geq (1-\delta)[\Pi_s^D + \sum_{t=1}^{N} E(\Pi^{NE})] + \delta^{N+1} E(\Pi^C)$$

After some manipulation, this condition becomes:

$$E(\Pi^C) \geq \frac{(1-\delta)\left[\Pi_s^D - \Pi_s^C + \frac{\delta(1-\delta^N)}{1-\delta} E(\Pi^{NE})\right]}{\delta - \delta^{N+1}}$$

Again, note that this condition is state dependent. In a similar fashion, a condition can be obtained for the (*L*,*H*) or (*H*,*L*) equilibria. Results of computations using the above formulae are summarized in table A.1 below.

[i] Stahl II shows that only player 1's condition is the relevant one.

A.2 The Imperfect Monitoring Treatment (GP theory)

As with the FI treatment, it is easy to show that the playing the one shot NE (*H*,*H*) in every period is an equilibrium of this game. Collusion, on the other hand, can be sustained through the use of trigger strategies; but, as opposed to the RS equilibrium, collusion here refers to a "collusive path" rather than observing (*L*,*L*) every period: reversion to the NE play is part of the equilibrium path for some period of time (finite or infinite). This is an important difference because observing reversion to the NE play in the RS model may be consistent with RS strategies but inconsistent with the RS equilibrium path, whereas in the GP model reversion is consistent with both equilibrium strategies and the equilibrium path.[ii]

Table A.1: Feasible Equilibria by Demand State and Parameterization in the FI Treatment, (strategy pair), [range of feasible punishment lengths "*N*"]

Demand State	Parameterization 1	Parameterization 2	Parameterization 3
High (*h*)	(*H*,*H*) (*H*,*L*)/(*L*,*H*) [2-∞]	(*L*,*L*) [3-∞]* (*H*,*H*) (*H*,*L*)/(*L*,*H*) [1-∞]	(*H*,*H*)
Medium (*m*)	(*L*,*L*) [2-∞] (*H*,*H*) (*H*,*L*)/(*L*,*H*) [1-∞]	(*L*,*L*) [1-∞] (*H*,*H*) (*H*,*L*)/(*L*,*H*) [1-∞]	(*L*,*L*) [5-∞] (*H*,*H*) (*H*,*L*)/(*L*,*H*) [3-∞]**
Low (*l*)	(*L*,*L*) [1-∞] (*H*,*H*) (*H*,*L*)/(*L*,*H*) [1-∞]	(*L*,*L*) [1-∞] (*H*,*H*) (*H*,*L*)/(*L*,*H*) [1-∞]	(*L*,*L*) [1-∞] (*H*,*H*) (*H*,*L*)/(*L*,*H*) [1-∞]

Note: This table corresponds to the case of risk neutrality. The table for the case of risk aversion (see Appendix B) is identical except for the entries marked with asterisks.
* For the case of risk aversion this entry has a minimal feasible punishment length of 2 (instead of 3).
** For the case of risk aversion this entry has a minimal feasible punishment length of 10 (instead of 3).

In this type of equilibrium, there are two variables that need to be calibrated: the punishment length and the threshold level. As shown by Abreu et al. (1990), the set of equilibrium payoffs can be obtained by considering a grim-trigger strategy and then finding the smallest public signal for which collusion is incentive compatible. Our design makes it straightforward to compute the feasible set of "collusive" equilibria (and their corresponding range of incentive compatible punishment lengths) using Abreu et al's insight. We proceed as follows: a) for each of the 5 possible threshold levels (one for each price, see Appendix C), we compute whether a grim-trigger strategy satisfies condition (4), and then b) if condition (4) is satisfied for a grim-trigger strategy, we find the minimum *N* for which (4) holds. Table A.2 presents the results of this calculation.

[ii] Because monitoring is imperfect, Stahl's results regarding the feasibility of (*H*,*L*)/ (*L*,*H*) do not apply to this case.

Table A.2: Feasible Equilibria in the IM Treatment, (choices),[range of feasible punishment lengths "*N*"]

Equilibrium Path	**Parameterization 1**	**Parameterization 2**
(*H*,*H*) every period	Yes	Yes
Trigger: (*L,L*) as long as observed *p* greater than:		
p_0	No	No
p_1	No	No**
p_2	No*	Yes [3-∞]***
p_3	No	No
p_4	No	No

Notes: Parameterization 3 is not considered as it does not have the imperfect monitoring feature. This table corresponds to the case of risk neutrality; the table for the case of risk aversion (see Appendix B) is identical except for the entries marked with asterisks.

* (L,L) is a feasible outcome in the case of risk aversion, with range [6-∞]

** (L,L) is a feasible outcome in the case of risk aversion, with range [6-∞]

*** (L,L) is also feasible outcome in the case of risk aversion, but the range is [2-∞]

Appendix B: Risk Aversion Estimate and Parameterization Details

B.1. Risk Aversion Estimate

The underlying assumption of the theories studied in this paper is that subjects are risk neutral, or that utility is of the form $u(x) = x$, where x is the monetary payoff. To allow for the possibility of risk averse (or risk seeking) behavior, we adopt the widely used constant relative risk aversion (CRRA) utility specification: $u(x) = x^{1-r} / 1 - r$, where r is the CRRA coefficient and nests risk neutrality ($r = 0$), risk aversion ($r > 0$) and risk seeking behavior ($r < 0$). To obtain an estimate of r for our sample, we employ the elicitation method proposed by Eckel and Grossman (2008). The method consists of asking subjects to choose from among six possible gambles the one they would most prefer to play (Table B.1); after choosing the gamble, a die is rolled to determine whether the subject obtains the high or the low payoff.[37]

Table B.1: Eckel-Grossman Gamble Choices

Choice (50/50 Gamble)	Low Payoff	High Payoff	Expected Return	Standard Deviation	Fraction of Subjects Choosing Gamble (%)
Gamble 1	18	18	18	0	9.2
Gamble 2	14	26	20	8.5	17.4
Gamble 3	10	34	22	17	27.2
Gamble 4	6	42	24	25.5	11.3
Gamble 5	2	50	26	34	28.7
Gamble 6	-2	54	26	40	6.2

Structural estimation assumes that utility is of the von Neumann-Morgentstern type and hence individuals evaluate alternatives based on the weighted average (expected utility): $EU_i = p_L u(i_L) + p_H u(i_H)$, where i denotes gambles 1 through 6, the subscripts *L* and *H* denote the Low and High payoffs, respectively, and *p* denotes the probability of occurrence. Individuals then choose the gamble that provides them with the highest utility level (with some econometric error). In order to obtain an estimate of *r*, we adopt the logit specification (see Harrison and Rustrom, 2008):

$$y_i^* = \frac{\exp\left(EU_i\right)}{\sum_{i=1}^{6} \exp\left(EU_i\right)}$$

where, y_i^* is a latent index based on the assumed preferences. The log-likelihood is then:

$$L(r) = \prod_j \prod_i (y_i^*)^{d_{ji}}$$

where, $d_{ji} = 1$ if individual *j* chooses gamble *i*, and zero otherwise.

---

[37] The instructions, available at http://www.umass.edu/resec/faculty/rojas/z-tree.html, contain details of the protocol used.

B.2. Parameterization Details

The estimated CRRA coefficient using the above procedure is r = 0.46 (SE=0.03), which means that subjects are risk averse. While this estimate is slightly smaller than those found by Harrison and Rustrom (2008) in their exercise of various elicitation methods (their estimates range from 0.51 to 0.86), it is in line with the common experimental finding that subjects are risk averse. Table B.2 contains details of the equilibrium conditions implied by both parameterizations under risk neutrality, as well as under risk aversion.

Table B.2: Implied Equilibrium Conditions for RS and GP Models

	Parameterization 1		**Parameterization 2**	
	$r = 0$	$r = 0.46$	$r = 0$	$r = 0.46$
Implied δ^* for RS Model (h, m and l demand states)	0.76	0.67	0.62	0.58
	0.48	0.51	0.32	0.42
	0.21	0.34	0.12	0.27
Implied (minimum) Punishment length for GP Model	N/A	6	3	2

Note: $r = 0$ denotes risk neutrality; $r = 0.46$ is the estimated CRRA coefficient.

For parameterization 1, note that the implied critical discount factor (δ^*) in the high demand state under risk neutrality is slightly above the "simulated" discount factor (0.75) (but below - δ^* =0.67- under risk aversion). For parameterization 2, δ^* is always below the simulated discount factor. Thus, we would expect that in parameterization 1, the predictions of the RS theory (i.e. breakdown of collusion when demand is high) would hold under risk neutrality (but not under risk aversion), while in parameterization 2 collusion should be an equilibrium in all demand states. Risk aversion implies a shorter optimal punishment period for the GP model in both parameterizations.

Thus, risk aversion may pose a problem for our test of the RS theory, but note that the implied δ^* in the high demand state (parameterization 1) is relatively close to the simulated discount factor and hence the intended incentives of our design may (weakly) work in this case (as shown by the results in the paper). To be sure, however, we carried out additional sessions for the FI treatment with a third parameterization that addresses this problem; Appendix D shows the parameters used in these sessions. The main difference between parameterization 3 and the other two parameterizations is that it has a much smaller difference between the collusive outcome and the non-cooperative outcome (especially in the high demand state); this effectively reduces the incentives to collude thereby increasing δ^*. Parameterization 3 has implied (critical) discount factors (for each demand state) of 0.88, 0.68 and 0.37 under risk neutrality and of 0.77, 0.62 and 0.44 under risk aversion (r = 0.46).

On the other hand, risk aversion is not a problem for testing the GP theory; on the contrary, in both parameterizations risk aversion implies a shorter optimal punishment length, giving the GP theory a better chance of occurrence (given the length of 33 periods).

APPENDIX C: Imperfect Public Monitoring

Consider parameterization 1. Without loss of generality, let's assume that the collusive quantity (*L*) is equal to 1; given our design, *L=1* automatically defines *H* (the high quantity), as well the prices (the public signal) in all the cells of the three payoff tables (high, medium and low demand) as follows:

Table C.1: Implied High Quantity (*H*) and Prices (p_i) in Parameterization 1 when *L*=1

		High Demand		**Medium Demand**		**High Demand**	
		Player 2		Player 2		Player 2	
		L=1	*H=5.73*	***L=1***	*H=5.73*	***L=1***	*H=5.73*
Player 1	***L=1***	$p_4 = 26$	$p_3 = 7.5$	$p_3 = 7.5$	$p_2 \approx 2.1$	$p_2 = 2.1$	$p_1 \approx 0.6$
	H=5.73	$P_3 = 7.5$	$p_2 \approx 2.1$	$p_2 \approx 2.1$	$p_1 \approx 0.6$	$p_1 \approx 0.6$	$p_0 = 0.17$

Notes:

a) *H=5.73* is obtained by dividing 43 (player 2's profit when player 1 chooses *L*, player 2 chooses *H* and high demand occurs) by 7.5 (the price that must hold in the *L,L* cell of the medium demand state and in the *L,H* cell of the high demand state). A slightly different *H* (because of rounding of profit numbers) can be obtained by dividing 12.5 (player 2's profit when player 1 chooses *L*, player 2 chooses *H* and medium demand occurs) by 2.1 (the price that must hold in the *L,L* cell of the low demand state and in the *L,H* cell of the medium demand state).

b) Because of rounded profit (which is easier for subjects to understand), the implied prices in some instances have also been approximated (the ones with a " ≈ ")

Note that the noisy price signal (indirectly) received by both players is the same, and hence it is public. For example, suppose that player 1 chooses to play the collusive outcome (*L*) and that player 2 chooses to defect (*H*). Further, suppose that medium demand occurred. Then, player 1 receives the noisy price signal of p_2 (upper right cell of the medium demand matrix or the upper left cell of low demand matrix) and so does player 2 (upper right cell of the medium demand matrix or the lower right cell of high demand matrix). A similar exercise can demonstrate the public nature of monitoring for parameterization 2.

Importantly, the public nature of the noisy signal is achieved for any number one wishes to choose for *L*. More generally, one can think of this design as a symmetric two-player Cournot game with two strategies and stochastic demand. This game maps choices $y = (y_1, y_2)$ onto price through a stochastic (demand) function, generating a conditional distribution *f*(*p*|*y*); in our case *f* has discrete support, that is $p \in (p_0, ..., p_4)$.

APPENDIX D: Results of Robustness Checks

First, we constructed an additional parameterization (table D.1), to address the potential drawback caused by subjects' risk aversion in the analysis of the evidence for the RS theory (see appendix A). To achieve the desired critical discount factors in the presence of risk aversion, however, the imperfect monitoring characteristic could no longer be maintained; thus, sessions were run only for the FI and M treatments.[i] A total of 102 subjects from the University of Massachusetts participated in 3 sessions (2 for treatment FI and 1 for treatment M); mean earnings (excluding show up fee and risk task payments) were \$19.12 for the FI treatment and \$19.40 for the M treatment. Further, parameterization 3 also serves as a robustness check for our other main finding (demand information removal does not decrease collusion).

Second, to check the robustness of the GP results, we ran additional sessions with parameterization 2 but varied the random draws that determine the demand states. These draws can be seen in figure D.1 below; we call it parameterization 2b. A total of 74 subjects from the University of Massachusetts participated in 4 sessions (2 for the IM treatment and 1 for each of the other two treatments). Average earnings were \$31.06, \$33.42 and \$24.73 for the FI, M and IM treatments, respectively. In addition, parameterization 2b also allows us to further check the results obtained in the analysis of evidence for the RS treatment.

Table D.1: Parameterization 3

		High Demand		**Medium Demand**		**High Demand**	
		Player 2		Player 2		Player 2	
		L	H	L	H	L	H
Player 1	L	17.00, 17.00	2.00, 31.00	5.00, 5.00	0.50, 9.00	1.40, 1.40	0.20, 2.50
	H	31.00, 2.00	12.50, 12.50	9.00, 0.50	3.50, 3.50	2.50, 0.20	1.00, 1.00

D.1 Effect of Demand Information on Collusion

We use parameterizations 2b and 3 to check the robustness of whether removal of demand information does not decrease collusion; table D.2 confirms this finding. In both parameterizations, cooperation is not statistically different between the FI and the M treatments, whereas collusion is statistically larger when demand information is removed. Figure D.1 confirms (for each parameterization) the similar frequency of cooperation in the two treatments.[ii]

[i] To minimize confounding effects, we made a special effort to change as few parameters as possible (i.e. the NE payoffs are the same as in parameterizations 1 and 2).

[ii] Parameterization 3 also confirms that the removal of monitoring (IM treatment) decreases the rate of cooperation and collusion with respect to both the FI and M treatments (results are not reported).

Table D.2: Frequencies of Cooperation and Collusion (standard deviation)

Treatment	**Parameterization**	**# Obs.**	**Frequency of Cooperation***	**Frequency of Collusion****
FI	2b	660	0.87 (0.33)	0.62 (0.49)
	3	*2112*	*0.64 (0.48)*	*0.48 (0.50)*
M	2b	792	0.90 (0.29)	0.83 (0.37)
	3	*1254*	*0.64 (0.48)*	*0.57 (0.49)*

* At least one player chooses *L*. ** Both players choose *L*.

Figure D.1: Frequency of Cooperation over 33 Periods of Stochastic Demand: h=high [---], m=medium or l=low [—]; (Parameterization)

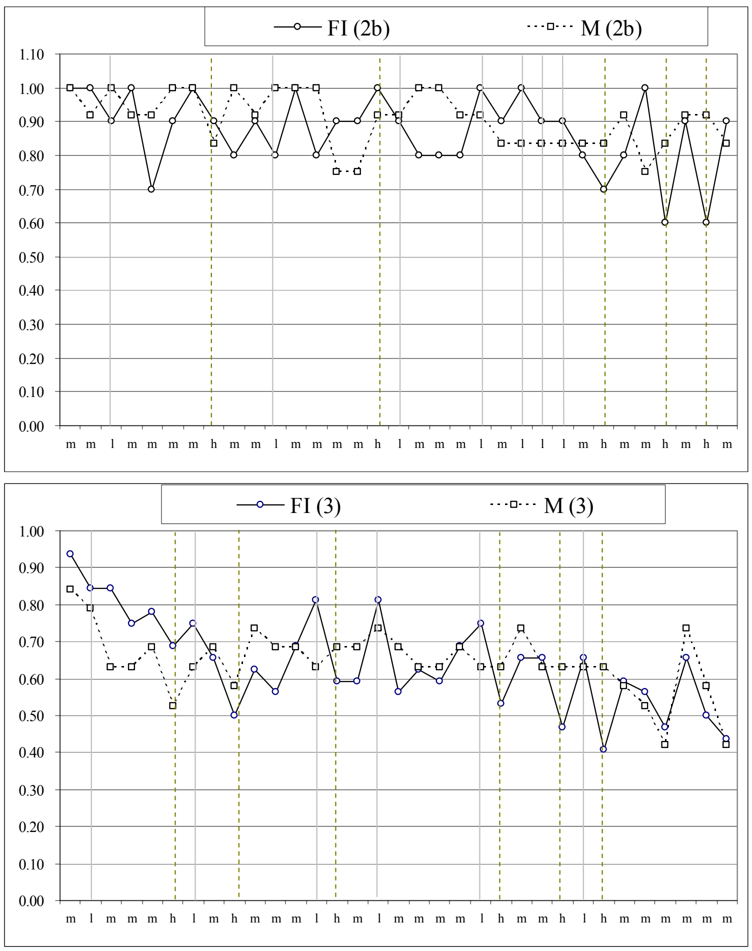

D.2 Evidence for the RS theory

Collusion frequencies shown in table D.3 are somewhat similar to those observed for parameterization 1 (table 6 of the paper). Collusion in parameterization 2b appears to be lower than in parameterization 2, whereas collusion in parameterization 3 appears to be similar as in parameterization 1. Consistent with the results reported in the paper, collusion is statistically lower when demand is high (with respect to the other two demand states) in parameterization 3, but not so in parameterization 2b. Similarly, the frequency of collusion in the medium demand state is not statistically different than that observed in the low demand state.

Turning to the analysis of individual strategies, the regressions reported in table D.4 still provide support for the RS strategy, but this evidence is not as strong as in parameterization 1. In particular, the TT strategy has a higher explanatory power than the RS strategy; the grim strategy continues to be the most significant (single) strategy and it can even explain data better than the combined RS+TT strategy.

Table D.3: Frequencies of Cooperation and Collusion in Full Information Treatment (St. Dev.)

Demand State	Param.	All Observations (Periods 1-33)		Periods 1-25	
		# Obs.	Freq. Collusion**	# Obs.	Freq. Collusion**
High (h)	2b	100	0.44 (0.50)	40	0.55 (0.50)
	3	*384*	*0.42 (0.49)*	*320*	*0.43 (0.50)*
Medium (m)	2b	420	0.65 (0.48)	320	0.66 (0.43)
	3	*1,344*	*0.49 (0.50)*	*960*	*0.54 (0.50)*
Low (l)	2b	140	0.64 (0.48)	140	0.64 (0.48)
	3	*384*	*0.52 (0.50)*	*320*	*0.53 (0.50)*

* Frequency of either player choosing L. ** Frequency of *both* players choosing L.

Table D.4: Probit Estimates of Different Strategies in the FI treatment, Parameterization 3, Rounds 1-25

Parameter	Random	RS	TT	P-∞	RS + TT
α	-0.87***	-0.75	-1.12*	-0.76*	-1.02**
γ_1		0.83*			0.94*
γ_2			0.88*		0.99*
γ_3				2.90*	
ψ	2.21*	2.30*	1.70*	0.23	1.79*
ρ	0.78*	0.80*	0.70*	0.32*	0.73*
LL	-526.29	-505.08	-499.87	-465.45	-473.57
LR Test (p-value)†	N/A	42.42 (<0.01)	52.84 (<0.01)	121.68 (<0.01)	105.44 (<0.01)

* Significant at 1%. ** Significant at 5%. *** Significant at 10%. † Likelihood ratio test with respect to the random strategy.

Notes: # of Observations =1,536 in all models to keep number of observations comparable across strategies (first period is lost in TT strategy). RS=Rotemberg and Saloner equilibrium strategy, TT=Tit-for-Tat, P-∞= infinite punishment. LL=Log-likelihood. Finite punishment strategies also have significant explanatory power but much smaller than the strategies displayed.

The frequencies of observed outcomes in parameterization 3 (displayed in table D.5) are similar to those for parameterization 1 presented in table 9 of the paper; the main difference is that the (H,H) outcome now appears to be more frequent (especially in the medium state). When compared with parameterization 2, observed outcomes for parameterization 2b are somewhat different, however; the main difference is that here the collusive outcome is observed less frequently (especially in the high demand state), while the (*H*,*L*)/(*L*,*H*) outcome is now observed much more frequently. This is our least strong robustness result.

The non-parametric tests are, however, consistent with what was reported in the paper. The best fit in parameterization 3 is given by the RS equilibrium (51%) followed by the "always collude" outcome (48%), the "always defect" outcome (35.61%), and the (*H*,*L*)/(*L*,*H*) outcome (16.11%). On the other hand, the best fit in parameterization 2b is given by the "always collude" outcome (65%), the RS equilibrium (64%), the (*H*,*L*)/(*L*,*H*) outcome (25.76%), and the "always defect" outcome (12.73%). Again, the evidence from parameterzation 2b is not as conclusive as that of parameterization 2.[iii]

Table D.5: Frequencies of Observed Outcomes

Demand State (outcomes)		**Parameterization 2b**	**Parameterization 3**
High (*h*)	(*L*,*L*)	**44.00%**	41.67%
	(*H*,*H*)	**24.00%**	**46.88%**
	(*H*,*L*)/(*L*,*H*)	**32.00%**	**11.46%**
Medium (*m*)	(*L*,*L*)	**64.76%**	**49.11%**
	(*H*,*H*)	**11.43%**	**36.01%**
	(*H*,*L*)/(*L*,*H*)	**23.81%**	**14.88%**
Low (*l*)	(*L*,*L*)	**64.29%**	**52.08%**
	(*H*,*H*)	**8.57%**	**22.92%**
	(*H*,*L*)/(*L*,*H*)	**27.14%**	**25.00%**

Notes: Bold numbers indicate that entry is a feasible equilibrium (see Appendix B, table B.1 for details)

D.3 Evidence for the GP theory

Table D.6 is consistent with the results obtained for parameterization 2 (reported in table 10 of the paper): large punishment lengths tend to explain cooperation and collusion better by than short ones. Similarly, regression results designed to study subjects' strategies (table D.7) are in line with those obtained for parameterization 2b (table 12 of the paper): a) the grim-trigger strategy appears to explain the data best, and b) the most likely threshold level for starting a price war appears to be p_l (regardless of the rule for returning to the collusive regime). Finally, results displayed in table D.8 below are also consistent with those reported for parameterization 2 in table 13 of the paper.

[iii] One reason for this could be the difference in number of observations: parameterization 2 has 1,584 observations, while parameterization 2b has 660.

Table D.6: Frequencies of Cooperation and Collusion in Collusive (C) and Reversionary (R) Regimes in IM treatment, Various Punishment Lengths, Rounds 1-25, Parameterization 2b

Punishment Length (N)	**Cooperation**			**Collusion**		
	R	C	*p-value**	R	C	*p-value**
2	0.68	0.67	*0.73*	0.18	0.24	*0.11*
3	0.69	0.67	*0.60*	**0.21**	**0.23**	***0.46***
4	0.66	0.70	*0.25*	**0.19**	**0.26**	***0.01***
14	0.66	0.78	*0.02*	**0.19**	**0.44**	***<0.01***
15	0.66	0.78	*0.02*	**0.19**	**0.44**	***<0.01***
∞	0.65	0.94	*0.02*	**0.19**	**0.44**	***<0.01***

Note: Bold numbers indicate that the entry entails a feasible punishment length in the GP equilibrium. The results are qualitatively similar if all rounds (1-33) are considered. Consistent with theory, the public signal assumed to trigger a price war is p_2.
* Pearson's Chi-Square statistic; p-values of other non-parametric tests (Wilcoxon, Kolmogorov-Smirnov, and Epps-Singleton) and the parametric t-test produce similar p-values.

Table D.8: Fraction of Times the Equilibrium Path Correctly Predicts Outcomes (Predictive Power)

Equilibrium Path	**Parameterization 2**
(H,H) every period	**37.98%**
GP_3	**40.40%**
GP_4	**50.51%**
GP_5	**60.61%**
GP_{14}	**75.96%**
GP_{15}	**78.18%**
GP_∞	**81.21%**

Notes: Bold numbers indicate a theoretically feasible equilibrium (see Appendix B, table B.1 for details). The GP_N path takes a value of 1 when collusion is predicted and 0 when a price war is predicted; a price war is assumed to be triggered by a low signal (price$\leq p_2$) which lasts N periods.

Table D.7: Probit Estimates of Different Strategies in the IM treatment, Rounds 1-25, Parameterization 2b

	Random	GP_N			$T1_N$ $\{N\}$			$TT1_N$ $\{N\}$	$T2$		$TT2$	
		N=5[1]	N=13	N=∞	$k = p_1$ [2] $\{5\}$	$k = p_1$ [2] $\{\infty\}$	$k = p_0$ [2] $\{\infty\}$	$k(L) = p_1$, $k(H) = p_0$ [3] $\{\infty\}$	$k^{down} = p_1$, $k^{up} = p_3$	$k^{down} = p_1$, $k^{up} = p_4$	$k^{down}(L) = p_1$, $k^{down}(H) = p_0$; $k^{up}(L) = p_4$, $k^{up}(H) = p_2$	$k^{down}(L) = p_1$, $k^{down}(H) = p_0$; $k^{up}(L) = p_4$, $k^{up}(H) = p_3$
γ	N/A	0.39*	0.50*	0.61*	0.63*	0.92*	0.48*	0.60*	0.90*	0.92*	0.55*	0.62*
LL	-437.66	-431.13	-431.1	-430.0	-420.4	-415.9	-430.98	-425.8	-416.0	-415.9	-426.90	-424.89
LR†	N/A	13.04	13.15	15.28	34.38	43.35	13.35	23.68	43.23	43.35	21.52	25.54
p-value		<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01

Notes: Estimates of α, ρ and ψ are significant at the 1% level in all specifications (not shown). Number of observations: 1,200 in both parameterizations. N=number of punishment periods. LL= Log-likelihood. GP=Green and Porter theoretical prediction; strategies T1, TT1, T2 and TT2 are defined in table 11. k are threshold levels, also defined in table 11.

[1] GP equilibrium is feasible for punishment lengths 3-∞ (2-∞ for the risk aversion case). Shown are only the three punishment lengths with the highest explanatory power.

[2] Other finite punishment lengths also had statistical significance, but were less significant than the one reported.

[3] Strategies with threshold levels $k(H) = p_2$ and $k(H) = p_1$ also had statistical significance, but were less significant than the one reported.

† Likelihood ratio test with respect to the random strategy

* Significant at 1% level