INFORMATION AND COST ASYMMETRY IN EXPERIMENTAL DUOPOLY MARKETS

Charles F. Mason and Owen R. Phillips*

Abstract—We analyze data from experimental duopoly markets to assess the role information plays in facilitating collusion. In these markets, profitability can be common knowledge or private information. Market outputs are estimated in structures with symmetric and asymmetric costs under the two information conditions. Symmetric markets are more cooperative when profitability is common knowledge; asymmetric market outputs are unaffected by information differences. However, common knowledge in asymmetric markets increases the share of the output produced by the low-cost producer, and therefore increases industry efficiency.

I. Introduction

THERE is substantial public policy concern over oligopoly firms having information on the prices, costs, and sales of rivals. It is generally understood that there are some benefits from the availability of such information. For example, the publicized prices of vendors can reduce price dispersion, which is beneficial to buyers who are searching for the lowest price (Stigler (1961)). But it is also recognized that when markets are concentrated, the availability of such information may facilitate collusion, by revealing focal points for collusive arrangements between firms (Schelling (1960)) and by homogenizing perceptions of reasonable outcomes (Clarke (1983) and Grether and Plott (1984)). Once an agreement is reached, improved information helps rivals monitor each other (Stigler (1964)); the capability to detect defection is crucial to a stable agreement (Green and Porter (1984)).

Several game-theoretic models have investigated the impact of information when players are in a noncooperative setting. A representative list includes Novshek and Sonnenshein (1982), Vives (1984), Li (1985), Gal-Or (1985), Shapiro (1986), and Kirby (1988). Agents in these models know their own payoff functions but they do not necessarily know the payoffs of other agents, and payoff functions need not be the same. This literature typically analyzes how making payoffs common knowledge influences actions. Agents are duopoly producers, which makes the models computationally convenient and gives firms substantial market power. Also, the games are of a multiple-stage, one-period variety, with a focus on the noncooperative equilibrium. In this literature, the effect on equilibrium behavior from an improvement in information depends on

the nature of the uncertainty (demand or cost), the choice variable (prices or quantities), the cross-price relation between the products (complements or substitutes), and the degree to which payoffs are nonlinear in one's own output. In contrast to the one-period models, Green and Porter's (1984) seminal work analyzes the effect of payoff uncertainty in a repeated game. While some collusion can still be supported when payoffs are uncertain, these outputs are larger than the cooperative outputs that could be supported when payoffs are common knowledge.

In this paper we empirically investigate the impact of improved information in experimental duopoly markets. We contrast behavior in markets with common knowledge on all payoffs, a treatment we also refer to as full information, against behavior in markets where players know only their own payoffs, a condition we call private information. Our experimental design has two types of players: one with low costs and one with high costs. In our first set of treatments, payoffs are common knowledge, and we estimate steady-state choices in two market structures, one where payoff functions are the same for both players and another where they are different. In a second set of treatments, payoffs are private information. Here again we evaluate outcomes when payoffs are the same and when they are different. When agents have the same payoff functions, they operate in a symmetric market structure, much like a duopoly producing a homogeneous good with the same cost function. When subjects have different payoff functions, they operate in an asymmetric market structure. This duopoly structure has a relatively large and a small seller because one firm has a cost advantage.

The theoretical models cited above provide useful insights, but they carry assumptions that may limit their ability to predict the influence of information in actual markets. These assumptions are (1) agents interact only once and (2) the probability distribution function governing the source of uncertainty is known. In reality, interactions are likely to be better represented as a repeated game, and the relevant probabilistic rules determining a rival's profitability are unlikely to be understood. Because we would like our experiments to reflect likely effects in naturally occurring market environments, we do not invoke assumptions (1) and (2) above. Rather, our subjects play against the same opponent for the duration of the experiment, which runs for an indeterminate length; this mimics a supergame. Furthermore, subjects are given no probabilistic information on their rivals' payoffs in the private information treatments, which resembles the realistic absence of information.¹

Received for publication June 21, 1994. Revision accepted for publication May 10, 1995.

* University of Wyoming.

We thank seminar participants at the University of Michigan for lively discussion, and two anonymous referees for helpful comments. This material is based on work supported by the National Science Foundation under grant RII-8610680 and the College of Business, University of Wyoming. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies.

¹ Our experimental design is similar to that used by Fouraker and Siegel (1963). Their experiments used payoff tables based on linear demand and

Our results show that information has a greater impact on duopolies that have symmetric payoffs than on those with asymmetric payoffs. Applying maximum-likelihood techniques, we find a significant behavioral difference in symmetric markets under the two types of information. In particular, symmetric markets are more cooperative when payoffs are common knowledge. By contrast, subject choices in the asymmetric treatments are well explained by Nash behavior under both information treatments. We also present two sets of results that support the hypothesis that improved information is welfare enhancing. First, a significant reduction in the length of time it takes subjects to stabilize is associated with common knowledge, as compared to private information. This is consistent with Stigler's (1961) argument that improved information stabilizes markets. Second, in the asymmetric sessions we find a significant increase in low-cost subject outputs, and a significant decrease in high-cost subject outputs. This result corroborates Shapiro's (1986) prediction that making quantity-choosing firms' cost information common knowledge will rationalize production by shifting output from higher cost firms to those with lower costs.

The remainder of the paper is organized as follows. In section II we discuss the role information sharing can have as a collusion-facilitating device, and we analyze other theoretical motives for information sharing. Section III describes our experimental treatments and the data. Section IV contains our econometric analysis and the statement and testing of various hypotheses on the behavior for the two market structures with either common knowledge or private information. Section V offers some concluding remarks.

II. Information and Collusion

It is well known that outcomes more cooperative than the Nash equilibrium in the underlying one-shot stage game can be supported as a subgame perfect equilibrium in infinitely repeated games, or in games with unknown and randomly determined end points (Fudenberg and Tirole (1991), Kreps (1990), and Rasmusen (1989)). Equilibria in these types of

games are based on trigger strategies. Firms select a cooperative output level until defection is detected; if defection occurs, a "punishment" phase ensues (Abreu (1986)). Punishment may last indefinitely, entailing perennial play at the Cournot-Nash level (i.e., firms play the "grim reaper" strategy), or punishment could occur for a finite number of periods, with firms then reverting to their cooperative outputs (i.e., firms play the "carrot-and-stick" strategy). Common knowledge of payoffs is important in this context, since it reveals output combinations that can be supported as an equilibrium, and because it facilitates the determination of each agent's incentives to participate in a tacit agreement.

While common knowledge of payoffs makes collusion easier, it is not required. Green and Porter (1984) show that some collusion can result when demand is stochastic, and therefore payoffs to both agents are random. A cooperative output level is maintained until price falls below some critical level, either because demand is unusually low or because some firm has cheated on an implicit agreement. At that point, firms select the noncooperative output for a fixed period of time, and then revert to the cooperative regime. In this story cheating cannot be directly observed; otherwise the punishment phase would only follow cheating, and not periods of low demand. Since low prices may come from either low demand or cheating, the frequency of losses to a cooperating rival becomes relatively high, and therefore firms have an incentive to play less cooperatively than in a scenario where low prices can only occur because of cheating, e.g., when payoffs are common knowledge. This more conservative play implies that when payoffs are private information, market outputs will be larger and prices will be lower.

Extending the Green and Porter model to situations where costs are private information and individuals choose quantity requires that the trigger strategies be conditioned on something besides market price. Stigler (1964) suggests that firms may use their market share as the triggering device.² Each firm chooses some cooperative level, presumably reflective of the expected value costs can take, so long as the firm's market share does not fall below a critical level. If the firm's share drops below this threshold, then the firm produces a high level of output for some punishment period.

With asymmetric costs, market shares will differ. In contrast to Green and Porter's model, the high-cost and low-cost firms experience differential gains from cheating on an implicit agreement, which implies that agents in the asymmetric markets have relatively different incentives to cooperate. Nevertheless, an abnormally low market share for one agent results from unexpectedly large output choices by one or more other agents, so that lower shares result if a rival firm defects from the cooperative regime. Therefore, for either price choosers or quantity choosers, firms have less incentive to play cooperatively without full information, so

constant symmetric marginal costs for a homogeneous good. Subjects made quantity choices for 22 periods without knowing the last period until period 21. While data were not analyzed as a time series, median choices by period do not appear significantly different between full and private information, nor do they appear different from the Nash equilibrium (see Smith et al. (1990)). However, the percentage of subject pairs at the Cournot equilibrium with private information are more than twice the percentage under full information (39% versus 81%). Dolbear et al. (1968) also studied the importance of information in duopoly. Subjects chose prices for goods that were not perfect substitutes. Marginal costs were constant. The same holds for sellers. Experiments ran 15 periods, with data reported for periods 8 through 12. During this subset of periods, there is no statistical difference in average choice behavior between the six duopoly markets with private information and the twelve pairs with full information. Both averages appear close to the Bertrand equilibrium solution. Both the Fouraker and Siegel and the Dolbear et al. results suggest that information does not affect choices. However, while reliance on median choices provides unbiased evidence on central tendencies, it is statistically inefficient. Similarly, using a subset of choices, instead of the entire sample, can taint results (Alger (1987)).

² Indeed, Green and Porter (1984) state that Stigler (1961) provided much of the motivation for their analysis.

the observed outputs will be larger than the collusive outputs supported when payoffs are common knowledge. Improved information about profits can engender more collusive behavior because cheating is less likely to occur. One then expects that observed market outputs will tend to be smaller when payoffs are common knowledge.

Effects other than the facilitation of collusion may result when firms obtain information about a rival's payoffs. Shapiro (1986) considers a model with linear demand, homogeneous products, and quantity-choosing duopolists whose costs are initially private information. When firms share information on their costs in this context, total surplus rises because of increased production efficiency and a smaller variance in outputs. Lower industry costs are realized because the mix of output shifts toward lower cost producers. When firms know they are relatively efficient sellers, they are induced to produce more; high-cost sellers have an incentive to produce less in order to maximize profits. Also, since firms know more about their relative efficiency, production mistakes are less likely to occur, causing the variance in output to fall.

Altogether, these remarks point to welfare ambiguities associated with increased information about the payoffs of rivals. Information can facilitate collusion, which generally reduces total surplus. On the other hand, information can increase welfare by lowering the variance in outputs and by shifting output toward the lower cost producer. Policy treatment toward information sharing arrangements surely depends on the degree to which improved information facilitates collusion. If information generally facilitates collusion in oligopoly, then regardless of the reasons firms share information, and regardless of the presence or absence of explicit price-fixing discussions, the appropriate policy approach should be hostile. If instead information is not a facilitating device for collusion, but has desirable welfare effects, a passive approach toward its dissemination would be in order. Given the theoretical ambiguities addressed above, resolving policy concerns becomes an empirical question. We submit that a natural way to analyze the influence of information on cooperation is by observing behavior in laboratory markets.

III. Experimental Markets and Data

Subjects in our experimental design make output choices from a payoff table that represents the strategic form of the game. This design is motivated by general game-theoretic models of oligopoly for which buyers have no market power, and so are passive, but sellers can have substantial market power.³ The use of payoff tables in experiments,

³ In addition to the various papers cited in the introduction, see Abreu (1986), Friedman (1983), Fudenberg and Tirole (1991), Kreps (1990), Kreps and Spence (1985), Rasmusen (1989), and Schelling (1960). Industries where buyers possess some market power are not well represented by payoff matrix games. Experimental representations of such markets generally include both buyers and sellers, e.g., by appealing to double oral auctions (see Davis and Holt (1993) or Friedman and Sunder

however, has a history that predates their description of oligopoly markets. Extensive surveys of earlier papers that used payoff tables to study rivalry or bargaining are contained in Colman (1982) and Rapaport et al. (1976). A more recent discussion of papers that investigate noncooperative behavior in general, and oligopoly behavior in particular, through the use of payoff tables is provided by Davis and Holt (1993, chap. 2) and Friedman and Sunder (1994, chap. 9).

The role of information about a rival's payoffs in determining behavior has not been given extensive study. While Fouraker and Siegel (1963) and Dolbear et al. (1968) provide pioneering work, more recent work analyzing the effect of information on the behavior of duopoly agents, in particular when there are asymmetric costs, is virtually nonexistent.⁴ In an earlier study (Mason et al. (1992)) we found that payoff symmetry was a powerful facilitator of collusive behavior. Since many naturally occurring oligopolies are characterized by both payoff asymmetry and uncertainty about rivals' payoffs, analyzing the interaction between these two structural elements is likely to enhance the understanding of oligopoly behavior. Data analysis in previous experimental studies of oligopoly tends to focus on summary statistics during the middle or end of the experiment; in particular this is true of Fouraker and Siegel (1963) and Dolbear et al. (1968). Friedman and Sunder (1994, p. 137) argue that convergence paths are more important measures of behavior than these relatively unsophisticated techniques. Because past work has not focused on convergence paths, for example, by using the time-series techniques we propose below, it is plausible that subtle differences in behavior across information treatments have gone undetected.

The payoff functions in our two-person repeated games are constructed using a linear-quadratic framework. The static inverse demand function is

$$P(q_1, q_2) = 4 - (q_1 + q_2)/24 \quad (1)$$

where agents choose quantities q_i , $i = 1, 2$, and P is the price of a homogeneous good. Marginal costs are 0 for low-cost sellers and $1/2$ for high-cost sellers. Thus the static payoff for player i , π_i , may be calculated from

$$\pi_i(q_1, q_2) = \begin{cases} P(q_1, q_2)q_i, & i \in L \\ [P(q_1, q_2) - 0.5]q_i, & i \in H \end{cases} \quad (2)$$

(1994) for details). Finally, because we are interested in assessing the effect of information on behavior, our control design needs to make subjects' payoffs common knowledge. The most efficient way to do this is to represent payoffs as a reduced form of the underlying demand and cost conditions, i.e., to show them the strategic form of the game. Placing subjects in experimental duopolies is a simple way to give them substantial market power.

⁴ Notable examples include Cason (1994), Cason and Mason (1995), and Cooper et al. (1992). The first two analyze the impact of demand uncertainty on duopoly behavior when costs are symmetric, whereas the third paper considers the importance of information in coordination games.

where L (H) is the set of low (high)-cost producers. These functions are presented to subjects in the form of payoff tables which show the dollar profit accruing from various output combinations. We use the linear-quadratic profit functions in equation (2) because they facilitate the computation of equilibria under the various cost combinations, which simplifies our exposition. In addition, the specific payoff tables we use provide a wide separation between the static Cournot–Nash outputs and the static collusive equilibria.⁵ Finally, payoff uncertainty can be introduced into our design in a very simple way, as we indicate below.

Subjects are paired in one of three market designs. If both players are low-cost producers, we label the duopoly LL. If both are high-cost producers, the market is called HH, and duopolies with asymmetric cost are termed LH. In a session subjects faced one of two information conditions. In one information treatment, all subjects were given the payoff table both for themselves and for their rival, and we announced that all subjects possessed the payoff table of their counterpart.⁶ This was the common knowledge, or full information, condition. In a second treatment, subjects were not told anything about the payoff table of their counterpart.⁷ This was the private information condition. In total, there are three experimental market structures, two having symmetric payoffs and one asymmetric payoffs, and two information treatments.

All subjects were recruited from beginning economics classes at the University of Wyoming. They reported to a reserved classroom, where a set of instructions was read aloud as they followed along on their own copy. Questions were taken, and one practice period was conducted. In the

practice period, subjects simultaneously chose a row value from a sample payoff table. Then each subject calculated and recorded the profits found at the intersection of the subject's row and a specified column. Each person was checked during the practice period to ensure that everyone understood a payoff table and how to keep a record of their choices and earnings.

Even though all choices were in a quantity-choosing market environment, subjects were simply told they were simultaneously choosing an X row value. For both the low-cost and the high-cost table, earnings were presented in a fictitious currency called tokens.⁸ At the end of the experiment, tokens were exchanged for cash at the rate of \$1.00 for every 1000 tokens.

Once the experiment got under way, each person was assigned an anonymous counterpart in another room. Subjects were told they would be paired with the same person for the duration of the experiment. In a market period, a subject wrote his or her output on a record sheet and a colored piece of paper. A monitor then exchanged the pair's colored slips, after which each player calculated earnings from the payoff table. Subjects had many more record sheets and colored slips of paper than were required for a session. The number of periods and the exact time an experiment would run were not known by any of the participants during a session. Subjects were recruited for a length of time 30 to 45 minutes greater than we expected an experiment to actually run.⁹

Every subject was given a starting cash balance of \$3.00 to cover potential losses, and was told that if their balance went to zero, they would be dismissed from the experiment with a \$5.00 participation fee. The remaining participant would then be allowed to operate as an unfettered monopolist. Although it was feasible for low-cost players to dispatch a high-cost opponent without suffering losses, such predatory behavior never occurred in any of our sessions; indeed, the lowest recorded cumulative earnings were near \$17.00. Some low-cost subjects earned as much as \$37.00 in a 2-hour setting.

We analyze data from 12 experimental sessions. Eight of these sessions had common knowledge on payoffs, and four had private information on payoffs. In the eight common-knowledge sessions there were a total of 130 subjects, or 65 subject pairs. In these treatments, three sessions with 25 subject pairs had asymmetric payoff functions, and five sessions with 40 subject pairs had symmetric payoffs. In the four private-information treatments there were 74 subjects, or 37 subject pairs. Here there were two sessions with

⁵ It is easy to verify that Cournot outputs for a model with inverse demand $P = A - BQ$ and marginal costs c_1, c_2 are $q_i = (A - 2c_i + c_j)/3B, j \neq i = 1, 2$. With $A = 4, B = 1/24$, and $c_1 = c_2 = 0$, these outputs are $q_1 = 32 = q_2$; $q_1 = 36$ and $q_2 = 24$ if $c_1 = 0, c_2 = 1/2$; and $q_1 = 28 = q_2$ if $c_1 = c_2 = 1/2$. Thus the static noncooperative Cournot–Nash equilibria entail each firm producing 32 in the LL design, for a total of 64; the low (high)-cost firm producing 36 (24) in the LH design, for a total of 60; and each firm producing 28 in the HH design, for a total of 56. In the symmetric designs, where $c_1 = c_2 = c$, the fully collusive regime entails the paired outputs summing to the monopoly output; this equals 48 when $c = 0$ and 42 when $c = 1/2$. With equal costs, the natural division of output entails $q_1 = q_2$. In general, the collusive regime is described by the locus of iso-profit tangencies. Using our parameters in Schmalensee (1987, eq. (6)), this locus is given by $q_2 = 69 - q_1 - [(729 - 6q_1)^{1/2}]/6$. Profits are maximized subject to the constraints that each subject's static collusive profits exceed static Cournot profits, which yields a market output of 45 after rounding. This may be further refined to $q_1 = 28$ and $q_2 = 17$ (after rounding) by utilizing either the Nash bargaining solution, Friedman's (1983) concept of balanced temptation, or several of the alternative collusive concepts discussed in Schmalensee (1987).

⁶ If payoffs were symmetric, one payoff table provided information for both agents. In this case, each participant received the one table, and we announced and demonstrated that either player's payoffs could be calculated from the one table.

⁷ In particular, subjects were not given any information on the distribution of payoffs a rival could earn. This absence of common priors contrasts with the analytical assumptions in the already cited theoretical literature. This is in keeping with our argument that firms in natural markets probably do not have good information about the profitability of rivals. This complicates the firm's decision problem, but Kalai and Lehrer (1991) have shown that repeated interaction in such an environment will converge to a Nash equilibrium.

⁸ Copies of the low- and high-cost payoff tables are available upon request from the authors. Rounding the token payoffs in some cases generated the same profits over two or three row choices in a column. Because of this, the unique analytical Cournot–Nash solution is not necessarily unique in the payoff table.

⁹ The actual end of the experiment was randomly determined. In keeping with our goal of extending Fouraker and Siegel (1963) and Dolbear et al. (1968), we adopt their convention of not discussing the rule used to terminate the session in the experimental instructions.

asymmetric payoffs with a total of 18 subject pairs, and two sessions with symmetric payoff functions with 19 subject pairs. Every experiment lasted at least 35 periods, with some going as long as 46 periods. Graphic descriptions of the data are presented in Mason and Phillips (1994b).

There are some notable choice patterns. Specifically, average outputs as a percentage of the *market Cournot output* in the symmetric (LL and HH) full-information sessions are below outputs in the symmetric private-information sessions for more than 80% of the choice periods. Average behavior is similar in the HH and LL market structures in both the common-knowledge and private-information treatments. In light of this similarity, we combine the LL and HH observations into a single symmetric data set in the econometric analysis conducted below.¹⁰ Private-information outputs are generally near the Cournot output levels for all market structures. The information treatment seems to have no impact on *average* choice behavior when there is asymmetry in the market structure. In general, outputs in all sessions tended to fall over the course of the experiment.¹¹

IV. Econometric Analysis

We use three distinct econometric models to evaluate the impact of information in the symmetric and asymmetric market structures. In the first model, we interpret the data set as a pooled cross-section time-series sample, where the dependent variable is subject pair output. The focus here is on testing for differences in ultimate market output between information treatments for each market structure. In the second model, we conduct a time-to-failure analysis. In this model we determine the influence private information has on the length of time it takes choices to stabilize. The third model analyzes individual outputs in the asymmetric sessions. Our aim here is to contrast the behavior of subjects when there is full information on payoffs against behavior when payoffs are private information, in order to test Shapiro's (1986) prediction that making agents' payoffs common knowledge has the effect of rationalizing outputs.

¹⁰ Outputs in the LL full-information treatment are below outputs in the LL private-information treatment for 29 of the first 35 periods. Average outputs in the HH treatment with full information are below the private-information outputs for 32 of these 35 periods. While the HH and LL sessions exhibit some slight differences during the middle periods, they trend together (converging toward 85% of the Cournot output in the full-information sessions and 95% of the Cournot output in the private-information sessions). Indeed, HH and LL behavior is statistically indistinguishable (cf. footnote 13).

¹¹ Average subject pair output, as a percentage of the associated Cournot level, ranged from 1.07 to 0.83 during the first 35 periods of the experiment; standard deviations ranged from a high of 0.12 to a low of 0.012. This phenomenon of emerging cooperation is common in oligopoly experiments. We believe it reflects the subjects' learning about their rivals' strategies (Mason and Phillips (1994a)) and the signaling of collusive desires (Kreps and Spence (1985)).

A. Paired Output as the Dependent Variable

In this subsection we regard our data set as a pooled cross-section time-series sample, and so we require an equal number of observations from each pair. As the shortest experimental session went 35 periods, we consider the first 35 observations on each pair from all experiments. With 65 subject pairs in the common-knowledge treatments, this gives us 2275 data points; with 37 subject pairs in the private-information treatments, we have 1295 data points.

The structural model we utilize in our analysis is

$$Q_k(t) = \alpha_k + \epsilon_k(t) \quad (3)$$

where $Q_k(t)$ is pair k 's period t output as a fraction of the associated Cournot output, α_k is the steady-state output (i.e., the output the pair will ultimately converge to), and $\epsilon_k(t)$ is a residual capturing variations about the equilibrium. We assume the disturbance follows a second-order autoregressive process, so that¹²

$$\epsilon_k(t) = \rho_{k1}\epsilon_k(t-1) + \rho_{k2}\epsilon_k(t-2) + \mu_k(t) \quad (4)$$

where the residual $\mu_k(t)$ is white noise (i.e., $E\mu_k(t) = 0$, $E\mu_k(t)\mu_k(s) = 0$ for $t \neq s$, and $E\mu_k(t)^2 = \sigma_k^2$). Equations (3) and (4) may be combined to yield

$$Q_k(t) = \beta_k + \rho_{k1}Q_k(t-1) + \rho_{k2}Q_k(t-2) + \mu_k(t) \quad (5)$$

where $\beta_k = \alpha_k(1 - \rho_{k1} - \rho_{k2})$. Dynamic stability requires that $|\rho_{k1}| < 1$, $|\rho_{k2}| < 1$, and $|\rho_{k1} + \rho_{k2}| < 1$ (Fomby et al. (1988)). Dynamic stability allows us to interpret α_k as the steady-state market output, and so it is the natural parameter to focus on when asking questions about ultimate behavior.

To test for differences in the steady-state behavior of subjects between the two information treatments, we assume that α_k , ρ_{k1} , and ρ_{k2} are the same for all subject pairs in a given design, and that any variation across pairs is captured by the respective residual terms. We allow for variance to differ across subject pairs, but assume that there is no cross-equation covariance: $E[\mu_k(t)\mu_h(s)] = 0$ for $k \neq h$. We denote treatment 1 as common knowledge with symmetric costs, treatment 2 as common knowledge with asymmetric costs, treatment 3 as private information with symmetric costs, and treatment 4 as private information with asymmetric

¹² There are a host of reasons to expect serial correlation in this structure. Each subject's output in period t is likely to be linked to both subjects' outputs in $t-1$ (Friedman (1983)). Any attempts at signaling a desire to collude hinge on an intertemporal connection (Shapiro (1980)). Similarly, any learning implies a connection between current and preceding outputs. Taken together, these argue for a time-series structure in the $\epsilon_k(t)$, which is likely to be more complicated than an AR(1) process. Indeed, in a similar experimental setting Mason et al. (1991) argue that this time-series structure is best described by an AR(2) process.

TABLE 1.—PARAMETER ESTIMATES UNDER DIFFERENT TREATMENTS

Treatment	Parameter	Estimate	Standard Error
1. common knowledge, symmetric payoffs	α_1	0.8441	0.0098
	ρ_{11}	0.3772	0.0254
	ρ_{21}	0.3419	0.0251
	$R^2 = 0.9281$		
	Durbin-Watson statistic = 2.1531		
2. common knowledge, asymmetric payoffs	α_2	0.9807	0.0056
	ρ_{12}	0.2858	0.0330
	ρ_{22}	0.2168	0.0318
	$R^2 = 0.9865$		
	Durbin-Watson statistic = 2.0847		
3. private information, symmetric payoffs	α_3	0.9702	0.0062
	ρ_{13}	0.2008	0.0394
	ρ_{23}	0.0927	0.0365
	$R^2 = 0.9190$		
	Durbin-Watson statistic = 2.0307		
4. private information, asymmetric payoffs	α_4	1.0177	0.0083
	ρ_{14}	0.1998	0.0391
	ρ_{24}	0.1780	0.0387
	$R^2 = 0.8880$		
	Durbin-Watson statistic = 2.0227		
Test statistic on $H_0: \alpha_1 = \alpha_3 = 10.874$; 5% critical value = 1.96			
Test statistic on $H_0: \alpha_2 = \alpha_4 = 1.3437$; 5% critical value = 1.96			

costs. The assumed parameter vectors are thus¹³

$$(\alpha_k, \rho_{k1}, \rho_{k2}) = \begin{cases} (\alpha_1, \rho_{11}, \rho_{12}) & \text{if pair } k \text{ is in treatment 1} \\ (\alpha_2, \rho_{21}, \rho_{22}) & \text{if pair } k \text{ is in treatment 2} \\ (\alpha_3, \rho_{31}, \rho_{32}) & \text{if pair } k \text{ is in treatment 3} \\ (\alpha_4, \rho_{41}, \rho_{42}) & \text{if pair } k \text{ is in treatment 4.} \end{cases}$$

This system may be efficiently estimated by feasible generalized least squares, yielding the estimator vector (b_p, r_{p1}, r_{p2}) for $(\beta_p, \rho_{p1}, \rho_{p2})$, $p = 1, 2, 3, 4$. We may then consistently estimate α_p by

$$a_p = b_p / (1 - r_{p1} - r_{p2}). \quad (6)$$

The results of this estimation procedure are given in table 1. We report for each experimental treatment the estimates a_p , r_{p1} , and r_{p2} , along with their standard errors. We observe that in each treatment, the conditions for dynamic stability are met, so that a_p may be regarded as the maximum-likelihood

¹³ The symmetric subsets combine LL and HH subject pairs. This is acceptable if the behavior of the two groups is indistinguishable. To test this, we ran a variant of the regression based on equation (5), where the LL and HH pairs were grouped together, and a dummy variable was included for the HH pairs. We conducted this analysis for each of the information treatments. If the intercept term in equation (5) does not differ systematically between LL and HH treatments, the coefficient on this dummy variable will be insignificant. For the full-information sessions, the coefficient on the HH dummy is 0.0053 and has a t -statistic of 0.8854. For the private-information sessions, the coefficient on the HH dummy is 0.0241 and has a t -statistic of 0.7090. Thus we accept the hypothesis that LL and HH subjects behaved alike, for both informational treatments.

estimators of the ultimate market outputs α_p in each of the four treatments. The two hypotheses we wish to test are:

HYPOTHESIS 1: When payoffs are symmetric, ultimate subject behavior is not significantly different when costs are common knowledge than when costs are private information, i.e., $\alpha_1 = \alpha_3$.

HYPOTHESIS 2: When payoffs are asymmetric, ultimate subject behavior is not significantly different when costs are common knowledge than when costs are private information, i.e., $\alpha_2 = \alpha_4$.

Each hypothesis may be tested by comparing the t -statistic for the relevant difference ($\alpha_1 - \alpha_3$ for hypothesis 1 and $\alpha_2 - \alpha_4$ for hypothesis 2) against the critical point on a student's t -distribution.

The test statistics for both hypotheses are reported at the bottom of table 1. These test statistics indicate that the null hypothesis—that the information treatment does not affect subject behavior—can be rejected with great confidence when payoffs are symmetric, but cannot be rejected when payoffs are asymmetric. Improved information facilitates more cooperative behavior only when payoffs are symmetric.

B. Time as a Dependent Variable

In this subsection we interpret a subject pair's output as having stabilized when their outputs do not vary by more than three units in either direction for the remainder of the experiment.¹⁴ Previous studies of similar issues have typically used either contingency table analysis or time-to-failure analysis.

To apply the contingency table methodology, we tabulate the number of subject pairs whose choices have stabilized by the end of the time frame we used in our regression analysis of section IVA, period 35, for each of the four combinations of information and market structure. The hypotheses of interest in this approach are:

HYPOTHESIS 3: When payoffs are symmetric, the frequency of output stabilization is not significantly different when costs are common knowledge than when costs are private information.

¹⁴ In general, the effect on payoffs when an agent changes his or her output slightly is relatively small. For example, if agent i increases output one to two units above the Cournot equilibrium level, payoffs will decrease on the order of two tokens. Correspondingly, a small change in an agent's output could occur as his or her choices converged; the associated aggregate change in output of one to four units would remain small. By contrast, if an agent altered his or her output by four units or more, that would result in a substantially larger change in payoffs. Our notion of convergence is meant to capture the former example of trembles, while ruling out the latter. To this end, we settled on an aggregated change of no more than three units in either direction.

TABLE 2.—CONTINGENCY TABLE ANALYSIS: RELATION BETWEEN INFORMATION AND NONCONVERGENCE TO STABLE CHOICES

Treatment	Information Treatment	Convergence to Stable Choice		Column Sum
		Yes	No	
Symmetric	Full	29	11	40
	Private	6	13	19
	Row sum	35	24	
	Test statistic for hypothesis that information has no effect on convergence = 8.9388.			
Asymmetric	Full	14	11	25
	Private	2	16	18
	Row sum	16	29	
	Test statistic for hypothesis that information has no effect on convergence = 12.8346; 1% critical value = 6.63			

HYPOTHESIS 4: When payoffs are asymmetric, the frequency of output stabilization is not significantly different when costs are common knowledge than when costs are private information.

The relevant data are contained in table 2. Each cell of the contingency tables presented here shows the number of subject pairs for that combination of information treatment and stabilization or nonstabilization. We present contingency tables for both symmetric markets (part 1) and asymmetric markets (part 2). For each symmetry treatment, the test statistic will be distributed as a chi-squared variate with one degree of freedom if the null hypothesis is true. For both the symmetric and the asymmetric designs, the test statistic reported in table 2 greatly exceeds the critical point. We therefore reject both hypothesis 3 and hypothesis 4 in favor of the alternative hypothesis that improved information increases the chance that outputs will stabilize, in both the symmetric and the asymmetric cost designs.

To apply the time-to-failure analysis, we define the failure time for each subject pair as the period when their outputs first stabilize. The *hazard rate* is then the probability a pair's choices stabilize in the period, given that they have not yet done so. Statistical analysis proceeds by modeling the hazard rate as a function of a set of explanatory variables. We assume the random component of failure times follows a Weibull density function, which implies that the hazard rate changes monotonically over time. The explanatory variables we use include four dummy variables, listed as FI_S, FI_A, PI_S, and PI_A in table 3. These variables represent a combination of information treatment (FI versus PI for full versus private information) and market condition (S versus A for symmetric versus asymmetric); each dummy equals 1 if the subject pair was in the associated treatment, and 0 otherwise. The model also includes a shape parameter, listed as σ in table 3, which is included to allow for time trends in the hazard rate.¹⁵ An estimated coefficient which is signifi-

¹⁵ The Weibull model subsumes a third alternative, the exponential time-to-failure model. The exponential model is appropriate when the hazard rate is constant. This may be tested by comparing the shape parameter σ with 1. Values of σ that are less (greater) than 1 imply that the

TABLE 3.—REGRESSION RESULTS FROM WEIBULL TIME-TO-FAILURE MODEL

Regressor	Parameter	Estimate	Standard Error
FI_S	α_1	3.5808	0.0732
FI_A	α_2	3.5860	0.1038
PI_S	α_3	3.9331	0.1675
PI_A	α_4	4.3852	0.2952
		0.3868	0.0510
	$\alpha_3 - \alpha_1$	0.5808	0.2940
	$\alpha_4 - \alpha_2$	0.7992	0.3096
Test statistic on $H_0: \alpha_1 = \alpha_3 = 1.9756$; 5% critical value = 1.96			
Test statistic on $H_0: \alpha_2 = \alpha_4 = 2.5815$; 5% critical value = 1.96			
Test statistic on $H_0: \alpha_1 = \alpha_3 \text{ \& } \alpha_2 = \alpha_4 = 9.8451$; 5% critical value = 6.99			

cantly greater than zero indicates that increases in the associated explanatory variable lead to a greater expected failure time, i.e., a delay in the anticipated time when convergence to a stable choice pattern occurs.

The hypothesis of interest is that there are no significant differences in the time it takes subjects' outputs to stabilize between the two information treatments. We test this in three ways:

HYPOTHESIS 5: When payoffs are symmetric, the time it takes outputs to stabilize when costs are common knowledge is not significantly different from the time it takes when costs are private information.

HYPOTHESIS 6: When payoffs are asymmetric, the time it takes outputs to stabilize when costs are common knowledge is not significantly different from the time it takes when costs are private information.

HYPOTHESIS 7: The time it takes outputs to stabilize when costs are common knowledge is not significantly different from the time it takes when costs are private information, regardless of symmetry in payoffs.

Hypotheses 5 and 6 are tested by an asymptotic *t*-test on the difference of the parameter estimates; hypothesis 7 is tested using a likelihood-ratio test, for which the appropriate test statistic is distributed as a chi-squared variate with two degrees of freedom under the null hypothesis. The results of our Weibull time-to-failure analysis are reported in Table 3. The relevant test statistics are presented at the bottom of the table for each of the three hypotheses; all three hypotheses can be rejected at conventional levels. For both symmetric and asymmetric treatments, the coefficient on the common-knowledge dummy variable is smaller than the coefficient on the private-information dummy variable. We conclude that it takes pairs longer to stabilize in the private-information treatment than in the common-knowledge treatment. Since making payoffs common knowledge has little effect on outputs in asymmetric markets and it reduces the variance of

hazard rate is increasing (decreasing) over time. The estimated value of σ in table 3 is 0.3868. It is significantly less than 1 at the 1% level, which indicates that the hazard rate in our data set is increasing over time.

choices, improved information enhances asymmetric market performance.

C. Individual Outputs in Asymmetric Designs

Our final analysis focuses on *individual* output behavior in the asymmetric market structure, where we test the impact information has on high- and low-cost agents' behavior. The two hypotheses of interest in this subsection are:

HYPOTHESIS 8: Ultimate outputs for low-cost agents do not significantly differ between information treatments.

HYPOTHESIS 9: Ultimate outputs for high-cost agents do not significantly differ between information treatments.

The alternative hypotheses are that low-cost agents make significantly larger output choices, and high-cost agents make significantly smaller choices when payoffs are common knowledge.

The econometric model we employ to test these hypotheses is closely related to that in subsection A. An individual's output in period t is assumed to be related to the rival's output in period $t - 1$ via a dynamic reaction function. As above, variations between subjects are captured through the variance of the residual term. Our model may be written as a two-equation system:

$$q_{Lt} = \phi_L + \mu_L q_{Ht-1} + \eta_{Lt} \quad (7)$$

$$q_{Ht} = \phi_H + \mu_H q_{Lt-1} + \eta_{Ht} \quad (8)$$

where $\eta_{mt} = \rho_m \eta_{m,t-1} + \epsilon_{mt}$, for $m = L (H)$ for the low (high)-cost agent. As in subsection A, we want to allow for learning or signaling effects, and so we assume a first-order autoregressive process here. By combining a first-order serial correlated disturbance with the dynamic reaction function, the model in equations (7) and (8) uses information from each of the two preceding periods to predict current choices, which parallels the regression model in equation (5).

Consistent and asymptotically efficient estimates of the ϕ 's and μ 's in equations (7) and (8) can be obtained by adapting the two-step processes commonly employed for treating heteroskedasticity or serial correlation. We first use a Cochrane–Orcutt procedure to control for the assumed serial correlation. Using the resultant residuals, we then estimate the variances for each agent and perform weighted least squares. This process is conducted for each of four groups: high-cost subjects with private information, high-cost subjects with common knowledge, low-cost subjects with private information, and low-cost subjects with common knowledge.

Consider now the deterministic analogues to equations (7) and (8), and suppose that agents choose the values q_L^e and q_H^e for several consecutive periods. This gives a system of

TABLE 4.—REGRESSION ANALYSIS OF INDIVIDUAL OUTPUTS, ASYMMETRIC TREATMENTS

Parameter	Estimate	Standard Error
<i>1. Common Knowledge</i>		
Low-cost equation		
a_L	32.6728	0.8191
b_L	0.0469	0.0320
$R^2 = 0.9281$		
Durbin–Watson statistic = 2.1531		
High-cost equation		
a_H	23.3675	0.7971
b_H	0.0319	0.0229
$R^2 = 0.9530$		
Durbin–Watson statistic = 1.9630		
<i>2. Private Information</i>		
Low-cost equation		
a_L	28.6090	1.0815
b_L	0.1344	0.0364
$R^2 = 0.9694$		
Durbin–Watson statistic = 1.8275		
High-cost equation		
a_H	25.8064	1.1182
b_H	0.0744	0.0341
$R^2 = 0.7026$		
Durbin–Watson statistic = 1.9333		
<i>Implied Steady States</i>		
α_{L1}	33.8192	0.1861
α_{H1}	24.4475	0.1412
α_{L2}	32.4018	0.2731
α_{H2}	28.2174	0.2507
Test statistic on $H_0: \alpha_{L1} = \alpha_{L2} = -4.2893$		
Test statistic on $H_0: \alpha_{H1} = \alpha_{H2} = 13.1040$		
5% critical value = 1.96		

two equations in two unknowns. Solving this system provides maximum-likelihood estimates of the underlying steady-state values (Fomby et al. (1988)). These estimates and the ϕ 's and μ 's are related by

$$q_L^e = (1 - \mu^L \phi_H) / (1 - \mu_L \mu_H) \quad (9)$$

$$q_H^e = (1 - \mu^H \phi_L) / (1 - \mu_L \mu_H). \quad (10)$$

Covariance information from the maximum-likelihood estimates of the μ 's and ϕ 's can similarly be used to construct consistent estimates of the covariance structure for the steady-state values (Fomby et al. (1988, corollary 4.2.2)). The corresponding estimates of q_L^e and q_H^e , which we denote by α_{L1} and α_{H1} for the common-knowledge treatment and α_{L2} and α_{H2} for the private-information treatment, are given in table 4. Both hypotheses 8 and 9 are tested using an asymptotic t -test on the difference of the parameter estimates, $\alpha_{m1} - \alpha_{m2}$, $m = L, H$. The two test statistics are statistically significant, and so we reject hypotheses 8 and 9 in favor of the alternative hypotheses associated with Shapiro's (1986) proposition that improved information raises (lowers) low (high)-cost agents' equilibrium outputs.

Because a relatively greater share of the output is produced by the low-cost firm when there is full information, the average cost of all units sold in the market is lower.

On average, unit costs in our laboratory markets fall by about 10%, from 0.2327 with private information to 0.2098 when payoffs are common knowledge. Since convergence to these relatively efficient market shares is more rapid when payoffs are common knowledge, and since *total* market outputs are not significantly different under the two information treatments, the results in this section provide evidence on the potential welfare benefits of common knowledge in asymmetric markets.

V. Conclusion

In this paper, we have analyzed the collusive effects of common knowledge on duopoly behavior in laboratory markets. Our main findings are that the provision of information can facilitate cooperation and act as a restraint of trade if firms are sufficiently similar, though such facilitation is unlikely when payoffs are even moderately asymmetric. When there is cost asymmetry, making payoffs common knowledge generates important net benefits. Payoff information can change firms' outputs so that the relatively more efficient producers have a larger market share, without changing total industry output. Further, making payoffs common knowledge brings a market into equilibrium faster. These benefits stand in contrast to the potential danger of collusion, and generate welfare gains that until now have not been empirically recognized. Our findings add perspective to the policy issue of allowing firms access to information about their rivals' profits. The results bolster a rule of reason approach toward allowing increased information between market rivals. It may only be necessary to challenge the sharing of information among rivals when the industry is composed of a relatively small number of firms, similar in size, with very similar products and technologies.¹⁶

One explanation for agents' lack of success at colluding when payoffs are private information is that it is harder for agents to signal their willingness to collude (Kreps and Spence (1985)). A willingness to collude can be signaled by unilaterally producing less or raising prices. Such signaling can be costly because a unilateral reduction in output can sharply reduce static payoffs. When payoffs are not common knowledge, an agent considering reducing their output so as to signal a willingness to collude must realize that a rival does not know his or her intentions. Indeed, the agent's rival may interpret the action as a tremble. An absence of information muddles the signal and in general lowers the expected benefit from signaling. The presence of payoff asymmetries probably distorts the signaling process, too. When agents have identical payoff functions, identical reductions in quantity are a transparent and focal means of

cooperation in the absence of communication. But when there is cost asymmetry, equal reductions in quantity impose a greater sacrifice on the small seller. The rules of cooperation are less transparent and there are no obvious ways by which agents can become more cooperative.

REFERENCES

- Abreu, Dilip, "Extremal Equilibria of Oligopolistic Supergames," *Journal of Economic Theory* 39 (June, 1986), 191–225.
- Alger, Daniel, "Laboratory Tests of Equilibrium Predictions with Disequilibrium Data," *Review of Economic Studies* 54 (Jan., 1987), 105–146.
- Cason, Timothy, "The Impact of Information Sharing Opportunities on Market Outcomes: An Experimental Study," *Southern Economic Journal* 61 (July, 1995), 18–39.
- Cason, Timothy, and Charles F. Mason, "Uncertainty, Information Sharing and Collusion in Laboratory Duopoly Markets," University of Southern California, Working paper (1995).
- Clarke, Richard, "Collusion and the Incentives for Risk Sharing," *Bell Journal of Economics* 14 (Autumn, 1983), 383–394.
- Colman, Andrew, *Game Theory and Experimental Games* (New York: Pergamon, 1982).
- Cooper, Russell, Douglas DeJong, Robert Forsythe, and Thomas Ross, "Communication in Coordination Games," *Quarterly Journal of Economics* 107 (May, 1992), 739–771.
- Davis, Douglas, and Charles Holt, *Experimental Economics* (Princeton, NJ: Princeton University Press, 1993).
- Dolbear, F. T., L. B. Lave, G. Bowman, A. Lieberman, E. Prescott, F. Reuter, and R. Sherman, "Collusion in Oligopoly: An Experiment on the Effect of Numbers and Information," *Quarterly Journal of Economics* 82 (May, 1968), 240–259.
- Fomby, Thomas, R. Carter Hill, and Stanley Johnson, *Advanced Econometric Methods* (New York: Springer, 1988).
- Fouraker, Lawrence, and Sidney Siegel, *Bargaining Behavior* (New York: McGraw-Hill, 1963).
- Friedman, Daniel, and Shyam Sunder, *Experimental Methods* (Cambridge: Cambridge University Press, 1994).
- Friedman, James, *Oligopoly Theory* (Cambridge: Cambridge University Press, 1983).
- Fudenberg, Drew, and Jean Tirole, *Game Theory*, (Cambridge, MA: MIT Press, 1991).
- Gal-Or, Esther, "Information Sharing in Oligopoly," *Econometrica* 53 (Mar., 1985), 329–342.
- Green, Edward, and Robert Porter, "Noncooperative Collusion under Imperfect Price Information," *Econometrica* 52 (Jan., 1984), 87–100.
- Grether, David, and Charles Plott, "The Effect of Market Practices in Oligopolistic Markets: An Experimental Examination of the Ethyl Case," *Economic Inquiry* 22 (Oct., 1984), 479–507.
- Kalai, Ehud, and Ehud Lehrer, "Rational Learning Leads to Nash Equilibrium (A New Extended Version)," Northwestern University, Discussion Paper 925 (1991).
- Kirby, Alison, "Trade Associations as Information Exchange Mechanisms," *Rand Journal of Economics* 19 (Spring, 1988), 138–146.
- Kreps, David, *Game Theory and Economic Modeling* (New York: Oxford University Press, 1990).
- Kreps, David, and Michael Spence, "Modelling the Role of History in Industrial Organization and Competition," in G. Feiwel (ed.), *Issues in Contemporary Microeconomics and Welfare* (Albany, NY: State University of New York Press, 1985).
- Li, Lode, "Cournot Oligopoly with Information Sharing," *Rand Journal of Economics* 16 (Winter, 1985), 521–536.
- Mason, Charles, and Owen Phillips, "Dynamic Learning in Two-Person Experimental Games," University of Wyoming, Working Paper (1994a).
- , "Information and Cost Asymmetry in Experimental Duopoly Markets," University of Wyoming, Working Paper (1994b).
- Mason, Charles, Owen Phillips, and Cliff Nowell, "Duopoly Behavior in Asymmetric Markets: An Experimental Evaluation," this REVIEW 74 (Nov., 1992), 622–670.

¹⁶ The most recent policy statements made by the Department of Justice (DOJ) and the Federal Trade Commission (FTC), the 1992 *DOJ & FTC Horizontal Merger Guidelines*, do provide a basis for the implementation of this rule of reason. Section 2.11 states that "reaching terms of coordination may be facilitated by product or firm homogeneity," and that "[k]ey information about rival firms and the market may also facilitate reaching terms of coordination" (emphasis added).

- Mason, Charles, Owen Phillips, and Douglas Redington, "On the Roles of Gender in a Noncooperative Game," *Journal of Economic Behavior and Organization* 15 (Mar., 1991), 215–235.
- Novshek, William, and Hugo Sonnenschein, "Fulfilled Expectations Cournot Duopoly with Information Acquisition and Release," *Bell Journal of Economics* 13 (Spring, 1982), 214–218.
- Rapaport, Anatol, Melvin Guyer, and David Gordon, *The 2 × 2 Game* (Ann Arbor: University of Michigan Press, 1976).
- Rasmusen, Eric, *Games and Information: An Introduction to Game Theory* (Cambridge, U.K.: Basil Blackwell, 1989).
- Schelling, Thomas, *The Strategy of Conflict* (Cambridge, MA: Harvard University Press, 1960).
- Schmalensee, Richard, "Competitive Advantage and Collusive Optima," *International Journal of Industrial Organization* 5 (1987), 351–367.
- Shapiro, Carl, "Exchange of Cost Information in Oligopoly," *Review of Economic Studies* 53 (July, 1986), 433–446.
- Shapiro, Leonard, "Decentralized Dynamics in Duopoly with Pareto Optimal Outcomes," *Bell Journal of Economics* 11 (Autumn, 1980), 730–744.
- Smith, Vernon, Kevin McCabe, and Stephen Rassenti, "Lakatos and Experimental Economics," in Neil de Marchi and Mark Blaug (eds.), *Appraising Economic Theories: Studies in the Methodology of Research Programs* (New York: Edward Elgar, 1990).
- Stigler, George, "The Economics of Information," *Journal of Political Economy* 69 (June, 1961), 213–225.
- , "A Theory of Oligopoly," *Journal of Political Economy* 72 (Feb., 1964), 213–225.
- Vives, Xavier, "Duopoly Information Equilibrium: Cournot and Bertrand," *Journal of Economic Theory* 34 (Oct., 1984), 71–94.