Market power and mergers in laboratory markets with posted prices

Douglas D. Davis and Charles A. Holt*

Abstract

This paper uses laboratory methods to evaluate determinants of supra-competitive pricing. The experiment involves three treatments, each with the same market supply, demand, and competitive price. In the baseline treatment, capacity is divided among five sellers so that the competitive price is a Nash equilibrium. Market power is created in a second treatment by reallocating capacity among the sellers. This market power raises observed prices in all sessions. In a third treatment, the three smallest sellers are merged in a way that holds market power constant. The consolidation has little residual effect on prices.

1. Introduction

One of the longstanding embarrassments for economic theory is the lack of consensus about the likely outcomes in markets that are neither highly competitive nor highly monopolized. Even the substantial majority of industrial organization economists who rely on noncooperative game theory will disagree about modeling issues, about refinements that rule out "unreasonable" Nash equilibria, and about the extent to which collusive prices can be maintained by mutual trust or fear of retaliation. Due to the absence of a single unified theory of oligopoly, the theoretical assumptions underlying U.S. antitrust policy have varied over time and particularly across political regimes.

As evidenced by an extensive and often vociferously contested empirical literature, it is difficult to distinguish rival theories. Accurate cost, production and transactions price information

---

* Virginia Commonwealth University and University of Virginia respectively. This research was supported by grants from the University of Virginia Bankard fund, the Virginia Commonwealth University Faculty Excellence Fund, and the National Science Foundation (grants SES 9012694 and SES 9012691). We are indebted to student and faculty participants in the University of Virginia Microeconomics Workshop for useful suggestions.

is often scarce. Even when reliable data are available, too many theoretically relevant factors change at once to allow for a clean "natural experiment." A merger that subsequently generated supra-competitive prices, for example, might have simultaneously decreased the number of sellers, increased market concentration, decreased industry capacity and reduced product heterogeneity. Any one or combination of these effects might have caused the noncompetitive performance.

The careful application of laboratory methods can allow some insight into the appropriate theoretical basis for antitrust policy, by cleanly isolating the effects of alternative determinants of price increases. This paper investigates two of these determinants. Specifically, we examine behavior in an experiment that isolates the effects of market power, which is usually indicated by the capacity of a single seller to profit from a unilateral price increase above a common competitive price determined by the intersection of aggregate supply and demand. We also examine a pure numbers effect arising from a reduction in the number of competitors. All treatments involve an identical market supply, market demand, and competitive price. Market power is altered in a first pair of treatments by reallocating capacity among a group of five sellers in a way that causes the Nash equilibrium prices to exceed the competitive level. The pure-numbers effect is examined by reducing the number of sellers from five to three in a manner that increases measured concentration without altering market power in a theoretical sense.

Despite the wide chasm between the laboratory and the complex markets of antitrust concern, the experiment is designed to address issues of economic policy. Although we cannot hope to make claims from the laboratory regarding the competitiveness of particular markets in the economy, antitrust policy can be (and has been) based on a variety of different economic models. Experimental methods can be used to distinguish among competing models in starkly simple environments. If the predictions of a theory fail to organize and explain data in a very simple environment, the burden of proof for a particular explanation of behavior should shift to the proponent of the theory.[1]

---

[1] For example, Holt (1985) showed that the notion of a Consistent Conjectures Equilibrium has very little predictive power in a simple quantity-choice oligopoly experiment. Similarly, Grether and Plott (1984) report the results of an experiment designed to rebut an assertion by counsel for the plaintiffs that a variety of "facilitating factors" (such as meet-or-release and most-favored-nation clauses in sales contracts) could not increase market prices.

Our focus on isolating market-power effects also addresses an important issue in the experimental literature. Numerous experiments have been reported where prices diverge from the competitive prediction when market power is present. (See, for example, Dolbear et al. (1968), Ketcham, Smith and Williams (1984), Holt, Langan and Villamil (1986), Alger (1987), Isaac and Reynolds (1989), Davis, Holt and Villamil (1990), Davis and Williams (1991), Brown-Kruse, Rassenti, Reynolds and Smith (1990), and Welford (1990)). But it is not possible to conclude from these studies that the observed deviations were motivated by market power or by some other design characteristic coincidental to the introduction of market power, such as few sellers or unusually-shaped aggregate supply and demand arrays.[2] It is not possible to isolate the separate effects of market power unless one holds such design characteristics constant when market power is introduced.

The remainder of this paper is organized as follows. Section 2 describes the experimental design. Theoretical Nash equilibrium price distributions for the treatments with market power are described in section 3. The experimental procedures and data are reported in sections 4 and 5 respectively. The final section contains a conclusion.

2. Experiment design

We will say that market power exists if the noncooperative equilibrium for the market game yields prices that exceed the competitive level.[3] In most market structures, this type of power is indicated if at least one seller can profit from a unilateral price increase above a common competitive price level (which implies that the competitive equilibrium is not a noncooperative equilibrium). This conception of market power is consistent with the notion of market power implicit in the 1984 U.S. Department of Justice Horizontal Merger Guidelines, which assesses market power in terms of the profitability of a small unilateral price increase.

[2] In fact, observed supracompetitive prices are often attributed to factors of this type. For example, Davis and Williams (1991) speculate that the supracompetitive prices observed in their market power sessions may be a consequence of an excess supply of only a single unit at prices close to the competitive level. Similarly, supracompetitive pricing may have been promoted by the fact that sellers earned zero profits in the competitive equilibrium in the design used by Brown-Kruse, Rassenti, Reynolds and Smith (1990), and in several of the designs used by Alger (1987).

[3] Holt (1989) discusses this and other definitions of market power.

As will be shown below, changing the allocation of industry capacity can have the theoretical effect of creating market power, since larger sellers with more inframarginal units may have a greater incentive to trade off sales of marginal units to obtain supra-competitive prices. Therefore, acquisitions that transfer capacity among firms can create market power, even though they have no effect on supply, demand, the number of sellers, or the competitive price.

If sellers recognize and respond strategically to unilateral incentives, then the addition of market power will (by definition) yield supra-competitive prices, even in the absence of cooperative behavior. Still larger price increases may be observed if sellers behave cooperatively, in the sense that they resist unilateral incentives to cut price. When prices exceed the levels determined by a unique noncooperative equilibrium, we will say that this is due to collusion, overt collusion if sellers are able to communicate directly, and tacit collusion otherwise. Since overt collusion is illegal, an evaluation of a proposed merger often revolves around assessing the likelihood of tacit collusion. A clear behavioral assessment of market power should distinguish between price increases due to noncooperative behavior and price increases due to tacit collusion.

We also consider the effects of changes in the number of sellers that change measured concentration, because structural measures are often used to evaluate the likelihood of supra-competitive pricing. One issue is whether the concept of market power we offer is more useful as a means of predicting noncompetitive behavior than standard alternative indicators.

Our designs are illustrated by the aggregate supply and demand arrays shown in the three panels of figure 1. The left and center panels involve five subjects, who are identified as sellers S1, S2, S3, S4, and S5. The panel on the right involves three subjects, identified as S1, S2 and S_m. The identities of sellers holding various units are listed below the corresponding units on each supply curve. The three designs in figure 1 share many common features. In each case, sellers as a group have the capacity to offer up to 11 units, arrayed in a two-step configuration of 5 low-cost units and 6 high-cost units. Similarly, each design uses the same two-step demand function, composed of 8 high-value units and 3 low-value units. The high reservation value is denoted by r, the low reservation value is denoted by P_c, the high cost is denoted by c, and the low cost is normalized to 0. The competitive equilibrium in each design involves a quantity of 11 and any price in the range [c, P_c] of vertical overlap of the demand and supply functions.

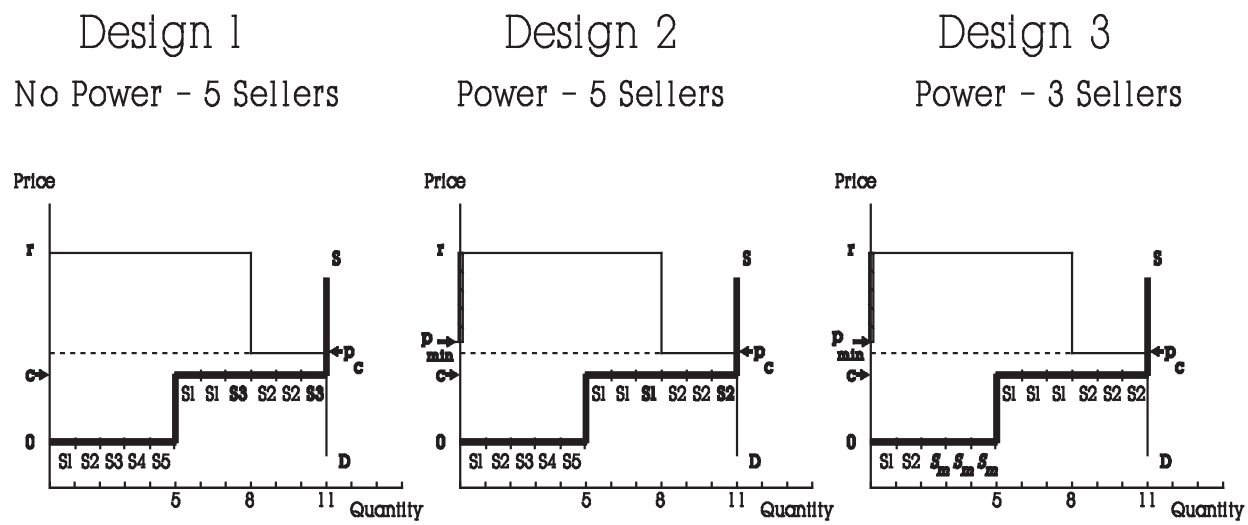

Figure 1. Supply and Demand Arrays

The focus of our analysis is on oligopoly situations in which a few sellers face passive buyers with no countervailing power. The market is assumed to be governed by posted-offer trading rules, which are a multi-stage version of a Bertrand pricing game. Each stage-game consists of two steps. First, the sellers simultaneously make price decisions, which are publicly displayed to all buyers and sellers. Second, the buyers shop, making purchase decisions on the basis of posted prices. To ensure that buyers fully reveal demand (i.e. to ensure that no buyer forgoes profitable purchases in an attempt to influence pricing decisions in subsequent periods), the demand side of the market in each design is represented by a simulated price-taking buyer.[4] The simulated buyer makes all profitable purchases, buying first from the seller(s) with the lowest posted price, then from the seller(s) with the next lowest price, and so forth. In the event of identical prices, the buyer equalizes purchases among the tied sellers as much as possible. When equal purchases are not possible, residual purchases are made randomly.

---

4 Although human buyers in posted-offer experiments typically reveal demand fully (Ketcham, Smith and Williams, 1984), complete public information about costs and values may encourage strategic withholding of purchases, especially with few buyers (Davis and Williams, 1989). As explained in the next section, complete information was provided, and therefore, we decided to simulate the demand side of the market.

The designs in figure 1 differ in the manner in which the units with bold-faced seller identities are assigned. In design 1, shown on the left side of figure 1, capacities are limited in the sense that no seller has a capacity that exceeds 3 units. This feature, combined with the excess supply of 3 units at supra-competitive prices, makes the highest competitive price, P_c, a pure-strategy Nash equilibrium. If all sellers post this price and offer all units for sale, no seller may unilaterally increase profits, since a unilateral price increase would result in no sales and a unilateral decrease would not increase sales.

Design 2, in the center of figure 1, involves a reassignment of units of capacity from seller S3 to sellers S1 and S2, a relatively minor change that causes the predictions of Nash and competitive theories to diverge. Sellers S1 and S2 may each offer up to 4 units in design 2, while S3, S4, and S5 may offer only 1 unit each. Unlike design 1, sellers S1 and S2 are each certain to sell at least one of their 4 units in design 2. Starting at a common price of P_c, a unilateral increase to r will be profitable for these large sellers if the increased revenue on the sale of the first unit, $r - P_c$, exceeds the lost profit of $3(P_c - c)$ on the 3 marginal units. This condition reduces to a requirement that the top of the competitive price range, P_c, be below a price, p_{min}, defined in (1):

$$p_{\min} = \frac{3c + r}{4}. \tag{1}$$

The selection of parameter values for which $P_c < p_{min}$ will ensure that a common competitive price of P_c is not a Nash equilibrium in design 2. It follows from the way p_{min} was calculated that it is the lower bound of an "Edgeworth cycle" for sellers S1 and S2. The range of the Edgeworth cycle is illustrated by the darkened section of the vertical axis in the graph for design 2. Finally, note that sellers S3, S4 and S5 face rather different pricing incentives than S1 or S2. The small sellers are certain to sell their single units only by posting a price slightly below p_{min}. Thus, a small seller will tend to cling close to p_{min}, and will price above this limit only if (s)he considers it sufficiently likely that at least one of the large sellers posts a yet higher price.

Design 3, shown on the right side of figure 1, differs from design 2 in that the 3 low-cost units given to sellers S3, S4 and S5 are consolidated into a single seller, S_m. This alteration is emphasized by the bolded seller identifier numbers shown under the supply function in design

3. Importantly, large sellers S1 and S2 retain the characteristic that they are each certain to sell at least one unit in design 3, so that P_c is not a Nash equilibrium in design 3 whenever P_c is not a Nash equilibrium in design 2. Moreover, S_m faces incentives imposed on each of the small sellers in design 2, since S_m is still not assured of selling any units at any price over p_{min}, and must consequently price below at least one of the large sellers to realize any sales. As indicated below, the consolidation of units from S3, S4 and S5 into a single firm does not affect the noncooperative equilibrium price outcomes in an essential way.

To summarize, aggregate supply, aggregate demand and the number of sellers remain fixed across designs 1 and 2, but the reallocation of units among sellers from design 1 to design 2 creates market power. Comparison of behavior across these designs will allow assessment of a "pure-power" effect. Similarly, comparison of behavior across designs 2 and 3 will allow assessment of a "pure-numbers" effect. The consolidation of low cost units from S3, S4 and S5 reduces the number of sellers, but aggregate supply, aggregate demand and the noncooperative equilibrium remain fixed across designs 2 and 3.

Two other features of our designs bear brief comment. The first is that concentration, of course, increases from design 1 to design 2, and then again from design 2 to design 3. To provide a frame of reference for concentration in our markets with those considered typically to be of policy concern in antitrust analysis, it is instructive to measure concentration in terms of the Herfindal-Hershmann Index (HHI), calculated as the sum of the squared market shares for each firm, multiplied by 10,000. The HHI (based on capacity) for design 1 is 2397. The reallocation of units among sellers in design 2 increases concentration by 496 points to 2893. The consolidation of capacity among S3, S4 and S5 in design 3 further increases concentration by 495, to 3388. Under the 1984 DOJ Horizontal Merger Guidelines, consolidations which change HHI values by at least 50 points, and which generate a post-consolidation HHI value in excess of 1800, are considered to be highly concentrated and are likely to be challenged. Thus, either consolidation examined here (from design 1 to design 2, or from design 2 to design 3) would quite likely be challenged under the 1984 Guidelines.

Second, notice that the three units that fail to trade when market power is exercised are low-value units for buyers, and are likely to be high-cost units for sellers. In the particular parameterization we employ, for example, less than 5% of the aggregate surplus is lost if sellers

1 and/or 2 fail to sell 3 high-cost units. Thus, in this design, market power primarily involves redistributional rather than efficiency consequences. As a general matter, market power exercise may affect efficiency more prominently. We elected to minimize these efficiency consequences in our designs, however, in order to get a clear separation in price predictions.

3. Calculation of equilibria in the five-seller and three-seller power designs[5]

A distinct advantage of our market power designs is that, despite the asymmetry of sellers' capacities, the noncooperative equilibria can be calculated and used as a basis for evaluating performance and measuring tacit collusion. As suggested by the persistent incentives for large sellers to shade on any common price in the above discussion, there is no pure-strategy Nash equilibrium in the power designs.[6] The subgame-perfect Nash equilibrium involves mixing. Derivation of the relevant mixing distributions is somewhat tedious. Nevertheless, the method of calculation may be succinctly outlined. The process involves three steps: 1) informed guessing about the range over which sellers randomize, 2) calculating sellers' expected profits at some convenient point in the range (typically the upper or lower bound), and 3) deriving price distributions that make each seller's expected profit constant over the support of the distribution, since otherwise the seller would not be willing to choose prices randomly.

The first two steps are straightforward for the five-seller power design shown in the middle panel of figure 1. For the large sellers, a reasonable guess is that the equilibrium involves randomization over the range of the Edgeworth cycle [p_{min}, r], and that security earnings are uniquely determined by the earnings on the zero-cost unit that is certain to sell at the limit price, r.[7] The small sellers are not certain to make any sales at the limit price r, so their security

---

[5] A detailed derivation of the noncooperative equilibria discussed in this section appears in a longer, unpublished version of this paper, which is available from the authors on request.

[6] The nonexistence proof uses standard arguments (see e.g., Holt and Solis-Soberon, 1992).

[7] Obviously, there is no incentive to ever price above r, and the large sellers can always sell at least one unit at prices below r, so security earnings can be no less than r. The only way that their expected earnings could exceed this minimum would be for both sellers to have a mass of probability at the upper bound, r, but in this case each seller would prefer a price slightly below r, which is inconsistent with the constant-expected-profit condition for randomization. The upper bound of the large sellers' randomization range cannot be below r, because the large sellers can always sell at least one unit at r with certainty. The uniqueness of the upper bound of the large sellers' distributions and the absence of a

earnings must be determined by the bottom of the large sellers' mixing distribution, p_{min}, which represents the highest price where the small sellers are certain to make a sale.[8] As will be seen presently, the small sellers will randomize over only a lower segment of the [p_{min}, r] range.

The third step, finding the pricing distributions for the mixed-strategy equilibrium, is more involved. These distributions are calculated by jointly considering for each seller the price distributions that the <u>other sellers</u> must follow in order for that seller's expected profits to equal a constant. Consider, for example, conditions under which S1 will randomize. Define G(p) as the distribution function for the price of S2, so G(p) is the probability that a price p posted by S1 is higher than the price posted by S2. Similarly, let F(p) denote the price distribution function for each of the three small sellers. For S1 to randomize, other sellers must post prices so that expected earnings are the same at any arbitrary price p, as they are at security price, r, or

$$r = p[G(p)F(p)^3] + (2p-c)[3G(p)F(p)^2(1-F(p))] + (3p-2c)[3G(p)F(p)(1-F(p))^2] \quad (2)$$
$$+ (4p-3c)[G(p)(1-F(p))^3] + (4p-3c)[1-G(p)].$$

Each term in the expected profit expression on the right side of (2) is simply the probability-weighted earnings possibilities for S1 at any price, p. For example, the first probability term, $[G(p)F(p)^3]$, is the probability that the other large seller and all three small sellers price below p; this is a product of probabilities since firms randomize independently. Only one unit is sold when a large seller has the highest price, so this probability term is multiplied by the earnings, p, from the sale of the zero-cost unit. Other terms are interpreted analogously. Due to the symmetry of incentives, (2) describes the equilibrium randomizing distribution for both large sellers.

---

spike of probability at this upper bound means that the expected payoff for the large sellers is uniquely determined.

[8] The small sellers would never wish to price below p_{min}, since they can sell all units whenever they price below the large sellers. For reasons parallel to those used in the previous note to show that no probability mass point exists at r, no probability mass point exists at p_{min}, and each small seller can be assured of selling its (zero-cost) unit at this price. Thus, the lower bound of the small sellers' mixed-strategy distribution must exactly equal p_{min}. The uniqueness of the security earnings at the lower bound implies that the expected payoff for the small sellers is uniquely determined.

The equation identifying the mixing distribution for the small sellers is more straightforward, since their sales are determined exclusively by the pricing choices of the large sellers. The capacity of each large seller exceeds excess supply at supra-competitive prices. Thus, the small seller's single unit will sell unless the seller is underpriced by both S1 and S2. The probability that a small seller with a price of p will sell its unit is 1 minus the probability that both large sellers price below p, which is $[1 - G(p)^2]$. The small seller only has one unit at zero cost, so its expected earnings expression is $p[1 - G(p)^2]$. Equating expected earnings to security profits, p_{min} given in (1), generates the randomizing condition for the small sellers:

$$p[1 - G(p)^2] = \frac{3c + r}{4}. \tag{3}$$

The equilibrium mixing distributions G(p) and F(p) are calculated as the simultaneous solution to equations (2) and (3):

Sellers S1 and S2:

$$G(p) = \sqrt{\frac{4p - 3c - r}{4p}} \quad \text{for} \quad p_{min} \leq p \leq p* \tag{4}$$

$$G(p) = \frac{4p - 3c - r}{3(p - c)} \quad \text{for} \quad p* < p \leq r.$$

Sellers S3 - S5:

$$F(p) = \frac{2}{3(p-c)}\sqrt{p(4p - 3c - r)} \qquad for \quad p_{\min} \leq p \leq p* \tag{5}$$

$$F(p) = 1 \qquad for \quad p* < p \leq r.$$

Notice that each distribution has two parts, distinguished by the relation of price to p*, which denotes the upper end of the mixing range for the small sellers.[9]

Two further features of the noncooperative equilibrium bear brief comment. First, the consolidation of small sellers S3, S4 and S5 into a single merged seller, S_m, does not essentially alter the equilibrium identified above in equations (4) and (5). Intuitively, this easily follows from the observation that S1 and S2 can still potentially supply all 8 units demanded at supra-competitive prices, despite a consolidation. Thus, the merged small seller faces incentives to price below the large sellers that are identical to those faced by the diffuse small sellers.[10] Similarly, randomizing conditions for the large sellers are unchanged.[11]

Second, finite repetition of the stage game does not create additional subgame perfect equilibria. This result easily follows since payoffs in the stage game equilibria are unique, and thus there are no alternative equilibria that could serve as threats to support a conspiracy.[12] For this reason, no trigger-strategy equilibria exist in any finitely repeated version of the games

---

[9] To verify these distributions, note that when F(p)=1 (as in the lower part of (5)), equation (2) is equivalent to the lower part of (4). When p is below p* and F(p)<1, the equivalence of (3) and the upper part of (4) is immediate, but showing the consistency of (2) with the upper parts of (4) and (5), involves rather tedious algebra. This latter calculation is simplified by adding and subtracting 4p−3c from the right side of (2), and exploiting the fact that the sum of the bracketed probability terms in (2) must equal 1. The simplified expression is readily converted to a product of the F(p) and G(p) expressions in the upper parts of (4) and (5).

[10] Specifically, since S_m has three units, this seller's security earnings are $3p_{\min}$, and expected earnings at any price below p* are $3p[1 - G(p)^2]$. Equating these expressions yields equation (3).

[11] Let $F_m(p)$ denote the price distribution for S_m. A large seller will sell 4 units at a price p, unless p is the highest price charged, which occurs with probability $G(p)F_m(p)$. Thus, the probability-weighted earnings expression for the large sellers is $(4p-3c)[1-G(p)F_m(p)] + pG(p)F_m(p)$, which must equal security earnings, r, in a mixing equilibrium. When $F_m(p) = 1$, this equation yields the lower part of (4). When p < p*, it is straightforward but tedious to show that this equation and (3) yield the upper parts of (4) and (5).

[12] Payoffs must be unique. As noted at the outset of this section, there can be no Nash equilibrium in pure strategies. Moreover, earnings in any mixing equilibrium must equal those for the mixed equilibrium characterized by (4) and (5), since they must be derived from the unique "security profits" at the relevant boundary of the support of the price distributions.

illustrated in designs 1, 2 and 3. The only subgame perfect equilibrium in each design is a repetition of the stage-game equilibrium.[13]

To summarize, a comparison of designs 1 and 2 in figure 1 reveals that the power treatment involves a relatively minor reallocation of capacity that does not affect aggregate demand and supply, but this reallocation changes the noncooperative equilibrium price from P_c in design 1 to randomization at prices above p_{min} in design 2. The consolidation of sellers S3, S4 and S5 into a single seller S_m does not alter the equilibrium price distributions calculated for design 2. Finally, the equilibria in designs 1, 2 and 3 are all unique subgame perfect equilibria for any finitely repeated version of the game.

4. Experiment procedures

We conducted twelve two-hour sessions, each with a different cohort of 5 subjects. Six sessions involved designs 1 and 2 (5 seller power/ 5 seller no-power), and six sessions involved designs 2 and 3 (5 seller power/ 3 seller power). Each session lasted for 60 periods, with a change of treatment at the midpoint. Sequence effects were controlled by reversing the order of treatments in every other session. The experiment was conducted at Virginia Commonwealth University using a networked-PC implementation of a posted-offer institution, written by Davis. Fully revealing demand behavior was simulated by the experiment monitor. Participants were undergraduate business students with previous experience as sellers in posted-offer markets (but with different supply and demand structures, and in different cohorts). Participants were paid a $3.00 appearance fee, plus money earned by selling units at prices above their costs. In sessions using the design 3 treatment, the two idle sellers were paid $6.00 for monitoring the 30 periods in which they were inactive.[14] Earnings ranged from $15 to $55 per subject.

With few exceptions, the procedures were standard for computerized posted-offer environments in which sellers are not able to communicate with each other (see Ketcham, Smith,

---

[13] Of course, multiple Nash equilibria exist for the finitely repeated game, but threats may not be equilibrium threats, and thus these equilibria are not subgame perfect. As is well known, subgame perfect trigger-strategy equilibria do exist for infinitely repeated versions of the game.

[14] The software allows sellers to monitor the market passively by posting a price and quantity of 0 in a period.

and Williams, 1984). There were three significant departures from the standard procedures. First, while participants proceeded though interactive instructions at their own terminals, the instructions were read aloud to them by the experimenter.[15] Second, since the software constrained prices to be integer valued, we increased all model parameters by a factor of 10 and divided laboratory earnings by 10 to calculate subjects' cash earnings. This 10-to-1 conversion rate generated a fine price grid that reduced the possibility of price ties. Third, participants were given complete information about the stopping rule for each treatment, about other sellers' costs for each treatment, and about the preferences and the shopping behavior of the simulated buyer. The relevant parts of the subjects' instructions are reproduced in the appendix. This information about the structure of the experiment was provided to match more closely the complete information (or clearly specified incomplete information) conditions typically assumed in noncooperative game theory. Our simple supply and demand structures are attractive in part because of the ease with which cost and value information can be presented.

Parameter choices were made to avoid focal numbers and to ensure reasonable earnings for small sellers in competitive outcomes. If r were 434 and c were 154, then p_{min} would be 224 and p^* would be 242.[16] Since the low-cost step has been normalized to zero, a parameter-shifting constant of 105 was added to all of the above parameters to raise the low cost step above 100; this ensures that all possible price choices involve the entry of three digits on the subject's keyboard. With this shift, $r = 539$, $c = 259$, $p_{min} = 329$, and $p^* = 347$. Recall that the choice of the competitive-price demand step, P_c, is arbitrary, subject to the constraint implied by (1) that it be between the high cost, c, and the lower bound of the mixed distribution, p_{min}. We selected $P_c = 309$, a value in the upper end of the acceptable range, to enhance competitive earnings. The medians of the mixing distributions in (4) and (5) for the large and small sellers are 371 and 332, respectively, after adding the shift parameter.

Given these parameters, a maximum of 3160 in surplus may be extracted from the market. In the no-power treatment, sellers S1, S2 and S3 each earn 304 from the sale of 3 units at the

---

[15] Reading instructions aloud facilitates learning and increases common knowledge.

[16] The value of c was chosen to be equal to 11r/31, which is convenient since it yields an integer value for the p^* that satisfies $F(p^*) = 1$.

competitive price (= [309 − 105] + 2[309 − 259]), while S4 and S5 each earn 204 on the sale of a single low-cost unit (= 309 − 105). In the mixing equilibrium for the power treatment, sellers S1 and S2 each earn, on average, security profits of 434 (= 539 − 105). Similarly, expected earnings for S3, S4 and S5, on average, equal 224 (= 329 − 105). Summing across sellers, it follows that under the competitive prediction, sellers earn 42% of the maximum available surplus, (= 1320/3160), while this ratio increases to 49% in the mixing equilibrium (= 1540/3160). Finally, the three units that fail to trade under the Nash equilibrium in the market power treatment generate a surplus loss of 150 (= 3[309 − 259]). Thus, given these parameters slightly more than 95% of maximum possible gains from exchange (= 3010/3160) are extracted in the mixed-strategy equilibrium with market power.

5. Results

Our primary focus is on price effects, so it is useful to begin by looking carefully at the dynamics of individual price decisions in a particular session, since several interesting patterns are observed in all twelve sessions. In figure 2, the first 30 periods, indicated on the horizontal axis in the center, involved the five-seller power treatment. The final 30 periods, to the right of the center vertical line, involved the no-power treatment. The supply and demand functions are reproduced on the left side to indicate the relative positions of the key price predictions; the competitive price is 309 (see the horizontal line), and the range of randomization for the power treatment is from 329 (the dashed horizontal line) to the demand intercept of 539 at the top. In each period, S1's posted price is represented by a diagonal cross, S2's price is represented by a vertical cross, and the 3 small sellers' prices are plotted as dots.

Notice that seller S2 prices at the upper bound and sells 1 unit for 8 of the first 9 periods. This "signaling" behavior brings all but one of the other sellers up to the reservation value by period 9. In period 10, seller S2 initiates a small and very lucrative price cut, which draws the others down into an orderly series of small price reductions. This decline is temporarily interrupted when S2 returns to the maximum price in periods 14-17. The tight, almost disciplined pattern of declining prices in periods 18-22 is especially striking. These sellers are nowhere near randomizing, but rather, there is controlled breakup of a high degree of tacit collusion. Subsequent signalling by S2 lifts prices prior to a more erratic downward spiral in the

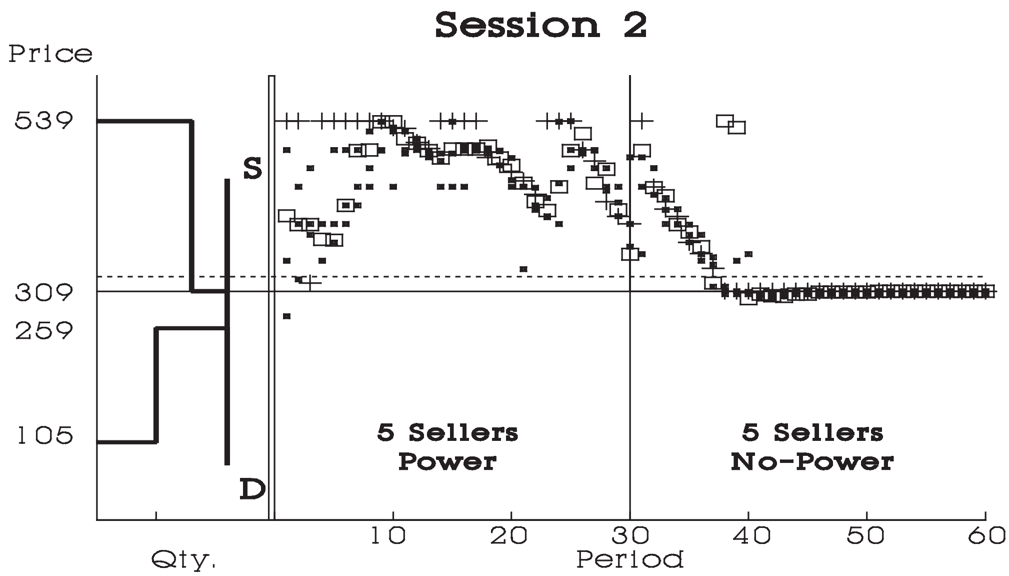

Figure 2. Price Series for a Power/ No-Power Sequence (Key: The box represents *S1*'s price, the cross represents *S2*'s price. Other prices are indicated by dots.)

final 5 periods of the treatment.[17, 18] Nevertheless, prices almost always exceed the theoretical medians for the mixed distributions (of 332 and 347 for the small and large sellers respectively).

After period 30, the announcement that describes the structure of the no-power treatment to follow was read. Prices start high in the no-power treatment, but decline steadily, falling to the competitive level by period 40. This time it is the other large seller, S1, who tries to reverse the decline with high prices in periods 38 and 39, but this signal is very costly in the no-power design, because unilateral price increases generate no sales, and hence, are not profitable. The prices lock on the Nash (and competitive) equilibrium price in periods 40-60.

The remainder of this section is organized about four observations. First, we consider more generally the effects of market power on prices in the five-seller designs. As will be seen,

---

[17] In our sessions we observed many price cycles of the type shown in figure 2. Maskin and Tirole (1988) have shown how price cycles can arise in price-setting games in which firms make price choices in alternating periods that must be maintained for 2 periods. Eaton and Engers (1990) show that the threat of a price-cycle war can sustain collusion in an alternating choice model. However, we know of no noncooperative theory to explain the types of cycles observed our simultaneous-price-choice markets.

[18] The downward trend in prices prior to period 30 suggests an "end effect," but this price pattern was not consistently observed in the power treatment periods of other sessions to be reported below.

the sequence of contracts illustrated in figure 2 is representative of the price effects of market power. The second observation pertains to the pure-numbers effect of reducing the number of sellers from five to three. The final two observations regard the causes and the consequences of the price increases prompted by market power.

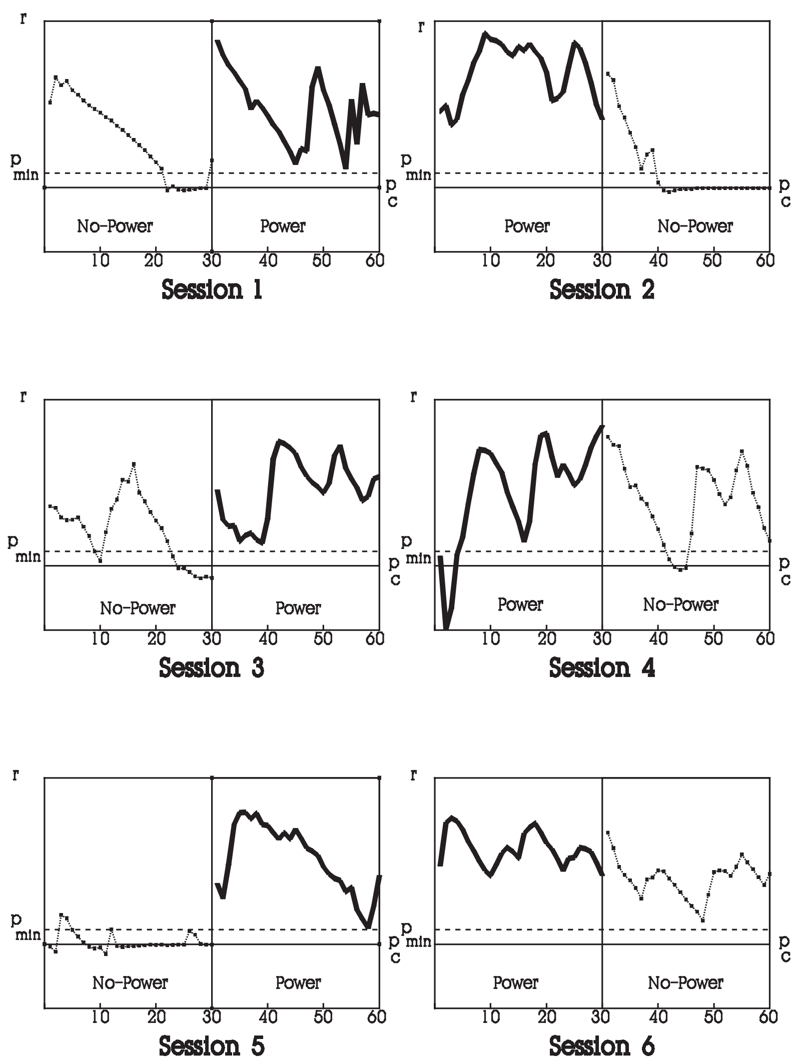

Figure 3. Average price sequences for 6 sessions using designs 1 and 2

Market power effects

The time series of average posted prices for the sessions conducted with designs 1 and 2 are illustrated in the six panels in figure 3. In each panel the average price path for the power treatment is represented by a bolded line, while the average price path for the no-power treatment is represented as a dotted line. Notice that the no-power treatment was first in the three left-hand panels and second in the three right-hand panels. Session numbers are located below the corresponding panels

Table 1. Average Prices*

	No-Power/Power Sessions (with five sellers)						Five-Seller/Three-Seller Sessions (with power)					
Session	1	2	3	4	5	6	7	8	9	10	11	12
No Power (Five sellers)	329	308	341	410	310	397	-	-	-	-	-	-
Power (Five sellers)	407	468	430	455	397	441	415	471	392	401	392	512
Power (Three sellers)	-	-	-	-	-	-	425	470	408	436	424	517
Treatment Effect	78	160	89	45	87	44	10	−1	16	35	32	5

* Reported prices are averages taken over the last 15 periods of each treatment in a session.

It is evident from figure 3 that prices were uniformly high and variable in the power treatment. In contrast, posted prices were much lower and more stable in the no-power treatment, with prices converging to the competitive prediction in four of six cases.[19] The effect of the power treatment is also seen table 1, which summarizes the average price for the last 15 periods of each treatment sequence. The bottom row of table 1 presents the difference between mean prices in the power and no-power treatments, by session. Note that the power treatment raised

[19] We include the no-power treatment in session 1 as an instance of price convergence, despite the small upward spike in period 30. The spike was due to a single high price posting that resulted in no sales.

average prices in all six sessions.[20] The uniformity of results allows us to feel fairly confident of a market power effect. For example, using the data summarized on the left side of table 1, the null hypothesis of no treatment effect can be rejected at a 95% confidence level, using the nonparametric Wilcoxon signed-ranks test.[21] The magnitude of the power treatment effect also deserves comment. While the effect of power varies widely across sessions, the average of the six differences listed in the bottom row on the left side of table 1 is 84 cents, more than one third of the entire range between p_c and r. On the basis of these observations, we draw our first conclusion: *All other things constant, market power tends to raise prices substantially.*

Of course, systematic sequencing effects would call for a qualification of this result. Fortunately, no such qualification is necessary here since there is no evidence of a correlation between the magnitude of the differences recorded in the bottom row of the table, and the order in which the treatments were presented.[22]

Despite the invariance of the power treatment effect to sequencing, note that absolute price levels are affected. Prices are lower when the power treatment comes second in sequence, suggesting that initial experience has a residual effect that persists in the subsequent treatment. In table 1, for example, power came second in sessions 1, 3 and 5, which have the three lowest price averages for the power treatment. Thus, it appears that prior experience in a competitive (no-power) environment exerts a residual tempering effect on pricing in subsequent power

---

[20] We report average prices for only the last 15 periods in each treatment sequence as a means of controlling for learning. The average prices over all 30 periods of each treatment were also higher in the power treatments in each of the six sessions. If the first 15 periods of each treatment are used instead, average prices are higher in the power treatments in five of the six sessions.

[21] The test statistic is 21, which exceeds the 95% critical value of 19. For a description of the test see, e.g., Conover, 1980, p. 280. The intuition behind the reported confidence level is easy to motivate for the special case observed in the table, where prices are higher under the market power treatment in all six sessions. This outcome is analogous to getting six heads in a row from a series of coin tosses. If the coin were fair (e.g., if there was no treatment effect) this event would occur with a probability of one half raised to the sixth power, or about .03.

[22] For example, divide the six price differences reported in the bottom row on the left side of table 1 into power/no-power and no-power/power groups. Since the observations are independent, a Mann-Whitney test is appropriate. To apply this test, the six mean price differences are ranked from highest to lowest. The test statistic is the sum of the ranks for one of the treatment orders, (say the power/no power sequences). The null hypothesis is rejected if the test statistic falls outside a critical value range determined by the number of observations in each group. In this case, the relevant test statistic is 9 and the null hypothesis cannot be rejected at a reasonable level of confidence (p = .35). For a more detailed discussion of the test, see e.g., Conover (1980), p. 218.

periods. Similarly, high prices in an initial power treatment sometimes raise prices in the subsequent no-power treatment. For example, the two instances where prices failed to converge in the no-power treatment, sessions 4 and 6, occurred after the power treatment. Sellers with no power in the last half of these sessions were apparently able to sustain patterns of tacit collusion that were learned in the power treatment in the first 30 periods.

Pure-numbers effects

Consider next the effect of reducing the number of sellers from five to three, holding other factors (including power) constant. The average price sequences for the six sessions conducted with this combination of designs 2 and 3 are shown in figure 4. In each panel, the average prices for the 5-seller power design are plotted as a bolded line, and the average prices for the 3-seller power design are plotted as a dotted line. As before, the order of treatments is reversed in every other session to control for sequence effects.

It is evident from figure 4 that a reduction in the number of sellers exerts a much less pronounced effect than the pure-power effect discussed above. Although there is a rather large difference in prices from one session to another, price performance within sessions appears largely unchanged by the alteration in the number of sellers. But inspection of average price differences across the five person and three person treatments in the right half of table 1 suggests that the reduction in the number of sellers raises prices to some extent.[23] Average prices over the last 15 periods of each treatment sequence increased with the reduction in sellers, in five of the six sessions. Using the data summarized on the right side of table 1, the null hypothesis of no treatment effect can be rejected at just below a 95% confidence level, using the Wilcoxon signed-ranks test.[24] This effect, however, is very small. The average price increase observed in the power/ no-power sessions (83 cents) is nearly 5 times greater than the average increase observed in 5-seller/3-seller sessions (17.3 cents). Moreover, the smallest price change observed

---

[23] Dolbear et al. (1968) and Isaac and Reynolds (1989) also find some evidence of numbers effects in posted-price auctions.

[24] The appropriate test statistic is 20, which exceeds the upper critical value for a 95% confidence level on a 1-tailed test, given a sample size of six.

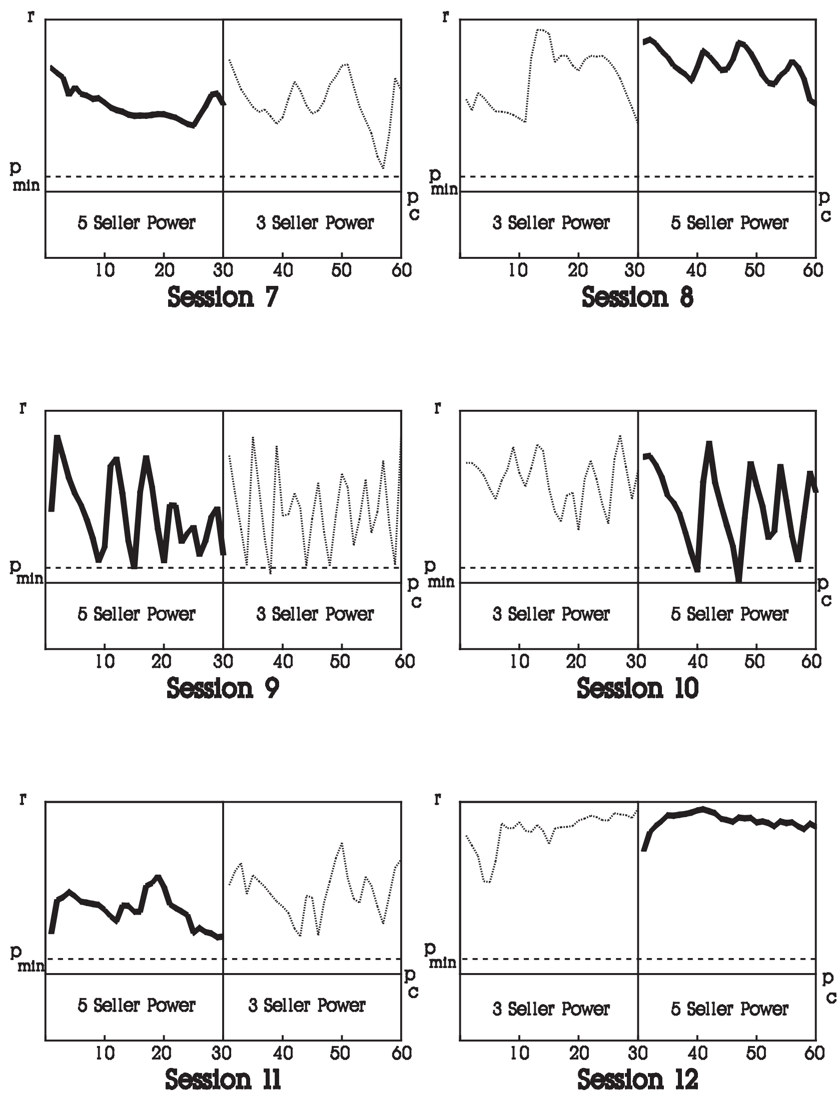

Figure 4. Average price sequences for 6 sessions using designs 2 and 3

in a power/no-power session exceeds the largest price change observed in the 5 seller/3-seller sessions. These observations motivate our second conclusion. *All other things constant, a reduction in the number of sellers from five sellers to three sellers increases prices somewhat. The effect, however, is very small compared to the effect of introducing market power.*

Like the power/no-power sessions, the treatment order has no effect on the magnitude of the price differences listed in the bottom row of table 1.[25] In contrast to the power/no power sessions, group effects are large relative to treatment effects in the five-seller/ three-seller power sessions. In these latter sessions, the range of treatment effects is 36 cents, while price averages vary across cohorts by more than $1.00. Nevertheless, it is reassuring to note that market power appears to have a fairly consistent effect on prices: There is no evidence of a systematic price difference across the 12 five-seller power sessions common to both treatment-pairs.[26]

Finally, we want to be careful to not dismiss too quickly the general possibility of a sizable pure-numbers effect. Although the effect of numbers alone is relatively small in the present design, numbers may play a more prominent role in other contexts, particularly in designs where there is no market power, and hence much less tacit collusion among a relatively large number of sellers. For example, it is easy to imagine duopolists colluding more readily than quadropolists, even if there is no power in either case.

Noncooperative equilibria, tacit collusion and market power

An important question that remains to be addressed regards the reason for the price effect prompted by market power. On the one hand, prices may change with the addition of market power, because of the change in the underlying noncooperative equilibrium. On the other hand, market power may facilitate tacit collusion, or the ability of sellers to maintain prices above the noncooperative equilibrium levels in the absence of explicit communications. Figure 5 presents evidence relevant to this question for the power treatment in which the most data were collected, the 12 five-seller power treatments. This figure shows the theoretical and observed price densities, aggregated across the last half of the five-seller power treatment sequences. Densities for the large sellers are shown in the left panel, while densities for the small sellers are shown

---

[25] The Mann-Whitney test, applied to the six differences reported in the bottom row on the right side of table 1, does not allow rejection of the null hypothesis of no sequencing effect in the three-seller/five-seller sessions. Coincidentally, the test statistic in this case is again nine (p = .35).

[26] Consider, for example, the null hypothesis that mean prices for the five-seller power treatments in sessions 1 to 6 do not differ from mean prices for the five-seller power treatment in sessions 7 to 12. Using data from the relevant row in table 1 and the Mann-Whitney test, the null hypothesis cannot be rejected at any conventional confidence level. The relevant test statistic is 36 (p = .349).

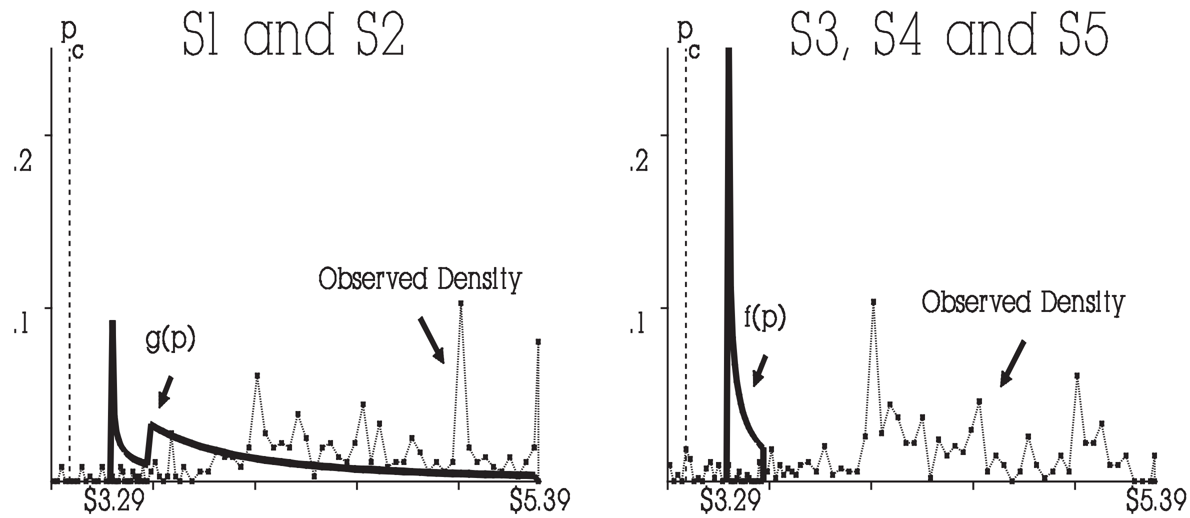

Figure 5. Predicted and observed densities for the 5-seller market power design

in the right panel. Within each panel, the density predicted under randomization with risk neutrality is plotted as a bolded line, while the observed density is plotted as a dotted line.

It is clear from figure 5 that participants did not randomize according to the predictions of the mixed strategy equilibrium, assuming risk neutrality. Although there is some separation between the pricing choices of the large and small sellers, far too much of the observed density weight in either distribution occurs at prices above those predicted under the mixing equilibrium. The observed medians for the large and small sellers are 448 and 415, respectively, far in excess of the medians of 371 and 332 predicted in the mixing equilibrium.[27, 28] Moreover, the

[27] Assuming risk neutrality, the differences in predicted and observed medians are so great as to make a statistical comparison almost superfluous. Any test of the null hypothesis that sellers priced according to the randomizing distribution would be easily rejected at conventional confidence levels. For example, the data summarized in figure 5 generate Kolmogorov-Smirnov test statistics of .265 and .86 for the large and small sellers, respectively. These test statistics easily exceed the respective 95% critical values of.07 and .058.

[28] Risk aversion represents a possible reason for prices to exceed the theoretic distributions that are consistent with a randomizing equilibrium. Holt and Solis-Soberon (1992) have shown that the effect of risk aversion is to produce a stochastic increase in the equilibrium price distributions in the interior of the support. The intuition is that a risk averter prefers the certain return available at the upper bound of the distribution to uncertain returns of equal expected value

obvious autocorrelation of prices suggested by figures 2-4 contradicts the temporal independence implied by randomization.[29] This motivates our third conclusion: *Most of the price increase generated by market power arises from the increased incidence of tacit collusion.*[30]

Finally, observe the variety of behavior consistent with tacit collusion in our markets. The 15-period price averages in table 1 range from 392 to 512 in the twelve 5-seller power sequences. Moreover, it is clear from figures 3 and 4 that prices in some of the tacitly collusive instances are reasonably stable (e.g., sessions 6 and 12) and very volatile in other instances (such as sessions 1, 9 and 10). The range of price patterns observed under tacit collusion highlights the difficulty of drawing inferences about the incidence of tacit conspiracy in natural markets. Given the volatility of prices in many instances, it is often unlikely that sellers would believe that they had engaged in tacitly collusive behavior of any kind. Moreover, no particular pricing pattern, stable or otherwise, separates tacit collusion from competitive behavior. For these reasons, it would be very difficult to detect noncompetitive performance in our markets absent rather detailed information regarding the underlying market structure.

Distributional effects of market power

---

available in the interior of the support under the risk neutral equilibrium. To make the risk averter indifferent over all prices in the randomization range, the probability of being a relatively high-price seller must decrease at each point in the interior of the support and, therefore, the equilibrium price distributions must shift upward. these observations do not change the important implication that randomization implies independent draws from a distribution, and there should be no serial correlation.

[29] Among other things, randomization implies that prices in one period should be independent of prices in the preceding period. But this is clearly not the case for the sessions summarized in the text. OLS regressions of each average price series on its one-period lagged value generate coefficient estimates that are significant at a 99% confidence level in 11 of the 12 sessions. In most instances the coefficient estimate is also quite large, exceeding .75 in 9 of the 12 sessions. For an interesting analysis of the implications of independence in an experimental context, see Brown and Rosenthal (1990).

[30] This observation of cooperative behavior in finitely repeated games is consistent with data from a wide variety of laboratory experiments in other contexts. For example, Andreoni and Miller (1991) found significantly more cooperative behavior in finitely repeated prisoner's dilemma games than in single-period prisoner's games. They attribute this to the presence of some subjects in a population who are genuinely cooperative. Cooper, DeJong, Forsythe, and Ross (1989, p. 580) report the results of single-period coordination games with a symmetric mixed equilibrium, and conclude: "We do find that the theory predicts 'tougher' play than was observed; the actual frequencies of play lie somewhere between those predicted by the theory and those consistent with players adopting cooperative mixed strategies."

Consider now the effects of the price increases on seller profits and on market efficiency. Information regarding the level of profits is summarized in table 2. Columns (1) and (2) of the

Table 2. Seller Profit Ratios (Seller Profit as a Percentage of Total Surplus)*

Power/ No Power			Five-Seller Power/ Three-Seller Power		
Session	(1) No Power (Five Sellers)	(2) Power (Five Sellers)	Session	(3) Power (Five Sellers)	(4) Power (Three Sellers)
1	40	58	7	61	63
2	42	73	8	59	69
3	44	67	9	52	55
4	55	70	10	72	63
5	40	55	11	49	61
6	53	67	12	85	89
Overall	46	65		63	67

* Seller Profit Ratio = $(\pi_{OS}/\pi_{CE})100$, where π_{OS} = observed seller earnings, and π_{CE} = 3160 is the maximum available aggregate surplus. The seller profit ratio is 42% under the competitive prediction, and 49% in the Nash equilibrium for the market power treatment. Ratios reported in the table are averages for the last 15 periods in each treatment.

table report seller profit ratios for the last half of each sequence in the power/ no power treatment. The numbers in these columns represent observed aggregate earnings for all sellers, expressed as a percentage of total possible surplus. (Recall that sellers are predicted to earn 42% of maximum possible gains from trade under the competitive prediction, and 49% of the same maximum in the market-power Nash equilibrium.) Consistent with the convention used in table 1, we report averages for the last 15 periods of each treatment sequence. As is clear from the summary measures printed at the bottom of columns (1) and (2), sellers tended to earn just slightly more than their predicted share in the no-power treatment, with an average ratio value of 46%. Seller earnings rose substantially in the power treatment, however, to an average of 65%, 16 percentage points more than predicted in the Nash equilibrium. Finally, note that the small price increases that result from the consolidation of the three small sellers in the power designs do not translate into substantial earnings increases, as can be seen by a comparison of columns (3) and (4) of table 2. Seller profit ratios for the 5-seller power and 3-seller power treatments average 63% and 67%, respectively, virtually the same as was observed in the five-seller power treatment summarized in column (2).

Table 3. Excess Earnings for Large Sellers; Excess Earnings for Small Sellers
(Expressed as a Percentage of Earnings Predicted in the Nash Equilibrium)*

Power/ No Power			Five-Seller Power/ Three-Seller Power		
Session	(1) No Power (Five Sellers)	(2) Power (Five Sellers)	Session	(3) Power (Five Sellers)	(4) Power (Three Sellers)
1	−3; −7	35; −2	7	31; 18	30; 26
2	0; −1	67; 27	8	15; 9	80; −6
3	−1; 10	42; 32	9	15; −2	31; −13
4	33; 32	64; 17	10	76; 12	37; 21
5	4; −5	22; 0	11	24; −32	34; 15
6	40; 13	51; 21	12	95; 49	87; 76
Overall	11; 7	47; 16		45; 9	50; 20

* Excess Earnings = 100 (π_{Oi} - π_{Ni})/π_{Ni}; i = L, S, where L = "large" sellers (S1 and S2), and S = "small" sellers (S3, S4 and S5, or S3). π_{Oi} = observed earnings for the type i sellers, π_{Ni} = predicted earnings for type i sellers under the relevant Nash equilibrium. Each entry lists excess earnings for both large and small sellers. Reported numbers are averages for the last 15 periods of each treatment.

It is interesting to note that the large sellers are quite clearly the primary beneficiaries of market power. This can be seen in table 3, which reports "Excess Earnings," or the excess of observed profits over profits predicted in the (relevant) Nash equilibrium, as a percentage of predicted profits. These ratios are presented for the large sellers (S1 and S2), and the "small" sellers (S3, S4, and S5, or S_m), for the last 15 periods of each treatment sequence.[31] From the overall averages, listed at the bottom of column (1), it is seen that S1 and S2 earned only a slightly higher proportion of profits than S3, S4 and S5 in the no-power treatment (11% vs. 7%). With the exception of sessions 4 and 6, (which exhibit a carry-over effect from the preceding

[31] In column 1 of table 4, earnings for S3 are combined with those of S4 and S5 in calculating the seller profit ratios for small sellers in the no-power design, despite the fact that S3's allocation is identical to the allocations for S1 and S2 in the no-power design. This was done to facilitate comparison with seller profit ratios for the five-seller power treatment, summarized in column (2).

power treatment sequence), neither group of sellers earned substantially more than those predicted under the competitive equilibrium in this treatment. In contrast, in the power sessions summarized in columns 2, 3 and 4, the proportional increase in earnings over the Nash prediction tended to double for the small sellers, while earnings for S1 and S2 quadrupled, soaring to roughly 50% in excess of profits predicted under the Nash equilibrium. Combined, these observations motivate our third conclusion: *Market power clearly increases seller earnings in this design, and the large sellers are the primary beneficiaries.*

Despite the prominent redistributional effects of market power, power generates little in the way of efficiency consequences here. As mentioned above, this is to be expected, since only 150 in surplus is lost when from the 3 units that fail to trade in the market-power Nash equilibrium. Average efficiency performance for the last 15 periods of each treatment is summarized in columns (1) - (4) of table 4. Efficiency is slightly below equilibrium predictions, and is slightly higher on average in the no-power than in the power sequences. Nevertheless, from the overall means listed at the bottoms of the four columns, mean efficiencies differed by less than 3% across treatments.

A methodological observation

We conclude this section by offering some comments on the benefits of blocking, which in this context involves subjecting each cohort to the two treatments in alternating order, as a means of controlling group and sequencing effects. Blocking is not always essential to identifying a treatment effect. The market-power effect, for example, is sufficiently strong that the null hypothesis of no effect could be rejected even if the power and no-power sequences had been obtained from separate cohorts of individuals in independent, single-treatment sessions.[32] Blocking, however, can greatly enhance the capacity to draw inferences regarding principal treatment effects from a data set with "noisy" differences between cohorts. Consider, for example, the five-seller power / three-seller power sessions. Rather than conducting six pairs of

[32] Suppose, for example, that the six power sequences and the six no-power sequences summarized in table 1 had been conducted as 12 independent trials, each with a different cohort of subjects. Then the null hypothesis of no power effect would be appropriately tested with a Mann-Whitney test. In this case, the test statistic is 54.5, which exceeds critical values at a 99% confidence level, of 24 and 54.

Table 4. Aggregate Efficiency (Surplus Extracted as a Percentage of Total Surplus)*

Power/ No Power			Five-Seller Power/ Three-Seller Power		
Session	(1) No Power (Five Sellers)	(2) Power (Five Sellers)	Session	(3) Power (Five Sellers)	(4) Power (Three Sellers)
1	94	93	7	93	94
2	100	92	8	94	89
3	94	95	9	92	93
4	90	92	10	91	93
5	98	92	11	88	94
6	91	93	12	93	95
Overall	95	93		92	93

* $E = 100(\pi_O - \pi_{CE})/\pi_{CE}$, where π_O = observed aggregate earnings, and π_{CE} = maximum aggregate earnings. E = 100% under the competitive prediction and E = 95.2% under the market-power Nash equilibrium prediction. Reported numbers average performance in the last 15 periods in each treatment.

30-period treatments, we might have conducted 12 independent 30-period sessions, six in the three-seller power design, and six in the five-seller power design. But if 12 independent sessions generated the data summarized on the right side of table 1, we would be unable to reject the null hypothesis of no treatment effect at standard confidence levels.[33] A paired-treatment design is especially useful if the difference across treatments within sessions is small relative to the variability in price sequences across sessions.[34]

[33] The relevant Mann-Whitney test statistic is 46 (p=.15).

[34] Kendall's τ, a nonparametric correlation coefficient based on ranks, may be used to analyze group effects. The coefficient is bound by one (in the event of a perfect direct relationship) and minus one (in the event of a perfect inverse relationship), and has an associated statistical test for the null hypothesis of no correlation. In the present case, analysis of the six data pairs in the three-seller power / five-seller power treatment suggests a group effect; $\tau = .87$ for these data, and the null hypothesis of no group effect may be rejected at above a 95% confidence level (two-tailed test). Group effects are much weaker in the power / no-power sessions. For the six data-pairs conducted in the power, no-power designs $\tau = .33$. For a more detailed discussion of this analysis, see Conover (1980), p. 257.

6. Conclusion

The strength of the correlation between market concentration and price performance has been hotly contested by economists for decades, on both theoretical and empirical grounds. Our data provide a starting point for a resolution of the debate. Concentration and pricing may be correlated because concentration often changes market power. In the laboratory, these conditions can be varied independently. Our data indicate the importance of market power. An increase in concentration that creates theoretic market power prominently affects observed prices. Further consolidation that does not alter market power appears to have little effect in the design used here.

In the treatments where market power was present, observed price increases are much greater than those implied by the change in the underlying noncooperative equilibrium. In every instance, far too much of the density in the observed pricing distribution was in the upper portion of the pricing range for pricing to be consistent with equilibrium randomization. Therefore, the main effect of market power in this design is to promote tacit collusion. Finally, the variability of persistently supra-competitive price series suggests that it is exceedingly difficult to distinguish tacit collusion from competitive performance without detailed information regarding the underlying market structure.

Appendix

Presentation of Market Information to Participants.

An effort was made to provide full information to participants in a way that was understandable. An initial announcement at the outset of the experiment described the number of periods under each treatment and the ten-to-one conversion of lab dollars to actual cash earnings. The following announcement described purchasing behavior:

> **Buyer Information**: The experiment monitor will simulate buyer decisions in today's experiment. The monitor will make purchases according to the following rules:
>
> a) **The monitor will purchase a total of 11 units, if all 11 units are available at prices of \$3.09 or less. The monitor will purchase a total of 8 units, if 8 units are available at prices of \$5.39 or below.** The monitor will purchase no units priced above \$5.39.
>
> b) The monitor will always purchase the lowest priced units first, followed by the next lowest priced units, and so on.
>
> c) In the event two or more sellers post the same price, the monitor rolls a die until the identity number of one of the tied sellers is tossed. The monitor then rotates purchases among the tied sellers, buying first from the seller whose identifier number was rolled, then from the tied seller with the next smallest identifier, and so forth, down to the tied seller with the smallest identifier. If possible, the monitor the makes a purchase from the tied seller with the largest identifier number, and the seller with the next largest identifier, etc., until purchases have been made from all tied sellers. The monitor repeats this rotation until all tied sellers are out of units, or until all available units have been purchased.
>
> An example may be instructive. Suppose *S2*, *S3*, *S4* and *S5* all post the same price. The monitor would then roll a die. If a 3 is rolled, the monitor purchases from *S3*, *S2*, then *S5* and *S4*. The monitor repeats this sequence if additional units are available.

The latter tie-breaking procedure was implemented to expedite administration of the sessions, and does not affect the equilibria.

The announcements pertaining to demand conditions and the number of periods were repeated prior to starting each 30-period treatment sequence. Supply conditions were also described prior to each treatment. The announcement for the structure of the no-power treatment read:

Information about other Sellers: There are five sellers. The sellers are not identical.

Seller Identity	Total Available Units	Cost of First Unit	Cost of Remaining Units
S1	3	$1.05	$2.59
S2	3	$1.05	$2.59
S3	3	$1.05	$2.59
S4	1	$1.05	
S5	1	$1.05	

Thus, a total of 11 units may be offered by all sellers combined: 6 units cost $2.59 each, and 5 units cost $1.05 each.

The announcement describing the market structure of the 5-seller, power treatment read:

Information about other Sellers: There are five sellers. The sellers are not identical.

Seller Identity	Total Available Units	Cost of First Unit	Cost of Remaining Units
S1	4	$1.05	$2.59
S2	4	$1.05	$2.59
S3	1	$1.05	
S4	1	$1.05	
S5	1	$1.05	

Thus, a total of 11 units may be offered by all sellers combined: 6 units cost $2.59 each, and 5 units cost $1.05 each.

References

Alger, D. "Laboratory Tests of Equilibrium Predictions with Disequilibrium Data." *Review of Economic Studies*, Vol. 54 (1987), pp. 105-145.

Andreoni, J., and Miller, J.H. "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence." University of Wisconsin-Madison, SSRI Workshop Paper 9102, March 1991.

Brown, J.N., and Rosenthal, R.W. "Testing the Minimax Hypothesis: A Re-Examination of O'Neill's Game Experiment." *Econometrica*, Vol. 58 (1990), pp. 1065-1081.

Brown-Kruse, J., Rassenti, S.J., Reynolds, S.S. and Smith, V.L. "Bertrand-Edgeworth Competition in Experimental Markets." Mimeo, University of Arizona, 1990.

Cooper, R., DeJong, D.V., Forsythe, R. and Ross, T.W. "Communication in the Battle-of-the-Sexes Game: Some Experimental Results." *Rand Journal of Economics*, Vol. 20 (1989), pp. 568-587.

Conover, W.J. *Practical Nonparametric Statistics*. New York: John Wiley and Sons, 1980.

Davis, D.D., Holt, C.A., and Villamil A.P. "Supra-Competitive Prices and Market Power in Posted-Offer Auctions." BBER Discussion Paper, University of Illinois, 1990.

______ and Williams, A.W. "The Effects of Market Information on the Emergence of a Pure Bargaining Equilibrium." working paper, Virginia Commonwealth University, 1989.

______ and ______ "The Hayek Hypothesis in Experimental Auctions: Institutional Effects and Market Power." *Economic Inquiry*, Vol. 29 (1991), pp. 261-274.

Dolbear, F.T., Lave, L., Bowman, L.B., Lieberman, A., Prescott, E., Rueter, F. and Sherman, R. "Collusion in Oligopoly: an Experiment on the Effect of Numbers and Information." *Quarterly Journal of Economics*, Vol. 82 (1968), pp. 240-259.

Eaton, J. and Engers, M. "Intertemporal Price Competition." *Econometrica*, Vol. 58 (1990), pp. 637-659.

Grether, D.M., and Plott, C.R. "The Effects of Market Practices in Oligopolistic Markets: An Experimental Examination of the *Ethyl* Case." *Economic Inquiry*, Vol. 22 (1984), pp. 479-507.

Holt, C.A. "An Experimental Test of the Consistent-Conjectures Hypothesis." *American Economic Review*, Vol. 75 (1985), pp. 314-325.

______ "The Exercise of Market Power in Laboratory Experiments." *Journal of Law and Economics*, Vol. 32 (1989), pp. S107-S130.

______, Langan, L. and Villamil, A.P. "Market Power in Oral Double Auctions." *Economic Inquiry*, Vol. 24 (1986), pp. 107-123.

______, and Solis-Soberon, F. "The Calculation of Mixed-Strategy Equilibria in Posted-Offer Markets." in R.M. Isaac, ed., *Research in Experimental Economics, Vol. 5*. Greenwich: JAI Press.

Isaac, R.M., and Reynolds, S.S. "Two or Four Firms: Does It Matter?" Mimeo, University of Arizona, June 1989.

Ketcham, J., Smith, V.L. and Williams, A.W., "A Comparison of Posted-Offer and Double-Auction Pricing Institutions." *Review of Economic Studies*, Vol. 51 (1984), pp. 595-614.

Maskin, E. and Tirole, J. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles." *Econometrica*, Vol 56 (1988), pp. 571-599.

Welford, C.P. "Horizontal Mergers: Concentration and Performance," in *Takeovers and Horizontal Mergers: Policy and Performance*, Ph.D. dissertation, Department of Economics, University of Arizona, 1990.