Vertical foreclosure in experimental markets

Stephen Martin[*]

Hans-Theo Normann[**]

and

Christopher M. Snyder[***]

We report the results of experiments designed to test recent theories of vertical foreclosure. Consistent with the theory, vertical integration improves the upstream firm's ability to commit to restricting output to the monopoly level, as does the use of public contracts. Public contracts are not a perfect substitute for vertical integration, however: integration allows more surplus to be extracted from the unintegrated downstream firm, a bargaining effect that has been underemphasized in the recent foreclosure literature. Motivated by some observations that are difficult to reconcile with existing theory, we extend the theory to allow downstream firms to have heterogeneous (rather than purely passive or symmetric) out-of-equilibrium beliefs.

1. Introduction

■   Antitrust treatment of vertical integration and vertical restraints has followed a policy pendulum.[1] In the early years of U.S. antitrust, little attention was paid to vertical relationships. From the late 1950s through the early 1970s, antitrust authorities took a more hostile view, blocking vertical mergers between General Motors and duPont, Brown Shoe and Kinney, and Ford and Electric Autolite.[2] The policy of the 1950s and 1960s was motivated by the idea that vertical

* University of Amsterdam; smart@fee.uva.nl.

** Humboldt University; normann@wiwi.hu-berlin.de

*** George Washington University; csnyder@gwu.edu.

[1] See Perry (1989), White (1989), and Kwoka and White (1999) for surveys. As Kwoka and White note, the policy shifts are evident in successive revisions of the U.S. Department of Justice *Merger Guidelines*: vertical mergers occupied an extensive section in 1968, less in the 1982 and 1984 revisions, and are absent entirely from the 1992 revision.

[2] See *U.S.* v. *E.I. duPont de Nemours & Co.*, 334 U.S. 495; *Brown Shoe Co.* v. *U.S.*, 370 U.S. 294 (1962); and *Ford Motor Co.* v. *U.S.*, 405 U.S. 562 (1972). During the same period, the Federal Trade Commission blocked a series of

integration can harm downstream firms by removing a source of inputs and upstream firms by removing an outlet for supply, an idea that has been labelled "naive foreclosure theory" because the linkage between harm to competitors and harm to consumers lacked a rigorous, formal basis. In the 1980s, antitrust authorities returned to a more benign view of vertical relationships. This view reflected the ascendancy of the Chicago School position, summarized by Bork (1978) among others, that market power in one market cannot profitably be leveraged into other markets.[3] The Chicago School approach meshed with a variety of efficiency explanations of tight vertical relationships.[4]

The benign view of tight vertical relationships has been challenged by a series of recent articles that have provided a formal basis for many of the conclusions of the naive foreclosure theory.[5] This literature develops game-theoretic models of vertically related oligopolies and shows that vertical integration may be associated with a variety of strategic effects, reducing output of the final good, increasing its price, and potentially reducing social welfare. A common theme of this work is that tight vertical relationships are not so much a device for leveraging existing market power to earn super-monopoly profits as a precondition for the exercise of market power at the upstream level. Without tight vertical relationships, a monopoly producer may be unable to make a credible commitment to downstream distributors that it will restrict output. Absent such a commitment, downstream firms will not, in equilibrium, accept terms of exchange that allow the producer to extract full monopoly profit. This commitment problem is similar to the situation of a durable-good monopolist that cannot commit to maintain high prices in later periods.[6] In that context, the inability to commit limits the price consumers will pay for the durable good in early periods. In the context of vertical relationships, the inability of a producer to commit not to sell to multiple distributors, or at least to restrict the output sold to each, limits the willingness of any one distributor to accept its portion of an overall package of offers that would yield the producer monopoly profit. Vertical integration—or any vertical contract that has the effect of credibly limiting the producer's actions—resolves the commitment problem and allows the producer to fully exploit its market power.

The theory has had only limited influence on antitrust enforcement until recently, perhaps because it took time to percolate from academic journals to policy circles.[7] Skepticism about the robustness of the models' results[8] and the paucity of direct empirical tests of the models' predictions[9] may have been factors. The experiments that are reported here are intended as a first attempt to provide such direct evidence.

We defer summarizing the results to the conclusion and instead emphasize the article's three novel contributions. First, the article contributes empirical evidence to the policy debate on vertical foreclosure, a debate that has not been informed by much empirical analysis, perhaps because of the difficulty in obtaining the necessary industry data. Second, the article contributes to the

---

vertical mergers between cement and concrete producers and undertook an investigation into the vertical structure of the petroleum industry.

[3] Martin (1999) provides a counterargument regarding horizontally related markets.

[4] See Vernon and Graham (1971), Williamson (1971), Schmalensee (1973), Warren-Boulton (1974), Gould (1977), and Klein, Crawford, and Alchian (1978).

[5] See Salinger (1988, 1989), Ordover, Saloner, and Salop (1990), Hart and Tirole (1990), O'Brien and Shaffer (1992), Bolton and Whinston (1993), McAfee and Schwartz (1994), Riordan (1998), Segal (1999), Choi and Yi (2000), Church and Gandal (2000), and Rey and Tirole (forthcoming).

[6] See Tirole (1988) for a summary of the theoretical literature on the durable-good monopolist and Güth, Ockenfels, and Ritzberger (1995) and Reynolds (2000) for experimental tests of the theory.

[7] See Klass and Salinger (1995) for a discussion of the relationship of the theory to recent antitrust actions against vertical mergers.

[8] See Carlton (1990), Williamson (1990), Reiffen (1992), Klass and Salinger (1995), Reiffen and Vita (1995), and Riordan and Salop (1995).

[9] Rosengren and Meehan (1994) and Snyder (1995) provide evidence from event studies. Other econometric studies (McBride, 1983; Grimm, Winston, and Evans, 1992; Waterman and Weiss, 1996; Chipty, 2001) examine the effects of vertical integration but cannot be interpreted as direct tests of the new foreclosure theories; see Snyder (1995) for a discussion. Allen (1971) and Pass and Hawkins (1972) provide two early case studies.

literature on experimental economics by constructing the first market experiments with a fully specified vertical industry structure. We are aware of only one other experimental article with a vertical industry structure, Mason and Phillips (2000). In their work, the link between the industry levels is incompletely specified in that upstream firms' profits are independent of downstream firms' actions.[10] Third, the article contributes to the theoretical literature on vertical foreclosure by developing a model in which parties' out-of-equilibrium beliefs are heterogeneous (rather than being purely passive or purely symmetric, as has traditionally been assumed). The new model succeeds at explaining some interesting experimental behavior that is at odds with the existing theory.

2. Commitment problem

■ In their survey of the literature on new theories of vertical foreclosure, Rey and Tirole (forthcoming) provide a model that can be used to illustrate the commitment problem. Consider the market depicted in panel A of Figure 1 in which an upstream firm U produces an input at constant average and marginal cost normalized to zero. Each unit of input is costlessly converted by downstream firms D_1 and D_2 into a unit of a homogeneous final good. The final good is sold to consumers with decreasing inverse demand $P(q)$, where $q = q_1 + q_2$ and q_i is the output of D_i. D_1 and D_2 engage in Cournot competition.[11] U makes take-it-or-leave-it contract offers (x_i, T_i) simultaneously to D_i, $i = 1, 2$, where x_i is the quantity of input offered and T_i is the payment required for the bundle.[12] This contractual form allows for general nonlinear tariffs.

If the contracts are public, meaning that D_1 observes the contract U offers to D_2 and vice versa, then U can obtain the monopoly profit for itself by offering contracts $(q^m/2, \Pi^m/2)$ to each D_i, where q^m and Π^m respectively denote the monopoly quantity and profit in the market. Each D_i earns zero profit net of the transfer to U if they both accept the contracts, convert all of their purchased inputs into final output, and sell all of this output; and they can earn no more than zero from any other strategy. Thus it is an equilibrium for them to do so.

If the contracts are secret, meaning that a downstream firm does not observe the contract U offers to its rival, it may no longer be possible for U to obtain the monopoly profit. Suppose U continues to offer $(q^m/2, \Pi^m/2)$ to D_2, and D_2 continues to accept the contract. U and D_1 can increase their joint profit by deviating from the contract $(q^m/2, \Pi^m/2)$. To see this, define $R(q_i, q_j) \equiv q_i P(q_i + q_j)$ and $b(q_j) \equiv \mathrm{argmax}_{q_i} R(q_i, q_j)$. $R(q_i, q_j)$ is the revenue (and also profit, since production is costless) generated by the sale of q_i units by D_i given that D_j sells q_j; $b(q_j)$ is D_i's best response (in terms of revenue maximization) to output q_j. The optimal deviation for U and D_1 would be a contract specifying $b(q^m/2)$ rather than $q^m/2$ units of input. It can be shown that $b(q^m/2)$ is strictly higher than $q^m/2$,[13] and thus that market output $q^m/2 + b(q^m/2)$ would be greater than q^m due to this deviation. Foreseeing that this deviation would reduce its gross profit below $\Pi^m/2$, D_2 would reject U's contract offer in the first place, since its profit net of the transfer to U would be negative.

With secret contracts, there are multiple perfect Bayesian equilibria (PBE). The equilibrium outcome depends crucially on how, in response to receiving an out-of-equilibrium contract offer, a downstream firm updates its beliefs concerning the contract offered to its rival. The two leading

---

[10] Another significant difference is that Mason and Phillips effectively restrict attention to linear contracts (appropriate for their focus on the double marginalization problem), whereas contracts are fully nonlinear in our setup. Theory suggests that vertical integration should have the opposite effect in the two settings, and this is indeed what a comparison of the results reveals: vertical integration increases output in their setting but reduces output in ours.

[11] Rey and Tirole note that the same outcome is obtained if D_1 and D_2 are assumed to engage in Bertrand-Edgeworth competition (competition in prices with homogeneous products and capacity constraints determined by the quantity of inputs purchased from U) as long as upstream marginal cost is high relative to downstream. Since—for simplicity—we have normalized upstream cost to zero, we adopt—for consistency—the assumption of Cournot competition.

[12] In equilibrium it will turn out that $x_i = q_i$, but there is a need to distinguish inputs from outputs out of equilibrium, a leading example being when a downstream firm rejects U's contract offer.

[13] The result is proved in Rey and Tirole for the case of concave inverse demand. The argument for general demand functions is similar to that found in Mas-Colell, Whinston, and Green (1995).

FIGURE 1

MARKET'S VERTICAL STRUCTURE
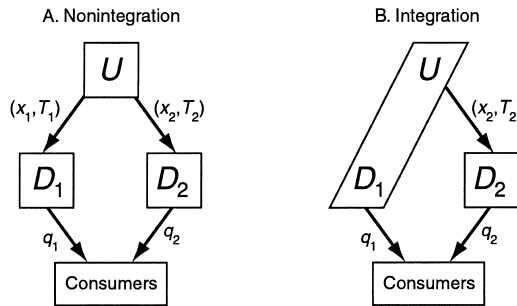
alternatives considered in the literature are passive beliefs and symmetric beliefs.[14] If a firm that holds passive beliefs receives an out-of-equilibrium offer, it conjectures that its rival receives the equilibrium offer. If a firm that holds symmetric beliefs receives an out-of-equilibrium offer, it conjectures that its rival receives the same out-of-equilibrium contract offer. Rey and Tirole show that in equilibrium with passive beliefs, downstream firms receive contracts (q^c, Π^c), where q^c and Π^c are, respectively, a firm's output and profit in a Cournot duopoly; e.g., q^c is the implicit solution to $q_i = b(q_i)$ and $\Pi^c \equiv R(q^c, q^c)$. The same logic as in the previous paragraph applies to show that U cannot commit to sell only the monopoly output to downstream firms. Thus, with passive beliefs, output is higher, and profits lower, than in the joint-profit-maximizing outcome. In equilibrium with symmetric beliefs, by contrast, downstream firms receive contracts $(q^m/2, \Pi^m/2)$, and the joint-profit-maximizing outcome is obtained. U's commitment power in this case stems from the fact that deviating to a contract with higher quantities is not as profitable with symmetric as with passive beliefs. Rey and Tirole argue that the assumption of passive beliefs is theoretically more sound.

The commitment problem arising in the passive-beliefs case can be solved if U integrates vertically with one of the downstream firms, say D_1, forming the firm U-D_1 as in panel B of Figure 1. The integrated firm can commit to sell the monopoly quantity through its downstream subsidiary and not supply D_2 at all, since it internalizes (through profit sharing) the loss in D_1's profit associated with an increase in D_2's supply. Other vertical restraints short of full integration would suffice to ensure that output is restricted to the monopoly level and that it extracts all of the monopoly profit. Exclusive-dealing contracts (see Hart and Tirole, 1990), most-favored-customer clauses (see DeGraba, 1996), and resale price maintenance (see O'Brien and Shaffer, 1992) can be sufficient. Indeed, as argued above, simple contracts of the form (x_i, T_i) are sufficient as long as they are public (and not secretly renegotiated). The fact that D_1 can see the contract offered D_2 before it decides whether to accept provides U all the commitment that is needed. The logic of folk theorems for repeated games with discounting or incomplete information (see Fudenberg and Maskin, 1986) suggest that U may have some commitment power even with secret contracts if the game is played repeatedly: in a finite game, U may be able to maintain a reputation for low output; in an infinitely repeated game, trigger strategies may be used to keep U from deviating.

3. Experimental design

■   To test the implications of the new foreclosure theories summarized in the previous section, we constructed an experimental market with two vertical levels, resembling Figure 1. We conducted treatments with and without vertical integration, holding all other elements constant, allowing us to determine whether vertical integration results in reduced output and increased

---

[14] See Hart and Tirole (1990), McAfee and Schwartz (1994), and Rey and Tirole (forthcoming).

upstream and industry profit, the traditional symptoms of vertical foreclosure. We conducted treatments in a nonintegrated market in which the nature of contracts (secret versus public) was varied, allowing us to focus more narrowly on the question of whether the secretness of contracts is the source of a commitment problem for upstream monopolists.

A host of design issues arose in testing the new foreclosure theories experimentally, including the number of upstream and downstream firms, the form of contract offers, the nature of the bargaining process over contracts, the mode of downstream competition, the form of costs, and demand. We designed the experiments reported here in the spirit of resolving each of these issues in the simplest possible way to provide the most straightforward test of the hypotheses under investigation.

□ **Nonintegrated treatments.** A market in a nonintegrated treatment involves three subjects: one playing the role of the upstream firm U and two playing the roles of downstream firms (D_i, $i = 1, 2$), as in panel A of Figure 1.[15] For simplicity, there are no production costs at either vertical level. At the start of each period, U makes take-it-or-leave-it offers (x_i, T_i) to the D_i, where $x_i \in \{1, 2, 3, \ldots\}$ is a discrete amount of input transferred to D_i and T_i is the lump-sum payment from D_i to U for the bundle.[16] U also has the option of offering no contract to D_i. In the next stage, downstream firms simultaneously choose whether to accept the offer ($a_i = 1$) or reject the offer ($a_i = 0$). If D_i accepts, each of the x_i units of input is automatically transformed into a unit of output and sold at the market-clearing price on the final-good market. Letting q_i be D_i's output, we have $q_i = a_i x_i$.[17]

The total quantity of the homogeneous final good produced by the D_i is denoted $q = q_1 + q_2$. Inverse demand for the final good, $P(q)$, is the discrete function given in Figure 2. The prices are denominated in ECU (experimental currency units), where we adopted an exchange rate of 20 ECU to one DM (or about 34 ECU to the U.S. dollar at the going exchange rate). The demand curve was designed so that the prediction for treatments in which U theoretically has commitment power (e.g., with vertical integration or restraints) differs markedly from that for treatments in which U does not have commitment power (e.g., with nonintegration and secret contracts). With our particular demand function, the joint-profit-maximizing outcome involves $q = 2$, but the PBE with secret contracts and passive beliefs involves $q = 4$.[18]

□ **Integrated treatment.** An experimental market in an integrated treatment involves two subjects: one playing the role of the integrated firm $U\text{-}D_1$ and one playing the role of the un-integrated downstream firm D_2, as in panel B of Figure 1. At the start of each period, $U\text{-}D_1$ makes a take-it-or-leave-it contract offer (x_2, T_2) to D_2. (It can also choose not to offer D_2 any contract at all.) D_2 then accepts or rejects the contract. As above, if D_2 accepts, it automatically produces $q_2 = x_2$ units of output. After observing D_2's accept/reject decision, $U\text{-}D_1$ decides the amount of output, q_1, to produce. The assumption of costless production and the discrete demand specification are the same as in the nonintegrated treatments.

□ **Other aspects.** We ran four different types of treatments, summarized in Table 1. For each treatment type, we ran three sessions, with a different group of subjects recruited for each session. Each session involved three markets and ten periods of play. A possible concern was that the repetition of play for ten periods—designed to allow the subjects an opportunity to learn about the game and strategic responses they typically receive—might introduce dynamic, strategic

---

[15] In the instructions, upstream firms were referred to as "manufacturers" and downstream firms as "retailers."

[16] We restricted x_i to be an integer between 1 and 9 and T_i to be an integer between 0 and 120. The constraints did not impair the generality of the results because both equilibrium contracts in theory and actual contracts in the experiments were well within the constraint set.

[17] We removed the option to purchase x_i units of input but sell $q_i < x_i$ units of output to simplify D_i's decision problem. This should not impair the generality of the results since, in equilibrium, $q_i = x_i$ if $a_i = 1$.

[18] These claims are proved formally in Appendix A.

FIGURE 2

DISCRETE DEMAND FUNCTION
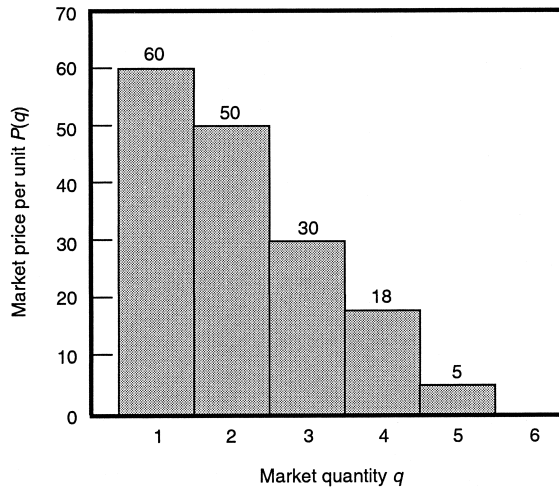
effects. Such effects are absent from the theory, since the theory is based on one-shot games. To mitigate this problem, we used a pseudo-random matching scheme in all but one treatment type. For the nonintegrated treatments, for example, nine subjects were recruited for a session. Upon arrival, three subjects were randomly designated as upstream firms and the remaining six as downstream firms. They were told that triples of one upstream subject and two downstream subjects would be scrambled from period to period. For the integrated treatment, six subjects were recruited for a session, three randomly chosen for the integrated firm and the remaining three for the unintegrated downstream firm. They were told that pairings of one integrated and one unintegrated subject would be scrambled from period to period. The matching process was random in that subjects interacted anonymously through computer terminals and were not told with whom they were matched in a given period as a result of the scrambling. We refer to the matching process as "pseudo" random because we preprogrammed the matching pattern rather than using random draws.[19]

To examine the interesting possibility that repeated interaction might lead to commitment power for the upstream firm through reputation or other dynamic, strategic effects, we introduced a treatment, *SECFIX*, in which the same group of participants met together in a market for all ten periods. This treatment could be compared with the *SECRAN* treatment with pseudo-random matching.

The experimental design does not exclude the possibility of negative profit for downstream firms, since D_i can agree to transfer more to U than the revenue D_i earns from the sale of the final good. The fact that D_i's revenue depends on the quantity D_j supplies, which in turn depends on secret negotiations between U and D_j in the *SECFIX* and *SECRAN* treatments, makes the possibility even more distinct. We took several steps to deal with the problem of downstream bankruptcy. Each subject playing a downstream firm received an initial endowment of 200 ECU. Subjects were told in advance that the experiment would last ten periods and that if a subject showed a negative balance for three successive periods, the session would be halted. In implementation, bankruptcy was never even remotely a prospect. Since negative profits for upstream firms were precluded by the experimental design, upstream subjects were given a smaller initial endowment of

---

[19] We programmed the matching pattern to yield the greatest amount of scrambling possible in a ten-period session. In nonintegrated (integrated) sessions, no two subjects participated together in the same market in more than two (four) periods.

TABLE 1   Treatment Types

Acronym for Treatment Type	Vertical Structure	Subject Matching	Contract Offers
INTEGR	U-D_1 integrated	Pseudo-random	—[a]
PUBLIC	Nonintegration	Pseudo-random	Public
SECFIX	Nonintegration	Fixed groups	Secret
SECRAN	Nonintegration	Pseudo-random	Secret

[a] Contract offers from U-D_1 to D_2 in INTEGR are public by definition, since all relevant players are parties to the contract. The fact that U-D_1 and D_2 choose final output simultaneously implies that internal contracts within U-D_1 are not observable by D_2.

60 ECU, both to conserve on resources and to equalize earnings across upstream and downstream subjects.[20]

The experiments were carried out at Humboldt University Berlin. In total, 99 subjects participated. Subjects were undergraduate economics majors, primarily in the early portion of their studies. The sessions lasted between 45 and 60 minutes, including the time for instructions. On average, subjects earned 21 DM, or about 12 dollars at the going exchange rate. The experiments were programmed using the software toolbox RatImage (Abbink and Sadrieh, 1995).

□ **Summary of experimental hypotheses.** In sum, the experiments were designed so that the theory from Section 2 predicts different outcomes in alternative treatments. In INTEGR, the monopoly outcome is predicted, involving market quantity $q = 2$ and market profit $\Pi = 100$. The same outcome is predicted in PUBLIC. In SECRAN, market output is predicted to be above the monopoly level ($q = 4$) and profit below the monopoly level ($\Pi = 72$). In all three treatments, theory predicts that the upstream firm (or integrated firm in INTEGR) obtains all market profit and unintegrated downstream firms obtain zero profit net of transfers to the upstream firm. The theory does not have direct predictions for SECFIX, since the theory is static, while SECFIX may involve dynamic effects stemming from repeated interaction. We will postpone the discussion of SECFIX until the end of Section 4. These predictions for INTEGR, PUBLIC, and SECRAN are summarized in the following three hypotheses:

Hypothesis 1.   In INTEGR, $q = 2$, $\Pi_U = \Pi = 100$, $\Pi_D = 0$.

Hypothesis 2.   In PUBLIC, $q = 2$, $\Pi_U = \Pi = 100$, $\Pi_D = 0$.

Hypothesis 3.   In SECRAN, $x = q = 4$, $\Pi_U = \Pi = 72$, $\Pi_D = 0$.

We have defined x to be the total amount of input offered in the market ($x = x_1 + x_2$), q to be the market output ($q = q_1 + q_2$), Π to be the market profit ($\Pi = \Pi_U + \Pi_D$), Π_U to be the upstream profit including transfers (integrated firm profit in INTEGR; U's profit otherwise), and Π_D to be the unintegrated downstream firm profit net of transfers (D_2's profit in INTEGR; D_1 and D_2's profit otherwise). Note that the theory does not have unique predictions concerning x in INTEGR and PUBLIC but does for SECRAN, and this is reflected in the statement of the hypotheses.[21]

---

[20] Initial endowments were private knowledge (on this point, see Smith, 1982). Our expectation was that upstream subjects would earn more than their downstream counterparts during the course of a session because upstream firms make take-it-or-leave-it offers. The higher initial endowment would compensate for this difference.

[21] It is plausible that $x = q = 2$ in equilibrium, but it is also possible for higher values of x to arise in equilibrium. In INTEGR, U-D_1 may transfer $x_1 = 2$ internally while offering (x_2, T_2) to D_2 with $x_2 \geq 1$ but with T_2 set so high that the contract is certainly rejected by a profit-maximizing D_2 that understands the structure of the experiments. In PUBLIC, U may offer two contracts, both involving positive inputs, one of which involves $x = 2$ and the other of which involves such a high transfer that, again, it is certainly rejected. In SECRAN, the theory predicts $x = 4$, reflected in Hypothesis 3.

4. Tests of foreclosure theories

■ **Central results.** Table 2 provides descriptive statistics for the variables involved in Hypotheses 1-3. To purge the data of noise in the early stages of each session as subjects became familiar with the experiment, we dropped observations from periods 1–5 before computing the statistics reported in the table. As will be seen below, the results are similar if all ten periods are analyzed. This leaves 45 observations per treatment: three sessions, three markets per session, and five periods (i.e., periods 6–10) per session. The first set of columns in the table are descriptive statistics for all 45 observations associated with each treatment type. Recalling that $q_i = a_i x_i$, the fact that q_i is substantially less than x_i on average for all treatment types indicates that the D_i rejected a substantial number of contract offers. Indeed, in 60% of the *INTEGR*, 40% of the *PUBLIC*, and 44% of the *SECRAN* observations, at least one downstream firm rejected an offer. One explanation for the high rejection rate in *INTEGR* is that U-D_1 intended its offers to produce the same outcome as no contract offer, fully expecting them to be rejected. Other explanations for rejections in *INTEGR*—which also apply to *PUBLIC* and *SECRAN*—are that (a) the upstream firm was playing suboptimally, (b) the upstream firm was trying (unsuccessfully, given the rejection) to earn super-equilibrium profits by exploiting mistakes by downstream firms, and (c) the downstream firms sometimes rejected offers out of "fairness" considerations (i.e., rejecting offers providing positive expected profit if their share of total expected profit were unacceptably low). We will investigate the validity of the various explanations in more detail below. To obtain some preliminary idea of the effect of these rejections on the descriptive statistics, the

TABLE 2   **Descriptive Statistics**

Treatment	All Observations					Excluding Observations in Which a Downstream Firm Rejects Contract Offer				
	x	q	Π	Π_U	Π_D	x	q	Π	Π_U	Π_D
INTEGR										
μ	2.98	2.20	94.3	90.4	3.9	2.67	2.67	92.4	82.6	9.8
σ	(.62)	(.63)	(11.4)	(13.9)	(9.3)	(.69)	(.69)	(8.9)	(11.9)	(12.7)
N	45	45	45	45	45	18	18	18	18	18
PUBLIC										
μ	2.31	1.71	86.0	56.6	29.4	2.04	2.04	99.6	71.7	28.0
σ	(.67)	(.63)	(27.8)	(23.5)	(17.0)	(.19)	(.19)	(1.9)	(12.8)	(12.6)
N	45	45	45	45	45	27	27	27	27	27
SECFIX										
μ	2.91	2.38	75.4	47.9	27.5	3.03	3.03	83.1	57.6	25.5
σ	(1.08)	(1.34)	(25.9)	(23.9)	(20.0)	(1.09)	(1.09)	(20.5)	(21.8)	(20.4)
N	45	45	45	45	45	30	30	30	30	30
SECRAN										
μ	3.44	2.47	69.9	53.0	16.9	3.44	3.44	75.8	70.5	5.3
σ	(1.12)	(1.52)	(33.6)	(26.7)	(31.1)	(1.16)	(1.16)	(26.0)	(26.0)	(32.1)
N	45	45	45	45	45	25	25	25	25	25

Notes: Means (μ), standard deviations (σ), and number of observations (N) for data after dropping first five periods to purge potential learning effects. Figures for x in *INTEGR* include both internal transfers for U-D_1 (measured by its final sales) and the quantity offered to D_2.

FIGURE 3

HISTOGRAMS FOR MARKET INPUTS OFFERED

Key: ☐ All periods ▨ Periods 6–10
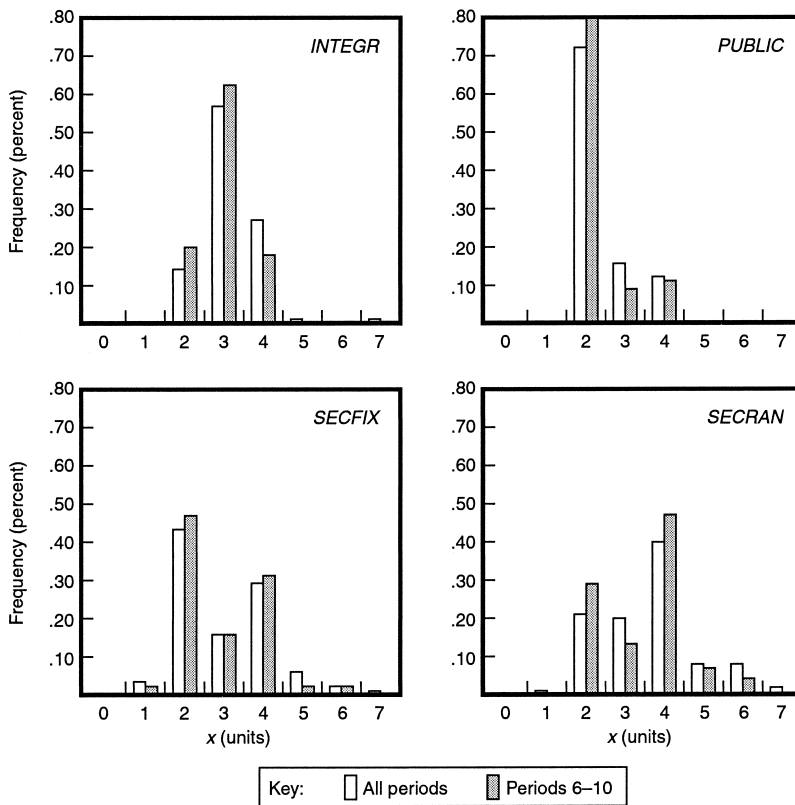
second set of columns excludes observations in which at least one downstream firm rejected its contract. Excluding rejections in this way resolves some apparent anomalies in the data. For example, in *SECRAN*, $q = 2.47$, close to the monopoly output of $q = 2$, though industry profit is only 69.9, far from the monopoly profit of 100. Upon removing rejections, we see in fact that U offered relatively high quantities to the D_i ($x = 3.44$); rejections led to both low realized output and low profit.[22]

As indicated by the high standard deviations, particularly for the variables associated with the *SECFIX* and *SECRAN* treatments, there is a substantial amount of noise in the data. In the presence of this noise, examining distributions, rather than simple means, is revealing. We depict the distributions of the outcome variables in the histograms in Figures 3 through 6. We then formally test the conclusions from our visual analysis of the histograms. The results are reported in Table 3. Such formal tests are of added value for two reasons: first, the histograms do not provide statistical significance levels, and, second, the histograms implicitly treat each observation as independent even though repeated observations for the same subjects over periods 1–10 may be dependent. We conducted two sorts of tests for the difference in outcome variables across treatment types, one sort involving means and one sort involving modes.

To test the significance of the difference in the mean of z (where z is one of the outcome variables q, x, Π, or Π_U) between treatment type A and B, we ran the regression

$$z_i = \beta_0 + \beta_1 TREATDUM_i + \epsilon_i, \tag{1}$$

---

[22] As will be seen in the histogram for q (Figure 4), both D_i rejected their contracts, resulting in a final output of zero, in about 10% of *SECRAN* observations.

FIGURE 4

HISTOGRAMS FOR MARKET OUTPUT

Key: ☐ All periods   ▦ Periods 6–10
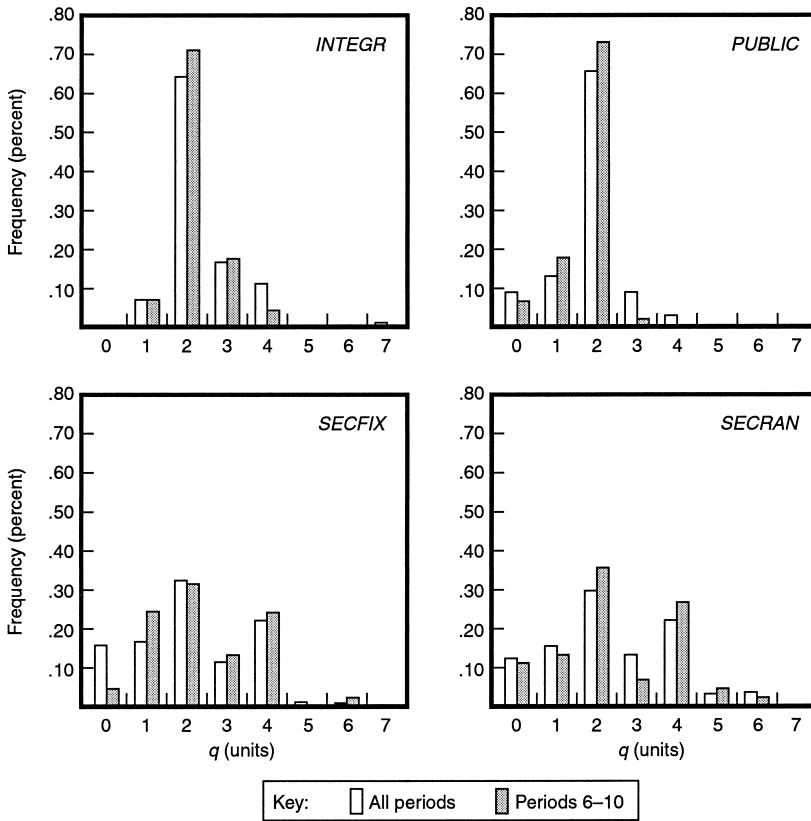
where i indexes observations from treatment type A or B only, *TREATDUM* is a dummy equalling one if the observation is from treatment type A and zero otherwise, β_0 and β_1 are coefficients to be estimated, and ϵ_i is an error term. The estimate of interest is $\hat{\beta}_1$, which can be interpreted directly as the difference in means in which we are interested. The computed standard errors are robust to general forms of heteroskedasticity and account for the possible dependence of repeated observations for subjects in the role of the upstream firm (formula provided in equation (A2) in Appendix A).

In what can be considered a test of the difference in modes across treatment types, we tested whether the probability that z took on a specific value, say $z = \bar{z}$, in treatment type A differed significantly from the probability in treatment type B. We ran a probit of the following form:

$$\mathbf{1}_{\{z_i = \bar{z}\}} = \beta_0 + \beta_1 TREATDUM_i + \epsilon_i, \tag{2}$$

where i indexes observations from treatment types A and B only, and $\mathbf{1}_{\{z_i = \bar{z}\}}$ is an indicator equalling one if $z_i = \bar{z}$ and zero otherwise.[23] The difference in probabilities, Δ, can be computed from the estimates using the formula $\Delta = \Phi(\hat{\beta}_0 + \hat{\beta}_1) - \Phi(\hat{\beta}_0)$, where $\Phi(\cdot)$ is the standard normal distribution function. The standard error computed for Δ is robust to general forms of heteroskedasticity and accounts for the possible dependence of repeated observations for subjects in the role of the upstream firm (formula provided in equation (A3) in Appendix A).

---

[23] To economize on notation, we use the same names for coefficients (β_0 and β_1) and error terms (ϵ_i) for regressions throughout the article, though of course there is no restriction that the values be the same across different regressions.

FIGURE 5

HISTOGRAMS FOR INDUSTRY PROFIT
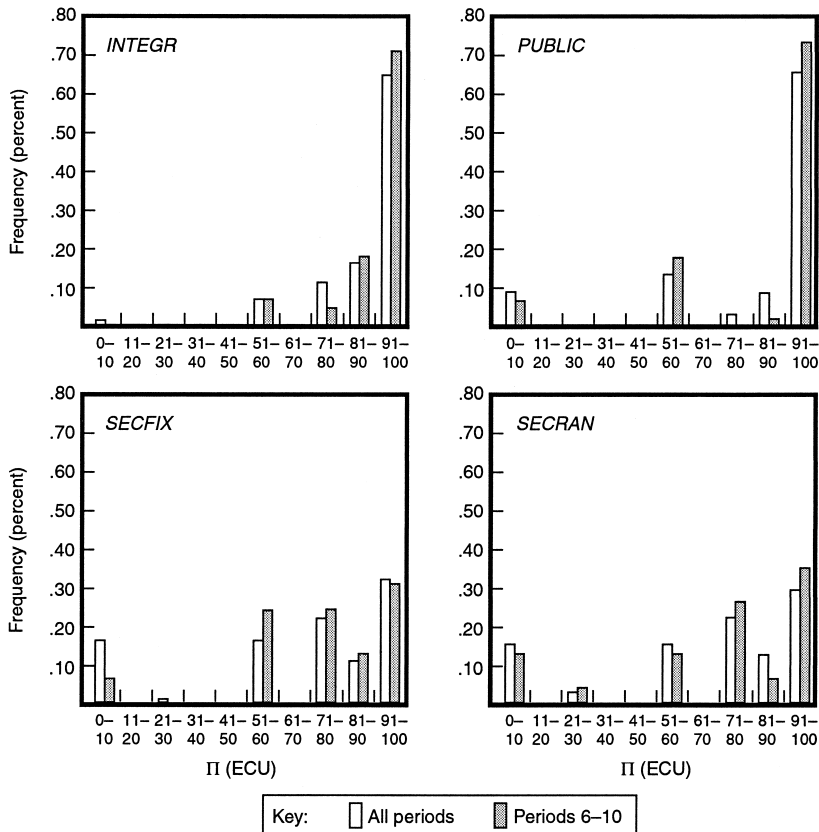
Consider the histograms for x and q in Figures 3 and 4. Throughout, the discussion will concentrate on the shaded bars, which omit data for periods 1–5, but it is evident from a comparison of the shaded and unshaded bars that the conclusions are the same if all periods 1–10 are considered. In *INTEGR*, in about 80% of the cases, three or more units were offered, where U-D_1's final sales (which implicitly is the quantity subdivision U_1 offers subdivision D_1) is added to the quantity U-D_1 offers D_2 to compute industrywide offered quantity. As mentioned in footnote 21, a finding of a greater quantity offered than the monopoly in *INTEGR* is not inconsistent with the theory: U-D_1 can still recover the monopoly profit as long as T_2 is high enough to induce D_2 to reject. In 91% of the observations in *INTEGR*, U-D_1 did offer D_2 a contract with a positive quantity, but in 80% of those, the offer included such a high tariff that we would expect it to be rejected by a sequentially rational D_2.[24] The prevalence of such contracts tends to inflate the value of x in *INTEGR*. In fact, 78% of such contracts were rejected by D_2, so that sales of the final good are in line with the theory: $q = 2$ in 71% of the *INTEGR* observations. In *PUBLIC*, two units were offered in 80% of the cases and two units sold in 73% of them (the discrepancy is due to a small rejection rate). For *SECRAN*, quantity offered has a bimodal distribution, with $x = 4$ the most

---

[24] For contracts with $x_2 = 1$, if D_2 accepts the contract and U-D_1 plays its best response on the continuation subgame, D_2 makes negative profit if $T_2 > 30$; for $x_2 = 2$, the associated condition is $T_2 > 36$. Half of the contracts were such that U-D_1 could guarantee itself at least the monopoly profit of 100—and would produce negative profit for D_2—if accepted (for $x_2 = 1$, the requisite condition is $T_2 \geq 40$; for $x_2 = 2$, the requisite condition is $T_2 \geq 64$). The other half would lead to lower-than-monopoly profits for U-D_1 in the event the contract were accepted, an event that theory predicts should only happen out of equilibrium.

FIGURE 6

HISTOGRAMS FOR UPSTREAM PROFIT

FIGURE 6

HISTOGRAMS FOR UPSTREAM PROFIT
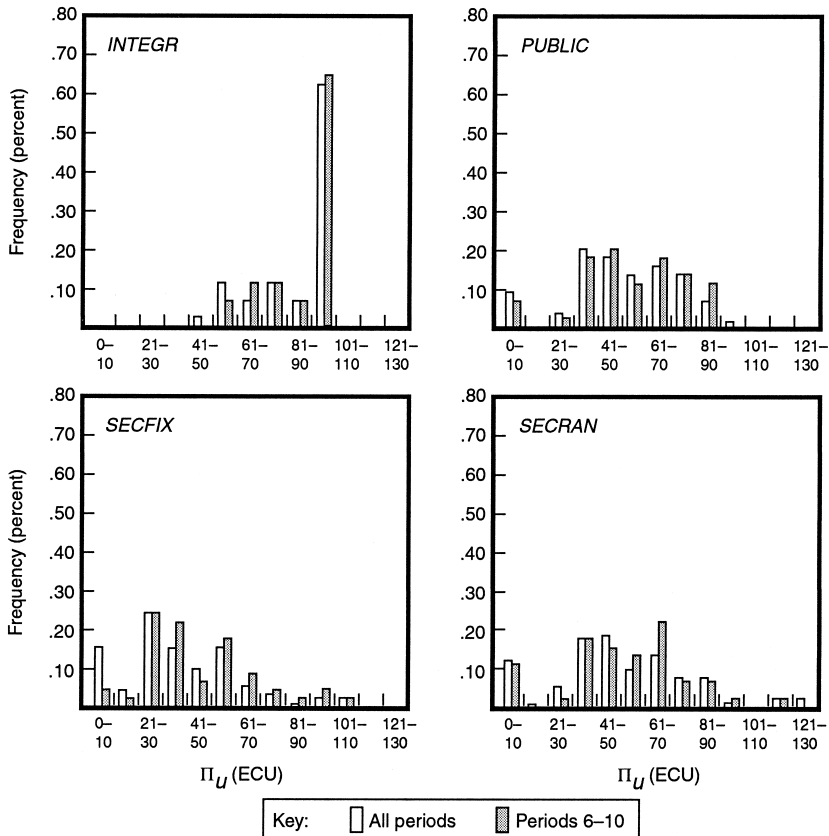
frequent case, as Hypothesis 3 predicts. A number of these offers were rejected; so though the distribution of q is still bimodal, $q = 4$ is less frequent than $q = 2$.

Turning to the formal results in Table 3, the means of x and q are lower for *INTEGR* and *PUBLIC* than for *SECRAN*, the difference between *PUBLIC* and *SECRAN* significant at the 5% level for both variables. The probability that $x = 2$ is 51% higher and that $x = 4$ is 36% lower for *PUBLIC* than for *SECRAN*, all differences significant at the 5% level. The probability that $q = 2$ is significantly higher for both *INTEGR* and *PUBLIC* than for *SECRAN* and the probability that $q = 4$ is significantly lower for both *INTEGR* and *PUBLIC* than *SECRAN*. There is no significant difference between *INTEGR* and *SECRAN* regarding the probabilities that $x = 2$ and $x = 4$—indeed, the probability that $x = 2$ is lower in *INTEGR* than in *SECRAN*—though, as explained above, this is not inconsistent with the theory and can be explained by contract offers from $U\text{-}D_1$ to D_2 in *INTEGR* that were nearly always rejected. The mean of q is significantly higher for *INTEGR* than for *PUBLIC*; the difference between *INTEGR* and *SECRAN*, though negative, is not significant. This can be explained by the fact that *PUBLIC* and *SECRAN* had a number of observations in which rejections led to $q = 0$ or 1; this happened much less frequently with *INTEGR*, since $U\text{-}D_1$ could always supply $q = 2$ itself if D_2 rejected its contract.

Given the mapping from q to Π, the histograms for industry profit, Figure 5, are similar to Figure 4: the bars have the same heights but simply have been rearranged along the axis. The histograms for *INTEGR* and *PUBLIC* have much of their mass at $\Pi = 100$ and little or none at $\Pi = 72$; the histogram for *SECRAN* has much less mass at $\Pi = 100$ and much more at $\Pi = 72$. Table 3 shows that the mean of Π for *INTEGR* is higher than for *PUBLIC*, but even more

TABLE 3   Tests of Differences Across Treatment Types

	Market Offered Input x			Market Output q			Market Profit Π	Market Profit Π_U
(1)	Mean of x (2)	Probability $x = 2$ (3)	Probability $x = 4$ (4)	Mean of q (5)	Probability $q = 2$ (6)	Probability $q = 4$ (7)	Mean of Π (8)	Mean of Π_U (9)
INTEGR vs. *PUBLIC*	.67**	−.60**	.07	.49**	−.02	.04	8.3**	33.8**
	(.26)	(.16)	(.13)	(.12)	(.12)	a	(3.9)	(6.3)
INTEGR vs. *SECFIX*	.07	−.27	−.13	−.18	.40**	−.20**	18.9**	42.5**
	(.35)	(.16)	(.18)	(.33)	(.15)	(.13)	(6.0)	(6.8)
INTEGR vs. *SECRAN*	−.47	−.09	−.29	−.27	.36**	−.22**	24.4**	37.4**
	(.37)	(.16)	(.17)	(.37)	(.14)	(.09)	(5.0)	(5.8)
PUBLIC vs. *SECFIX*	−.60	.33*	−.20	−.67*	.42**	−.24**	10.6*	8.7
	(.37)	(.18)	(.17)	(.32)	(.10)	a	(5.9)	(6.8)
PUBLIC vs. *SECRAN*	−1.13**	.51**	−.36**	−.76*	.38**	−.27**	16.1**	3.7
	(.39)	(.18)	(.15)	(.36)	(.08)	a	(4.9)	(5.8)
SECFIX vs. *SECRAN*	−.53	.18	−.16	−.09	−.04	−.02	5.5	−5.0
	(.46)	(.18)	(.20)	(.47)	(.12)	(.15)	(6.6)	(6.4)

Notes: An entry in column (2), (5), (8), or (9) is the difference between the means of the indicated treatments, computed as discussed in the text following equation (1). An entry in column (3), (4), (6), or (7) is the difference in the probability of the indicated event between the two treatments, computed as discussed in the text following equation (2). To purge learning effects, only last five periods are used to compute all estimates. White (1980) robust standard errors, adjusted for nonindependence of observations for the same upstream firm, are reported in parentheses.
* Significant at the 10% level; ** significant at the 5% level.
a *PUBLIC* has no observations in which $q = 4$, ruling out a probit; the reported entry is the difference in probabilities observed in sample, and the reported significance level is a conservative bound found by adding a fictitious observation for *PUBLIC* in which $q = 4$.

substantial is the difference in Π between these treatments and *PUBLIC*, 24.4 ECU for *INTEGR* and is 16.1 ECU for *PUBLIC*, both differences significant at the 5% level.

The main inconsistency with the theory revealed so far is that there are a substantial number of observations for *SECRAN* in which $x = 2$ and/or $q = 2$. These observations are inconsistent with hypothesis Hypothesis 3. Factors not accounted for by the theory seem to allow U to exercise greater commitment power than expected in a number of the *SECRAN* observations.

The histograms for upstream profit in Figure 6 present another interesting departure from the theory. The graph for *INTEGR* is close to the theoretical prediction: in a substantial majority of cases, 62%, the upstream firm obtained the monopoly profit 100. In *SECRAN*, upstream profits are distributed evenly along the horizontal axis. Much of the mass is below $\Pi_U = 70$. The remarkable fact is that the picture for *PUBLIC* is closer to *SECRAN* than to *INTEGR*. Focusing on the data for periods 6–10, there are no cases in which $\Pi_U = 100$; the mass is spread fairly evenly along the horizontal axis below $\Pi_U = 90$. The observations from the histograms are borne out in Table 3. The differences between the mean of Π_U for *INTEGR* and the other treatments are huge: mean Π_U was 33.8 ECU higher than for *PUBLIC* and 37.4 ECU higher than for *SECRAN*, both

differences significant at the 5% level. In contrast to the hypotheses, the mean of Π_U is essentially the same in *PUBLIC* as in *SECRAN*.

To summarize, the results provide partial support for the foreclosure theory. *SECRAN*, a treatment in which theory anticipates there should be a commitment problem, differs in the predictable way from *INTEGR* and *PUBLIC*, treatments in which theory anticipates no commitment problem. Though the commitment effect is apparent in the data, not all the observations follow the predictions of theory, so that *SECRAN*, while different from *INTEGR* and *PUBLIC*, is not as different as the theory would have. The experimental evidence departs from the theoretical predictions regarding the division of profits between upstream and downstream. Theory predicts that the upstream firm should have all the bargaining power, since it makes take-it-or-leave-it offers to the downstream firms. In *INTEGR*, consistent with the theory, the integrated firm is able to extract all of the industry profit a majority of the time. In *PUBLIC* and *SECRAN*, however, U only obtains a fraction of industry profits. In the nonintegrated treatments, the threat of D_i's rejecting U's offer seems to have disciplined U's bargaining power. This bargaining effect of vertical integration has real distributional consequences. The unintegrated downstream firm's surplus is effectively extracted by the merging firms.[25] Thus, the bargaining effect provides a rationale for vertical integration despite the existence of other vertical restraints (such as public contracts) that can solve the commitment problem as well: other vertical restraints may not allow the upstream firm to extract as much industry rent as full vertical integration. We will discuss the bargaining effect in further detail in Section 6.

□   **Repeated interaction.** As suggested by Macauley (1963) and later authors, reputation and repeated interaction can serve as a substitute for contracts. A comparison of *SECFIX*, in which upstream firms faced the same downstream firms for each of the ten periods, and *SECRAN*, in which the subjects were scrambled each period, will allow us to test whether repeated interaction is an adequate substitute for vertical restraints in solving the commitment problem. If so, the outcomes in *SECFIX* should depart from *SECRAN* in the direction of *INTEGR* and *PUBLIC*. An examination of Figures 3–6 and Table 3 shows only slight differences between *SECFIX* and *SECRAN*. The largest differences involve x: the mean of x is .53 units lower, the probability $x = 2$ is 18% higher, and the probability $x = 4$ is 16% lower with *SECFIX* than with *SECRAN*. These differences are not statistically significant, however. As the last row of Table 3 shows, there are no other discernible differences between *SECFIX* and *SECRAN*. There is little evidence that finite repetition of the game allows for the development of reputation.[26]

□   **Learning effects.** We have thus far focused on observations for periods 6-10 from each session, anticipating that subjects' strategies may converge during the course of the session to some stable point as they become more familiar with the game. In this subsection we present evidence of such learning effects based on two sets of formal tests, parallelling the tests discussed above in the subsection containing the central results. To test the significance of the difference in the mean of z (where z is one of q, x, Π, or Π_U) for periods 6–10 and the mean of z for periods 1–5 for a single treatment type A, we ran the regression

$$z_i = \beta_0 + \beta_1 PERIODUM_i + \epsilon_i, \tag{3}$$

where i indexes observations from treatment type A only and *PERIODUM* is a dummy equalling one if the observation is from periods 6–10 and zero otherwise. The estimate of interest is $\hat{\beta}_1$,

---

[25] How this additional surplus is split between the merging subsidiaries U and D_1 is a question that our experimental design does not address because we have a single party playing both roles in U-D_1. If one theorizes that U auctions the right to merge to D_1 and D_2, equilibrium will depend on how tough the bidding competition is between them. If, for example, they engage in Bertrand competition, U can extract 100% of the profit, implying that both downstream units, the integrated subsidiary and the unintegrated firm, would earn zero net surplus.

[26] It might be argued that reputation effects would arise more readily in sessions with random rather than certain ending points; however, Selten and Stoecker (1986) show that the distinction has little impact on the behavior of experimental subjects.

which can be interpreted directly as the difference in means in which we are interested. As before, the reported standard errors are robust to heteroskedasticity and nonindependence of repeated observations.

In the second set of tests, we sought to determine whether the probability that z took on a specific value, say $z = \bar{z}$, in periods 6–10 differed significantly from the probability in periods 1–5. We ran a probit of the following form:

$$1_{\{z_i = \bar{z}\}} = \beta_0 + \beta_1 PERIODUM_i + \epsilon_i, \qquad (4)$$

where i indexes observations from treatment type A only. The difference in probabilities Δ can be computed as described below equation (2). As before, the reported standard errors are robust to heteroskedasticity and nonindependence of repeated observations.

The results are presented in Table 4. There is some evidence of convergence to the theoretical equilibrium in *INTEGR* and *SECFIX*, but little evidence of movement in *PUBLIC* and *SECRAN*. In *INTEGR*, the mean of x fell by .42 units, the probability $x = 2$ rose by 11%, and the probability $x = 4$ fell by 18%. U-D_1 essentially learned that it was more profitable to offer D_2 one or no units rather than two. This translated into .36 fewer units sold on average, an increase in the probability $q = 2$ of 13%, and an equal fall in the probability $q = 4$. This also translated into an increase in mean Π of 5.7 ECU.

In *SECFIX*, mean q rose .44 units. The rise was due to a reduction in the number of cases in which both D_i rejected U's offer, resulting in no sales (see Figure 4). The reduction in the number of "double rejections" in turn translated into an increase in the mean Π of 13.8 ECU and an increase in mean Π_U of 16.6. These changes reflect movement toward the predictions of

TABLE 4   **Tests for Learning Effects**

	Market Offered Input x			Market Output q			Market Profit Π	Market Profit Π_U
	Mean of x	Probability $x = 2$	Probability $x = 4$	Mean of q	Probability $q = 2$	Probability $q = 4$	Mean of Π	Mean of Π_U
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
INTEGR	−.42*	.11*	−.18*	−.36	.13*	−.13**	5.7**	3.6
	(.21)	(.07)	(.08)	(.23)	(.07)	(.07)	(1.7)	(2.1)
PUBLIC	−.18	.16	−.02	−.27	.16**	.07	4.1	6.0
	(.11)	(.09)	(.04)	(.16)	(.06)	a	(3.7)	(3.4)
SECFIX	−.20	.07	.04	.44*	−.02	.04	13.8*	16.6**
	(.28)	(.12)	(.04)	(.23)	(.08)	(.07)	(6.5)	(3.1)
SECRAN	−.40	.16	.13	.11	.11	.09	3.4	3.5
	(.54)	(.17)	(.15)	(.47)	(.09)	(.12)	(5.3)	(5.3)

Notes: An entry in column (2), (5), (8), or (9) is the difference between the mean for periods 6–10 for the treatment in (1) and the mean for periods 1–5, computed as described in the text following equation (3). An entry in column (3), (4), (6), or (7) is the difference in the probability of the indicated event for periods 6–10 for the treatment in (1) and the probability for periods 1–5, computed as discussed in the text following equation (4). White (1980) robust standard errors, adjusted for nonindependence of observations for the same upstream firm, are reported in parentheses.

* Significant at the 10% level; ** significant at the 5% level.
a *PUBLIC* has no observations in which $q = 4$ in periods 6–10, ruling out a probit; the reported entry is the difference in probabilities observed in the sample, and the reported significance level is a conservative bound foundby adding a fictitious observation for period 6–10 for *PUBLIC* in which $q = 4$.

© RAND 2001.

theory: from periods 1–5 to periods 6–10, the mean of q rose from 1.93 to 2.38 (slightly closer to the theoretical prediction $q = 4$) and the mean of Π rose from 61.6 to 75.4 (closer to the theoretical prediction $\Pi = 72$).

5. Out-of-equilibrium beliefs

■ **Evidence from the experiments.** As indicated in the theoretical discussion of Section 2, the nature of downstream firms' out-of-equilibrium beliefs is crucial for the foreclosure theory. If the D_i hold symmetric beliefs, the theory implies that U is able to extract the monopoly profit with or without vertical restraints; if the D_i hold passive beliefs, U cannot extract the monopoly profit without vertical restraints. Indirect evidence on the nature of subjects' beliefs can be obtained by examining the outcomes in *SECRAN*. If beliefs are passive, the outcomes in *SECRAN* should differ from *INTEGR* or *PUBLIC*. In fact, we found differences among these treatments, but the differences were not as large as the theory with passive beliefs would suggest. In any event, this is only indirect evidence on the nature of downstream firms' beliefs, since outcomes depend on the actions of three parties rather than the beliefs of a single downstream subject.

We can construct a more direct test of the nature of beliefs by focusing on the acceptance/rejection decision a_i of a single downstream firm D_i in *SECRAN* treatments. The tests are based on the theoretical result that, in the *SECRAN* setup, there exist offers (x_i, T_i) for which a_i differs depending on whether downstream firms hold passive or symmetric beliefs. Mixed strategies are important in the subsequent discussion, so it is useful to define α_i to be the probability that D_i accepts U's offer; i.e., $\alpha_i = \Pr(a_i = 1)$. For both passive and symmetric beliefs, Figure 7 graphs the value of α_i that constitutes a PBE of the continuation game following offer (x_i, T_i). Verification of the results in Figure 7 are provided in Appendix A. Several remarks about the graphs are in order. First, (x_i, T_i) may well be out-of-equilibrium offers. Indeed, the equilibrium offer with symmetric beliefs—(1, 50)—is not an equilibrium offer with passive beliefs—(2, 36); all other offers are out of equilibrium for both belief structures. Second, in the case of symmetric beliefs, there is a large range of offers for which $\alpha_i \in (0, 1)$, implying that equilibrium in the continuation game involves mixed strategies. Third (and most important for our purposes in this section), there are distinct differences between the passive- and symmetric-beliefs cases over a broad range of offers (x_i, T_i); this will form the basis of our tests.

According to Figure 7, the continuation equilibria for passive and symmetric beliefs involve different pure strategies only following offers of the form $(1, T_i)$, $T_i \in (30, 50)$ (see the shaded portion of Figure 7). We will focus on offers in this class, for brevity labelled "dispositive contracts." Theory predicts that D_i should reject a dispositive contract with certainty if beliefs are passive but accept it with certainty if beliefs are symmetric. The reason D_i rejects a dispositive contract in the passive-beliefs case is that it continues to believe D_j receives—and accepts—the equilibrium contract (2, 36). If D_i were to accept, it would earn gross profit 30, which is negative net of the $T_i > 30$ transfer. The reason D_i accepts a dispositive contract in the symmetric-beliefs case is that it believes D_j receives the identical out-of-equilibrium offer. By accepting, D_i earns at worst gross profit of 50, which is still positive net of the $T_i < 50$ transfer.

Restricting attention to periods 6–10 of the *SECRAN* treatments, there were 27 dispositive contract offers. Two-thirds were accepted. That this proportion is substantially higher than zero is inconsistent with all D_i holding passive beliefs. That this proportion is substantially less than one is inconsistent with all D_i holding symmetric beliefs. It is possible that the nature of beliefs is heterogeneous among the D_i, with some holding passive and some symmetric beliefs.

A caveat regarding this conclusion is in order. The third that rejected a dispositive contract may have held symmetric beliefs but rejected the contract anyway because their expected share of surplus, though positive, was not as high as they considered "fair." (This is a facet of the bargaining issue treated in more detail in Section 6.) Alternatively, they may have made mistakes in reasoning. To investigate these possibilities, we examined *PUBLIC* observations in which both downstream firms received dispositive contracts. These observations serve as a control for our purposes, since the issue of out-of-equilibrium beliefs about rivals' contracts is rendered irrelevant

FIGURE 7

CONTINUATION EQUILIBRIUM ACCEPTANCE PROBABILITIES FOR DIFFERENT BELIEFS

Key: ——— Passive beliefs   - - - Symmetric beliefs
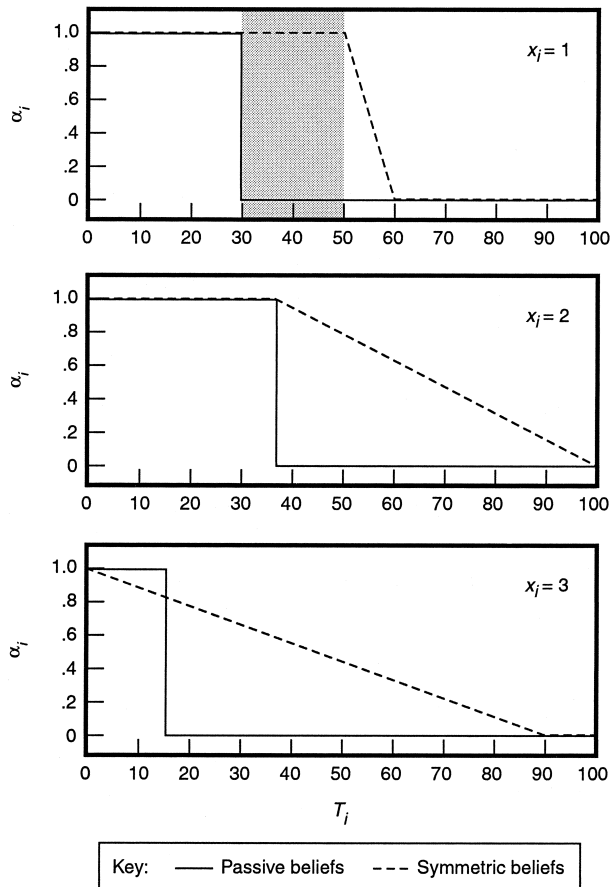
by the publicness of the contracts. With these observations, any rejection of a contract offer by a downstream firm must be for reasons other than out-of-equilibrium beliefs. D_i accepted 77% of such offers, 10% more than for $SECRAN$, though this difference was not found to be statistically significant.[27]

A further question regards whether U's offers justified either passive or symmetric beliefs: i.e., were U's out-of-equilibrium offers symmetric or not? Of the 27 dispositive contract offers to D_i, only 37% of the simultaneous offers to D_j were identical; the rest either involved the same input quantity and a different tariff (44%) or a different input quantity (19%).

□ **Evidence from postexperimental questionnaires.** A postexperimental questionnaire, in which we asked participants to describe how they made decisions in the experiements, provides more direct evidence on beliefs.[28] Though we did not specifically ask them about out-of-equilibrium beliefs, a few of the participants in $SECRAN$ treatments responded with some useful information on the issue, translated from the original German below.

Two subjects playing the role of a downstream firm made statements consistent with symmetric beliefs. One wrote, "In general (and this was confirmed over time by the results) it turned

---

[27] The statistical test was based on a probit similar to (2) and (4).

[28] We also asked the students questions for statistical purposes on major field and year of study and gender. Questionnaires are generally regarded as a useful method of gathering additional information in experiments (Friedman and Sunder, 1994). A weakness of this method is that there are no monetary incentives for answering the questions.

out that the other retailer received a roughly similar offer." There were several cases in periods 6–10 in which this subject received a dispositive contract; the subject accepted them consistent with symmetric beliefs. Another wrote, "I held the basic assumption that both retailers receive the same offer." This subject did not receive dispositive contract offers in periods 6–10, so we cannot determine whether or not his actions were consistent with his questionnaire statement. There was also a subject in the role of an upstream firm who suggested that downstream firms hold symmetric beliefs: "I was assuming that the retailers think that I offered the same quantity to both of them." Nine out of ten of this subject's contract offers in periods 6–10 were of the dispositive form, which are accepted with certainty in equilibrium with symmetric beliefs and rejected with certainty in equilibrium with passive beliefs.

Two subjects made statements consistent with passive beliefs. One wrote: "I calculated whether I would make a positive profit if the other retailer purchased 1 or 2 units." In periods 6–10, this subject rejected offers $(1, 35)$ and $(1, 40)$, consistent with passive beliefs, though in another instance he did accept contract $(1, 35)$. Another wrote: "Accept if, even when the other retailer buys 2 units, a profit remains." This subject rejected the dispositive contracts offered in periods 6–10, consistent with passive beliefs.

In sum, the direct evidence on beliefs suggests that there may have been a substantial amount of heterogeneity in beliefs among participants. D_i's accept/reject decision was consistent with passive beliefs in some cases and symmetric beliefs in others. The actual offers made by U were sometimes symmetric and sometimes not. Statements about beliefs in postexperimental questionnaires were sometimes consistent with symmetric and sometimes symmetric beliefs. These findings motivate our analysis in the next subsection, in which we extend the vertical foreclosure theory to take account of heterogeneity of beliefs. The modified theory is successful in explaining a number of instances of seemingly anomalous behavior by the experimental subjects.

☐   **Extending the foreclosure theory.** Rather than having all downstream firms hold passive beliefs or all hold symmetric beliefs, suppose the downstream firms are drawn from a population in which a fraction $s \in [0, 1]$ hold symmetric beliefs ("symmetric types") and $1 - s$ hold passive beliefs ("passive types"). The distribution of beliefs in the population is common knowledge, but D_1 and D_2's realized beliefs are private information.

The belief structure assumed in this section generalizes previous theoretical work; however, it is still quite restrictive compared to the set of possible belief structures. In general, beliefs may be neither passive nor symmetric.[29] In general, beliefs may be contingent on the out-of-equilibrium offer received. An analysis of equilibrium under general belief structures is beyond the scope of this article, though the initial steps taken in this section may at least serve to suggest the possible benefits from such a general analysis.

In all other respects besides the heterogeneity of beliefs, the model is as described in Section 2. In particular, inverse demand is $P(q_1+q_2)$, production is costless at all stages, U is unintegrated, and U makes secret offers to the D_i. It turns out that for some parameter values there are multiple PBE. In the following proposition, proved in Appendix A, we will focus on extremal PBE, i.e., PBE giving U its highest profit. Refer to Section 2 for definitions of the functions R and b used throughout the subsection below.

Proposition 1. In an extremal PBE, U offers contract (x_i^*, T_i^*) to D_i, $i = 1, 2$; and D_i, $i = 1, 2$ accepts with probability α_i^*, where $T_i^* = \alpha_i^* R(x_i^*, x_i^*) + (1 - \alpha_i^*)R(x_i^*, 0)$ and where x_i^* and α_i^* solve

$$\max_{x_i \in [0,\infty), \alpha_i \in [0,1]} \{\alpha_i[\alpha_i R(x_i, x_i) + (1 - \alpha_i)R(x_i, 0)]\}, \tag{5}$$

---

[29] For example, McAfee and Schwartz (1994) explore a third alternative, so-called wary beliefs. If D_i has wary beliefs and receives an out-of-equilibrium offer, it believes that D_j received the best response to D_i's contract. In our setting with secret offers, McAfee and Schwartz (1994) show that wary beliefs produce the same PBE as passive beliefs; so our restriction to passive and symmetric beliefs may be less restrictive than it appears.

subject to

$$\alpha_i[\alpha_i R(x_i, x_i) + (1 - \alpha_i)R(x_i, 0)] \geq \max_{x_i^d \in [0,\infty)} \{(1 - s)[\alpha_i R(x_i^d, x_i) + (1 - \alpha_i)R(x_i^d, 0)]\}. \quad (6)$$

For the fraction α_i^* of downstream firms that accept, T_i^* is set at just the right level to extract all of their expected surplus from accepting. In general, this tariff may be higher than $R(x_i^*, x_i^*)$—the expression we are used to from the model with 100% passive or 100% symmetric types—since the rival rejects with probability $1 - \alpha_i^*$, allowing the sole active downstream firm to earn higher net revenue $R(x_i^*, 0)$. The constraint in Proposition 1 ensures that U does not want to deviate by offering a contract with a higher quantity targeted at passive types. In fact, much of the proof of Proposition 1 is directed to showing that this is the only relevant deviation from a large set of possible deviations. It is particularly lucrative to target passive types with deviation contracts, since passive types do not revise up their conjectures about a rival's quantity when their own is higher than anticipated.

As a corollary of Proposition 1, we prove the following proposition in Appendix A:

Proposition 2. Define $\hat{s} \equiv 1 - R(q^m/2, q^m/2)/R(b(q^m/2), q^m/2)$. For $s \in [\hat{s}, 1]$, the solution to (5) involves $\alpha_i^* = 1$, $x_i^* = q^m/2$, and $T_i^* = \Pi^m/2$, i.e., the same outcome as in the PBE with 100% symmetric types. For $s = 0$, the solution to (5) involves $\alpha_i^* = 1$, $x_i^* = q^c$, and $T_i^* = \Pi^c$; i.e., the same outcome as in the PBE with 100% passive types.

Proposition 2 verifies that our extended model properly nests the results for the cases of 100% symmetric and 100% passive types. It also shows that it is not necessary to have 100% symmetric types to generate monopoly profit for U. It is sufficient to have $s \in [\hat{s}, 1]$. This is indeed a large range of values, since it can be proved that $\hat{s}$ is always less than $1/2$.[30]

It is possible to use Propositions 1 and 2 to compute the extremal PBE for all $s \in [0, 1]$ in our experimental setting. The results are graphed in Figure 8, with the associated calculations provided in Appendix A. For $s \geq \hat{s} = 1/6$, the extremal PBE—involving contracts $(1, 50)$ that are accepted with probability one—is indistinguishable from the outcome in the standard model with 100% symmetric types. For $s = 0$, the equilibrium outcome with 100% passive types is obtained, involving contracts $(2, 36)$ that are accepted with probability one. For $s \in (0, 1/6)$, we have a departure from the standard model with 100% symmetric or 100% passive types. The equilibrium involves $x_i = 2$, as does equilibrium with 100% passive types; unlike the case of 100% passive types, it involves a higher tariff and a positive probability of contract rejection. Though U is not directly able to restrict output to the monopoly level by restricting the inputs it offers, it is able to do so indirectly by offering a high-enough tariff that a fraction of the downstream firms reject the contract. For $s < \hat{s}$, T_i^*, $1 - \alpha_i^*$ (the probability of rejection), and Π_U^* (U's expected profit) are all increasing in s.

Our model of heterogeneous beliefs can thus be used to rationalize behavior observed in our experiments that cannot be rationalized in a model with 100% symmetric or 100% passive types. Of the 180 contract offers in periods 6–10 in *SECFIX* and *SECRAN*, 38 were of the form $(2, T_i)$, $T_i > 36$. Our theory implies that the acceptance rate should be at least 83%; the observed acceptance rate for these 38 offers (72%) was in this range, but not quite as high.

6. Bargaining issues

■ In this section we abstract from questions concerning beliefs that were the focus of the previous section and instead focus on the bargaining aspects of our game. In the following two subsections, we examine the *INTEGR* and *PUBLIC* treatments, treatments in which secret contracts—and hence out-of-equilibrium beliefs—are not an issue. As seen in Table 2, average downstream profit Π_D was positive for both *INTEGR* and *PUBLIC*; it can be shown that these means are significantly greater than zero at the 5% level. The theory, on the other hand, predicts $\Pi_D = 0$ for all treatment types. This contradiction is not surprising: such differences between

---

[30] This follows by noting $2R(q^m/2, q^m/2) = R(q^m, 0) > R(b(q^m/2), q^m/2)$ and rearranging.

FIGURE 8

EXTREMAL PBE FROM MODEL WITH HETEROGENEOUS BELIEFS IN
EXPERIMENTAL SETTING

Probability of symmetric type s
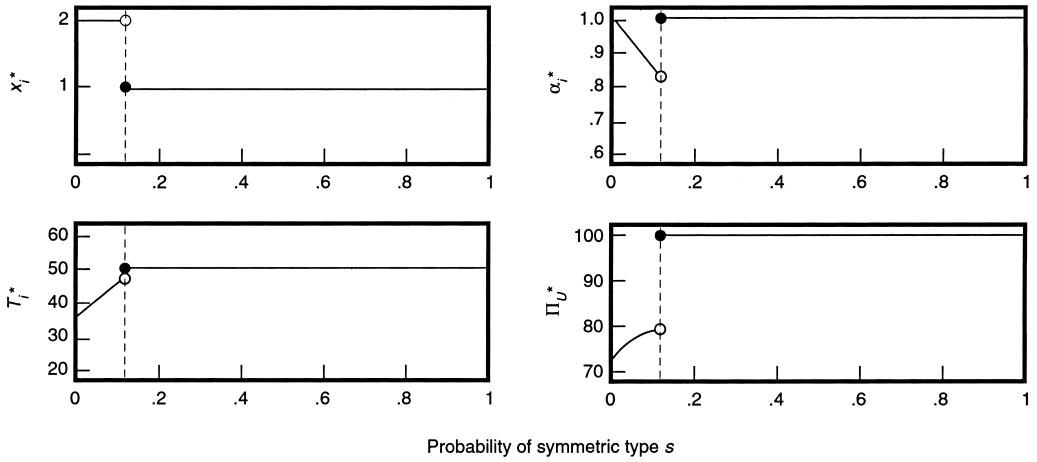
predictions of noncooperative game theory and experimental data are a robust finding in the experimental research on bargaining games (see Roth (1995) for a survey). The surprising feature of our results is that they are much closer to the theoretical prediction than those of previous studies.

The section continues with a demonstration that a concern for fairness on the part of downstream firms (meaning that they may reject offers that give them a positive surplus if their share of the surplus is not sufficiently large), which usually harms the party making take-it-or-leave-it contract offers, actually can increase the upstream firm's profit in our setting. The section concludes with a discussion of the Shapley value as a predictor of parties' surplus shares.

☐ **Bargaining and the *PUBLIC* treatment.** Our *PUBLIC* treatment type is similar to ultimatum bargaining experiments (Güth, Schmittberger, and Schwarze, 1982). In the ultimatum game, the first player (the proposer) offers a division of a pie of fixed size to the second player (the responder). The responder can accept the division, in which case it is implemented, or reject it, in which case both earn nothing. In the subgame-perfect equilibrium, the proposal gives all of the surplus to the proposer and is accepted by the responder. Typical experimental results for the ultimatum game involve a more equitable distribution of the surplus than the theory predicts. For example, in one treatment in Forsythe et al. (1994), specifically the "April–Pay" treatment in which a pie of five dollars was to be split, the modal proposal gave the responder a 50% share; the mean proposal gave the responder a 44% share. Results from other ultimatum experiments are similar (see Roth, 1995).

A number of our *PUBLIC* observations have a similar structure to the ultimatum game. For the last five periods, there were 13 *PUBLIC* observations in which $x_i = 2$ and $x_j = 0$. In such observations, 100 ECU were available to split between U and D_i, and this split could be accomplished by setting T_i at various levels. In four cases, $T_i = 70$, and this was never rejected. Eight times, $T_i = 80$, and this was accepted every time but two. In the remaining case, $T_i = 90$, and this was accepted. The mean of T_i was 77.7; the mean of Π_U was 65.4, and the mean of Π_D was 19.2. The mean offer was about 22% closer to the theoretical prediction of zero surplus for the responder than in the experiments of Forsythe et al. (1994) reported above.

Similarly, we can analyze the 23 *PUBLIC* observations in which $x_i = x_j = 1$.[31] In such observations, assuming D_j accepts its offer, 50 ECU is the size of the pie to be split between U

---

[31] Since there are two proposals in each of the 23 observations, our analysis is based on 46 proposals.

and D_i, with T_i set to divide the surplus appropriately. The mean of T_i was 35.6 and the mode was 45. The overall acceptance rate was 83%, and this acceptance rate was only slightly negatively correlated with T_i, so that the modal offer of 45 was the most profitable even after accounting for rejections. The mean offer was about 16% closer to the theoretical prediction of zero surplus for the responder than in the ultimatum experiments of Forsythe et al. (1994) reported above.

For the 23 *PUBLIC* observations in which $x_i = x_j = 1$, we may have found results closer to the theoretical prediction than the previous experimental literature due to the linkage between downstream firms. Since D_i and D_j are competitors, the pie to be split between U and D_i depends on D_j's action. In equilibrium, D_j accepts its offer and the pie to be split between U and D_i is 50. If D_j rejects, the size of the pie to be split between U and D_i increases from 50 to 60. The possibility of rejection by D_j may increase D_i's incentive to accept U's offer for a given T_i, and may in turn induce U's to increase T_i. This explanation is incomplete, however. In the 13 *PUBLIC* observations in which $x_i = 2$ and $x_j = 0$, the pie to be split between U and D_i was independent of D_j's action, yet we still found results closer to the theoretical prediction than the previous experimental literature.[32]

There are several other possible explanations. One possibility is that our experiments involved a buyer-seller exchange. Hoffman et al. (1994) found that ultimatum bargaining experiments framed as buyer-seller exchanges were closer to the theoretical prediction than standard ultimatum experiments, though not as close as our results. Another possibility is that the computations involved in making proposals was more complicated in our setup than in the typical ultimatum game and that subjects allowed U an additional rent for this. Hoffman et al. (1994) observed similar effects in experiments in which the role of the proposer had to be "earned" in a contest. A third possibility is that initial endowments in our setup were not public information (U was not told the D_i's initial endowment, nor were the D_i told U's). Franciosi et al. (1995) found that, in posted offer markets, treatments with private knowledge about payoffs led to outcomes that were closer to the theoretical prediction than treatments with public knowledge. Given that there was public information about all aspects of payoffs except for initial endowments, unless U thought that there was some extraordinarily high initial capital for the D_i, U had good reason to believe it would receive a much higher total payment than the D_i. Therefore, it is unlikely that initial endowments being private knowledge had a significant impact on the results, though this possibility cannot be excluded.

□ **Bargaining and the *INTEGR* treatment.** Our *INTEGR* treatment bears more similarity to the dictator game than the ultimatum game. In the dictator game, the proposer decides how to split a pie, but the responders have no veto power. The proposer's decision is final. Similarly, in *INTEGR*, the integrated firm U-D_1 can take actions to guarantee itself the monopoly profit regardless of D_2's actions, for example, by refusing to offer D_2 a contract or offering D_2 a contract in which T_2 is sufficiently high. In the subgame-perfect equilibrium of the dictator game, the proposer obtains all the surplus; but as with experiments with the ultimatum game, the experimental results diverge from theory. For example, in the "April–Pay" treatment with five-dollar pies in Forsythe et al. (1994), only 36% of the proposers took the entire pie and still 21% gave away the equal share. The mean share of the respondent's share in the proposals was 18%.

In 73% of the cases in *INTEGR*, by contrast, U-D_1 either refused to offer D_2 a contract or offered D_2 a contract in which T_2 was so high that the subgame-perfect continuation equilibrium involved D_2's rejecting the offer.[33] Averaging over all observations, D_2's profit was 4% of industry profit; even averaging over observations in which D_2 accepted the offer, its profit was only 14%

---

[32] Güth, Marchand, and Rulliere (1997) studied the effects of competition in bargaining using an ultimatum bargaining setup with one proposer and five responders. In one treatment, the responder with the lowest acceptance level became the decisive responder; in another, the decisive responder was chosen randomly. Both treatments produced results closer to the subgame-perfect equilibrium than standard ultimatum games, (the first treatment significantly closer than the second).

[33] See footnote 24 for a discussion relating to these observations.

of industry profit. The highest profit D_2 earned in any observation was 30, and the highest percentage of industry profit it earned was 30%.

One possible reason that our results are closer to the theoretical prediction than previous experiments is that surplus is freely transferable in the dictator game but not in *INTEGR*. In particular, it is impossible for U-D_1 to commit to transfer a surplus in the interval $(0, 20)$ to D_2 credibly while maintaining industry profit at 100. One way to accomplish this is to offer D_2 the contract $(1, T_2)$, $T_2 \in (30, 50)$, and have U-D_1 sell one unit. Another way is to offer D_2 the contract $(2, T_2)$, $T_2 \in (80, 100)$, and have U-D_1 sell no units. In the continuation equilibrium, D_2 would reject either of these offers, since U-D_1's commitment to sell zero or one unit is not credible.[34] Faced with the choice of giving D_2 zero surplus on the one hand or 20 or more on the other, U-D_1 may choose zero even though it would have given D_2 some small positive surplus if surplus were freely transferable.

The other explanations from the previous subsection for why the *PUBLIC* treatment is closer to theory than previous experiments may apply to *INTEGR* as well: the experiment involves a buyer-seller transaction, U has to engage in possibly difficult computations for which it may expect a rent, and initial endowments are private information. As above, none of these explanations seems to be satisfactory on its own.

☐ **Bargaining and secret contracts.** In *SECFIX* and *SECRAN*, the share of surplus accruing to the downstream firms was roughly the same (37% and 24%, respectively) as that in *PUBLIC* (34%). Since equilibrium in *SECFIX* and *SECRAN* does not involve splitting a pie of maximal size—industry profit is 72 rather than 100 in equilibrium assuming passive beliefs—the analogy between these treatments and the ultimatum game is not particularly close. Instead of comparing our experimental results to those for the ultimatum game, we will present a counterintuitive theoretical result that the existence of a "crazy" downstream type—a type that out of concerns for fairness rejects offers giving it a positive surplus if its surplus share is not sufficiently high—may actually benefit U. The reason is that this crazy type reduces the expected level of output, in effect an imperfect substitute for a commitment device.

Let θ be the probability a downstream firm is crazy in the sense that it must receive 100% of the surplus in its transaction with U or it will reject the offer. This is an extreme assumption concerning the nature of this type that serves to simplify the subsequent calculations; the results would be qualitatively similar if the crazy type accepted offers with more moderate divisions of surplus. "Sane" downstream firms accept any offer giving them nonnegative expected surplus. In equilibrium with passive beliefs, U offers the D_i contracts $(2, 36 + 64\theta)$. Crazy types reject the contract. Sane types accept. Note that sane types earn expected profit $100\theta + 36(1 - \theta) - T_i$ if they accept because, with probability θ, rival D_j is a crazy type and rejects, in which case $q = 2$ and D_i's revenue is 100. Substituting $T_i = 36 + 64\theta$ verifies that sane types make nonnegative profit by accepting. Thus U's expected profit is $2(1 - \theta)(36 + 64\theta)$, graphed in Figure 9. For $\theta \in (0, .43)$, U's expected profit is higher in the equilibrium with the crazy type than it is without (in which case it earns 72 in equilibrium).

☐ **Shapley value.** Although noncooperative game theory provides an accurate prediction of total industry surplus in all treatments and an accurate prediction of the shares of surplus accruing to the various parties in *INTEGR*, as we have seen, it is not as good at predicting the shares of surplus accruing to the various parties in the nonintegrated treatments. The upstream firm earns significantly less than the 100% share predicted by noncooperative game theory. The cooperative game-theoretic concept of Shapley value turns out to be a better predictor of surplus shares.[35]

---

[34] A third way would be to have U-D_1 offer D_2 a contract of the form $(0, T_i)$ with $T_i < 0$. Our experimental design ruled out such contracts because T_i was constrained to be an integer in $[0, 120]$.

[35] Being a cooperative game-theoretic concept, Shapley value posits that industry profit is maximal (100 in our setup) in all treatments and so does a worse job than the noncooperative theory in predicting total industry profit in *SECRAN* and *SECFIX*.

FIGURE 9

UPSTREAM PROFIT WITH SECRET CONTRACTS AND CRAZY DOWNSTREAM TYPE
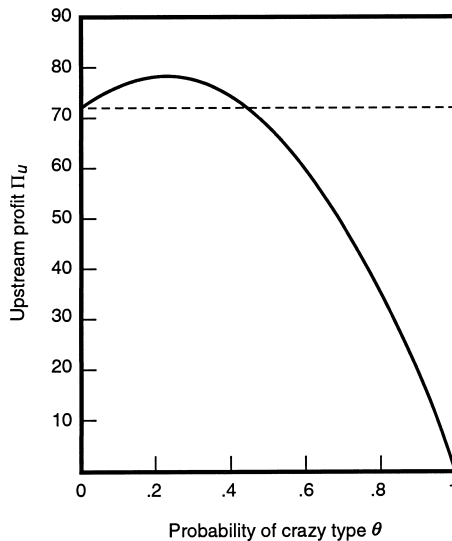
The Shapley value gives each player the average of its marginal contribution to a coalition of players including it and its predecessors in a given permutation of players assuming that each permutation is equally likely.[36] Appendix A provides the Shapley value calculations. For the nonintegrated treatments, the Shapley value involves shares 2/3 for U, 1/6 for D_1, and 1/6 for D_2. Intuitively, U makes a marginal contribution to all coalitions in which it is a member, since no sales can be made without its input. D_i only makes a marginal contribution to the coalition $\{U, D_i\}$; in coalitions with all three players, U can sell two units through D_j and generate industry profit of 100 without D_i. Our experimental results produced similar surplus shares, with U obtaining 66% in *PUBLIC*, 64% in *SECFIX*, and 76% in *SECRAN*.

The Shapley value also predicts surplus shares in *INTEGR*. In this treatment, the Shapley value involves a 100% share for U-D_1 and a 0% share for D_2. The prediction is close to the 96% share for U-D_1 observed in our experiments.

7. Conclusion

■ Much of the new foreclosure theory rests on the existence of a commitment effect: absent some vertical restraint, an upstream monopolist is not able to commit to restrict output when dealing with competing downstream firms. We observed the commitment effect in our experiments. In treatments with vertical restraints (vertical integration in *INTEGR*, public contracts in *PUBLIC*), the upstream subject was able to restrict output to the monopoly level significantly more frequently, around 40% more, than in treatments with no vertical restraints (*SECFIX* and *SECRAN*). The evidence suggests that out-of-equilibrium beliefs are not uniformly symmetric, as would be required to obtain commitment without vertical restraints.

Our results also suggest that the commitment effect alone is not sufficient to explain the existence of vertical integration. The upstream firm was able to obtain the same degree of output restriction with the simple vertical restraint of public contracts as with full vertical integration. The main advantage of integration over less severe vertical restraints to the upstream firm is that it confers a tremendous amount of bargaining power. That is, vertical restraints of various forms may succeed in maximizing the size of the industry profit pie; vertical integration allows the integrated firms to obtain a larger slice at the expense of the unintegrated firm. Therefore, the

---

[36] See Mas-Colell, Whinston, and Green (1995) for a discussion.

subset of the foreclosure literature that looks at bargaining effects (e.g., the "scarce needs" and "scarce supplies" variants in Hart and Tirole (1990), Bolton and Whinston (1993), and Segal (1999)) may be as important in explaining anticompetitive vertical integration as the larger literature on the commitment effect.

Our results suggested that the existing models in which it is common knowledge that players' out-of-equilibrium beliefs are homogeneous may not be sufficiently rich to capture the complicated strategic environment in practical settings. There were a number of instances in our experiments where downstream subjects received offers that were out of equilibrium in any standard model. The downstream subjects' responses were neither uniformly consistent with symmetric beliefs nor uniformly consistent with passive beliefs. This motivated us to construct a model with both passive- and symmetric-believing players in the population of downstream firms. We characterized the extremal perfect Bayesian equilibrium for arbitrary demand specifications. Consistent with the intuition from standard models, the upstream firm does better—in a sense has more commitment power—the greater the proportion of symmetric types. The comparative-static effects of an increase in this proportion on the equilibrium contract were quite surprising. For some parameters, the upstream firm does not have the commitment power to restrict output by directly restricting inputs; instead it indirectly restricts output by charging such a high tariff that it induces the downstream firms to randomize over accepting/rejecting; increasing the probability of rejection reduces expected output. We applied our theory to the experimental data and showed that some behavior, impossible to rationalize in existing models, could now be rationalized.

Given that our experimental setting had one subject making a take-it-or-leave it offer to another, it was natural to compare our findings on the division of surplus between parties to the existing experimental literature on ultimatum and dictator games. Our results were generally closer to the game-theoretic predictions involving 100% surplus for the offering party. One of the standard reasons given for why experiments deviate from the game-theoretic predictions is that subjects are concerned with "fairness" in addition to their monetary payoff in the game. We derived a somewhat counterintuitive theoretical result that a concern for "fairness" on the part of downstream subjects in treatments with secret contracts may increase the expected profit of the upstream subject. The reason is that the increased probability of rejection indirectly leads to the desired restriction of expected output.

Appendix A

■ Proofs of Propositions 1–2 and other mathematical details referred to in the text follow.

☐ **Verification of footnote 18.** To verify that $q = 2$ in the joint-profit-maximizing outcome, define industry profit as $\Pi(q) = q P(q)$. Then $\Pi(1) = 60$, $\Pi(2) = 100$, $\Pi(3) = 90$, $\Pi(4) = 72$, $\Pi(5) = 25$, and $\Pi(q) = 0$ for $q \geq 6$. $\Pi(q)$ is highest for $q = 2$. To verify that $q = 4$ in equilibrium with secret contracts, note that equilibrium outputs (q_1^*, q_2^*) must be best responses to each other:

$$q_1^* = b(q_2^*) \quad \text{and} \quad q_2^* = b(q_1^*). \tag{A1}$$

This rules out $(q_1 = 1, q_2 = 1)$ for equilibrium quantities because, given that D_2 is selling one unit, U and D_1 can increase their joint surplus by tranferring two units rather than one (generating revenue of 60 for D_1 rather than 50). Indeed, it can be shown that 2 is the best response to a rival's output of 0, 1, or 2. Thus $(q_1 = 2, q_2 = 2)$ satisfies (A1) and hence is an equilibrium quantity vector.

It can be checked on a case-by-case basis that for any other quantity vector satisfying (A1), each component must be 5 or more. Such equilibria arise in our setup due to the assumption of zero production costs. They would be ruled out by assuming positive production costs (no matter how small) or by imposing the trembling-hand perfection refinement. In any event, such outcomes did not arise in our experiments.

☐ **Robust variance estimators.** The formula for the robust variance matrix estimator for regressions (1) and (3) is

$$\left(\frac{N-1}{N-k} \right) \left(\frac{M}{M-1} \right) (X'X)^{-1} \left(\sum_{m=1}^{M} u_m' u_m \right) (X'X)^{-1}, \tag{A2}$$

where X is the matrix of regressors, N is the number of observations, M is the number of clusters G_m (a "cluster" being

the group of observations for a single upstream firm), k is the number of regressors, $u_m = \sum_{j \in G_m} \hat{\epsilon}_j x_j$, $\hat{\epsilon}_j = y_j - x_j \hat{\beta}$, y_j is the dependent variable for observation j, x_j is the vector of independent variables for observation j, and $\hat{\beta}$ are the coefficient estimates. See StataCorp (1997) for further details.

The formula for the robust variance of $\Delta = \Phi(\hat{\beta}_0 + \hat{\beta}_1) - \Phi(\hat{\beta}_0)$ in probits (2) and (4) is

$$\left(\frac{M}{M-1} \right) d(X'X)^{-1} \left(\sum_{m=1}^{M} u'_m u_m \right) (X'X)^{-1} d', \tag{A3}$$

where

$$u_m = \sum_{j \in G_m} \left[\frac{\phi(x_j \hat{\beta})}{\Phi(x_j \hat{\beta})} \mathbf{1}_{\{z_j = \bar{z}\}} - \frac{\phi(x_j \hat{\beta})}{1 - \Phi(x_j \hat{\beta})} (1 - \mathbf{1}_{\{z_j = \bar{z}\}}) \right],$$

d is the row vector $(0 \quad \phi(\hat{\beta}_0 + \hat{\beta}_1))$, $\hat{\beta}$ are the coefficient estimates, and the rest of the variables are defined as in the previous paragraph. See StataCorp (1997) for further details.

☐ **Verification of Figure 7.** *Case 1 ($x_i = 1$, passive beliefs).* D_i believes D_j receives, and accepts, equilibrium contract $(2, 36)$. If $a_i = 1$, D_i's beliefs imply $q = 3$, $P = 30$, and D_i earns a gross profit of 30. Hence $\alpha_i = 1$ if $T_i < 30$ and $\alpha_i = 0$ if $T_i > 30$.

Case 2 ($x_i = 2$, passive beliefs). D_i believes D_j receives, and accepts, equilibrium contract $(2, 36)$. If $a_i = 1$, D_i's beliefs imply $q = 4$, $P = 18$, and D_i earns a gross profit of 36. Hence $\alpha_i = 1$ if $T_i < 36$ and $\alpha_i = 0$ if $T_i > 36$.

Case 3 ($x_i = 3$, passive beliefs). D_i believes D_j receives, and accepts, equilibrium contract $(2, 36)$. If $a_i = 1$, D_i's beliefs imply $q = 5$, $P = 5$, and D_i earns a gross profit of 15. Hence $\alpha_i = 1$ if $T_i < 15$ and $\alpha_i = 0$ if $T_i > 15$.

Case 4 ($x_i = 1$, symmetric beliefs). D_i believes D_j receives the same contract as itself, $(1, T_i)$. If $a_i = 1$, D_i's gross profit is at least 50 (exactly 50 if $a_j = 1$, 60 if $a_j = 0$). Hence $\alpha_i = 1$ if $T_i < 50$. Suppose $T_i > 50$. The unique continuation equilibrium is in mixed strategies. For D_j to be indifferent between $a_j = 1$ and $a_j = 0$, it must earn zero expected net profit from $a_j = 1$: i.e., $50\alpha_i + 100(1 - \alpha_i) - T_i = 0$. This implies

$$\alpha_i = \max \left\{ \frac{60 - T_i}{10}, 0 \right\}.$$

Case 5 ($x_i = 2$, symmetric beliefs). D_i believes D_j receives the same contract as itself, $(2, T_i)$. If $a_i = 1$, D_i's gross profit is at least 36 (exactly 36 if $a_j = 1$, 100 if $a_j = 0$). Hence $\alpha_i = 1$ if $T_i < 36$. Suppose $T_i > 36$. Following the logic of case 4, α_i must solve $36\alpha_i + 100(1 - \alpha_i) - T_i = 0$ in the unique continuation equilibrium, implying

$$\alpha_i = \max \left\{ \frac{100 - T_i}{64}, 0 \right\}.$$

Case 6 ($x_i = 3$, symmetric beliefs). D_i believes D_j receives the same contract as itself, $(3, T_i)$. If $a_i = 1$, D_i's gross profit is either 0 (if $a_j = 1$) or 90 (if $a_j = 0$). Following the logic of case 4, α_i must solve $(0)(\alpha_i) + 90(1 - \alpha_i) - T_i = 0$ in the unique continuation equilibrium, implying

$$\alpha_i = \max \left\{ \frac{90 - T_i}{90}, 0 \right\}.$$

Proof of Proposition 1. An extremal PBE is the solution to the problem of maximizing (A4) subject to (A5) and (A6). Let $M_1(s)$ be the value function associated with the problem

$$\max_{x_i, T_i \in [0, \infty); \alpha_{si}, \alpha_{pi} \in [0,1]} \{ [s\alpha_{si} + (1 - s)\alpha_{pi}] T_i \} \tag{A4}$$

subject to

$$[s\alpha_{si} + (1 - s)\alpha_{pi}] R(x_i, x_i) + [s(1 - \alpha_{si}) + (1 - s)(1 - \alpha_{pi})] R(x_i, 0) \geq T_i \tag{A5}$$

$$[s\alpha_{si} + (1 - s)\alpha_{pi}] T_i \geq M_2(x_i, \alpha_{si}, \alpha_{pi}, s), \tag{A6}$$

where (A5) must hold with equality if $\alpha_{si} \in (0, 1)$ or $\alpha_{pi} \in (0, 1)$, and where M_2 is itself a value function, associated with the optimization of (A7) subject to (A8):

$$\min_{\alpha_{si}^d, \alpha_{pi}^d \in [0,1]} \left\{ \limsup_{x_i^d, T_i^d \in [0, \infty)} \{ [s\alpha_{si}^d + (1 - s)\alpha_{pi}^d] T_i^d \} \right\} \tag{A7}$$

subject to

$$\alpha_{ji}^d = \begin{cases} 1 & \text{if } \xi_{ji}^d > T_i^d \\ 0 & \text{if } \xi_{ji}^d < T_i^d \end{cases} \qquad \text{for } j = s, p, \tag{A8}$$

where

$$\xi_{si}^d \equiv [s\alpha_{si}^d + (1-s)\alpha_{pi}^d]R(x_i^d, x_i^d) + [s(1-\alpha_{si}^d) + (1-s)(1-\alpha_{pi}^d)]R(x_i^d, 0) \tag{A9}$$

$$\xi_{pi}^d \equiv [s\alpha_{si} + (1-s)\alpha_{pi}]R(x_i^d, x_i) + [s(1-\alpha_{si}) + (1-s)(1-\alpha_{pi})]R(x_i^d, 0). \tag{A10}$$

Expression (A4) is U's expected profit from selling to D_i given that a symmetric (respectively passive) type accepts the contract with probability α_{si} (respectively α_{pi}). Individual-rationality constraint (A5) ensures that D_i would not strictly gain by rejecting the contract. With probability $s\alpha_{si} + (1-s)\alpha_{pi}$, rival D_j accepts its equilibrium contract, in which event D_i earns $R(x_i, x_i)$ by accepting; with the complementary probability, D_j rejects its equilibrium contract, in which event D_i earns $R(x_i, 0)$. Constraint (A5) does not depend on D_i's type since both types share the same conjectures about x_j in equilibrium. Incentive-compatibility constraint (A6) ensures that U does not strictly gain by offering a deviating contract to D_i. To understand constraint (A6), consider the optimization problem associated with value function M_2. In (A7), U offers the best possible deviating contract in the face of acceptance probabilities making deviation as unattractive as possible (to ensure the PBE is extremal for U). The probabilities can be chosen freely in [0, 1] if D_i is indifferent between accepting and not. If D_i strictly prefers to accept (respectively reject) given its beliefs, the acceptance probability must equal one (respectively zero), captured in constraint (A8). Implicit in expression (A9)—D_i's expected gross surplus from accepting conditional on being a symmetric type—is that D_i conjectures that D_j receives the same deviating contract (involving input quantity x_i^d) and responds with either acceptance probability α_{si}^d or α_{pi}^d depending on its type. Implicit in expression (A10)—D_i's expected gross surplus from accepting conditional on being a passive type—is that D_i continues to conjecture that D_j receives the equilibrium contract (involving input quantity x_i) and D_j responds with either acceptance probability α_{si}^d or α_{pi}^d depending on its type.

The key insight in the proof is that if $\alpha_{si}^d > 0$ in the optimal deviating contract associated with M_2, then (A6) does not bind, implying we can take $\alpha_{si}^d = 0$ without loss of generality. To see this, let $\tilde{M}_1$ be the value function associated with the otherwise similar problem as M_1 except that (A6) is omitted. Let $\tilde{M}_2$ be the value function associated with the otherwise similar problem as M_2 except α_{si}^d is constrained to be positive. We proceed by showing $\tilde{M} \leq \tilde{M}_1$ by considering similarities in the associated optimization problems. The objective functions (A4) and (A7) are similar except that the probabilities are chosen to minimize rather than maximize in (A7). If α_{si}^d is set at some positive value in the problem associated with M_2, constraint (A8) becomes $\xi_{si}^d \geq T_i^d$ for $j = s$, which can be shown to be equivalent to (A5). The remaining constraint in the problem associated with M_2—namely constraint (A8) for $j = p$—has no analogue in the problem associated with $\tilde{M}_1$. In sum, all factors in the comparison of the associated optimization problems point to $\tilde{M}_1$ weakly exceeding M_2. But then (A6) must be satisfied: the left-hand side equals $\tilde{M}_1$, and the right-hand side equals $M_2 = \tilde{M}_2 \leq \tilde{M}_1$, where the first equality holds setting α_{si}^d to be positive in the problem associated with M_2.

Without loss of generality, then, we can substitute $\hat{M}_2$ for M_2, where $\hat{M}_2$ is the value function associated with the otherwise similar optimization problem as M_2 except that we set $\alpha_{si}^d = 0$. That is,

$$\hat{M}_2(x_i, \alpha_{si}, \alpha_{pi}, s) \equiv \min_{\alpha_{pi}^d \in [0,1]} \left\{ \limsup_{x_i^d, T_i^d \in [0,\infty)} \{(1-s)\alpha_{pi}^d T_i^d\} \right\} \tag{A11}$$

subject to

$$\alpha_{pi}^d = 1 \text{ if } \xi_{pi}^d > T_i^d. \tag{A12}$$

It is evident that $\alpha_{pi}^d = 0$ if $\xi_{pi}^d \leq T_i^d$; so $\xi_{pi}^d > T_i^d$, implying $\alpha_{pi}^d = 1$. Hence

$$\hat{M}_2(x_i, \alpha_{si}, \alpha_{pi}, s) = \max_{x_i^d \in [0,\infty)} \{(1-s)\xi_{pi}^d\}. \tag{A13}$$

The original problem can be simplified by making the change of variables $\alpha_i \equiv s\alpha_{si} + (1-s)\alpha_{pi}$. It can be simplified further by noting that constraint (A5) binds at an optimum, since T_i should be as high as possible subject to (A5). After these changes, and after substituting from (A13) into constraint (A6), our original problem reduces to the one in the statement of Proposition 1.   Q.E.D.

Proof of Proposition 2. If constraint (6) does not bind in the maximization of (5), it is evident that the solution involves $\alpha_i^* = 1$ and $x_i^* = q^m/2$. Substituting these values into the expression for T^* in Proposition 1, we obtain $T_i^* = R(q^m/2, q^m/2) = \Pi^m/2$.

Constraint (6) binds if it is violated by this solution. Substituting $\alpha_i = 1$ and $x_i = q^m/2$ into (6), it binds if and only if

$$R(q^m/2, q^m/2) < \max_{x_i^d \in [0,\infty)} \{(1-s)R(x_i^d, q^m/2)\}$$

$$= (1-s)R(b(q^m/2), q^m/2).$$

Rearranging, it is evident that (6) does not bind if and only if $s \geq \hat{s}$.

To solve the maximization problem in Proposition 1 for the case $s = 0$, note that constraint (6) becomes

$$\alpha_i [\alpha_i R(x_i, x_i) + (1 - \alpha_i) R(x_i, 0)] \geq \max_{x_i^d \in [0, \infty)} \{\alpha_i R(x_i^d, x_i) + (1 - \alpha_i) R(x_i^d, 0)\}. \tag{A14}$$

Note that $x_i^d = x_i$ is a feasible solution to the maximization problem on the left-hand side of (A14). Substituting $x_i^d = x_i$ reveals that the left-hand side of (A14) is greater than the right-hand side unless $\alpha_i = 1$. Thus, a necessary condition for constraint (A14) not to be violated is $\alpha_i = 1$. Substituting $\alpha_i = 1$ into (A14), we obtain

$$R(x_i, x_i) \geq \max_{x_i^d \in [0, \infty)} \{R(x_i^d, x_i)\}$$
$$= R(b(x_i), x_i).$$

For the constraint to be satisfied, $x_i = b(x_i)$; i.e., $x_i = q^c$. It is easy to show from here that the equilibrium contract is (q^c, Π^c) and is accepted with probability one by downstream firms, all of whom are passive types.   *Q.E.D.*

☐ **Verification of Figure 8.** The extremal PBE cannot involve $x_i = 0$ or $x_i \geq 6$ in our experimental setting since U's profit would be zero in those cases. The extremal PBE can thus be computed by examining the remaining cases— $x_i = 1, 2, 3, 4, 5$—individually. Tedious calculations show that constraint (6) can only be satisfied if $x_i = 2$. Hence $x_i^* = 2$. The optimal deviation also involves two units; i.e., $x_i^d = 2$. Substituting these values into (5) and (6), α_i^* in the extremal PBE is the solution to $\max_{\alpha_i \in [0,1]} \{100 - 64\alpha_i^2\}$ subject to $100\alpha_i - 64\alpha_i^2 \geq (1 - s)(100 - 64\alpha_i)$. The solution is $\alpha_i^* = 1 - s$.

☐ **Computation of Shapley values.** In the nonintegrated treatments, the six permutations of the players are

$$\{U, D_1, D_2\}, \{U, D_2, D_1\}, \{D_1, U, D_2\}, \{D_2, U, D_1\}, \{D_1, D_2, U\}, \{D_2, D_1, U\},$$

all equally likely. These permutations correspond to the following coalitions including U and its predecessors:

$$\{U\}, \{U\}, \{D_1, U\}, \{D_2, U\}, \{D_1, D_2, U\}, \{D_2, D_1, U\},$$

again, all equally likely. There is no surplus produced in coalition $\{U\}$, since there is no downstream outlet, so U's marginal contribution must be zero. The remaining coalitions can produce industry profit 100, and U is pivotal for this surplus to be realized. The Shapley value allocates surplus $(0 + 0 + 100 + 100 + 100 + 100)/6$ to U. As a fraction of the total surplus available, 100, this equals 2/3. The remaining 1/3 of the surplus is divided equally between the symmetric downstream players.

In *INTEGR*, there are two permutations of the players: $\{U\text{-}D_1, D_2\}$ and $\{D_2, U\text{-}D_1\}$. These permutations lead to the following coalitions including $U\text{-}D_1$ and its predecessors: $\{U\text{-}D_1\}$ and $\{D_2, U\text{-}D_1\}$. $U\text{-}D_1$ can generate surplus 100 by selling two units through its downstream subsidiary, so it is allocated surplus $(100 + 100)/2$ by the Shapley value, 100% of the total surplus available. D_2 thus is allocated none of the surplus.

Appendix B

■ This Appendix is an English translation of the instructions to the experimental subjects. The original German-language instructions are available upon request from the authors.

☐ **Introductory instructions.** Welcome to our experiment! In the next hour you will make decisions at a computer. One thing is important right from the start: please be quiet during the entire experiment and please do not talk to your neighbors. The experiment runs over 10 periods. Before the first period starts, the use of the computer will be explained in detail in a trial round.

In the experiment we will use a fictitious currency called ECU. In the beginning you will get a starting capital in ECU. During the experiment you can earn some real money, but losses are also possible. Should it happen that some participant loses the entire starting capital and that this participant has a negative total profit for more than three periods, we have to stop the experiment.

After the last period, you will be paid 1 DM for every 20 ECU you earned during the experiment. Concerning the payment, there is strict anonymity with respect to the other participants as well as with respect to us. We will record no data in connection with your name.

☐ **Further instructions for *PUBLIC*, *SECRAN*, and *SECFIX*.**

What is the experiment about? The experiment is about decision making in a market with one manufacturer and two retailers. Some of you will make decisions for a manufacturer, others for a retailer. You will be a manufacturer or a retailer

for all 10 periods of the experiment. Consumers in the market are simulated by the computer program. You will be told whether you are a manufacturer or a retailer during the trial period. Currently, you are all reading the same instructions.

(This paragraph only for *PUBLIC*, *SECRAN*.) Note that in every period the manufacturer-retailer groups change. You do not know which retailer or manufacturer you will meet.

(This paragraph only for *SECFIX*.) Note that in each period the manufacturer-retailer groups remain the same. You are always together with the same retailer and manufacturer in the market.

The basic structure of the market is such that the manufacturer produces the product and sells it to the two retailers. In their stores, retailers sell the product on to the consumers.

What are you supposed to do as a manufacturer or retailer? A manufacturer has to decide how many units of the product he wants to sell at which price to the two retailers. This decision has the form of an offer to the retailers: each retailer is offered a specified quantity of the product at a specified total price. The manufacturer may also decide not to offer the product to one or both retailers.

If a retailer receives an offer, he has to decide either to accept the offer or to reject it. If he accepts the offer, he receives the number of units of the product specified in the offer and has to pay the total price. If he rejects the offer, he does not receive the product and does pay anything to the manufacturer.

What price do retailers get for the product in their stores? The market price paid by the consumers is determined by the computer program in the following way. The market price per unit depends on the total quantity supplied together by both retailers. Here the following relationship between the quantity supplied and the market price holds:

Total quantity	Market price
1	60
2	50
3	30
4	18
5	5
6 or more	0

The table reads as follows. In the left column, one finds the total quantity of the product supplied by both retailers. For each total quantity there is exactly one market price. Take an example: Suppose retailer 1 received 2 units from the manufacturer and retailer received 1 unit. As the total number of units is 3, the market price per unit is 30 ECU.

Retailers' revenues are the number of units supplied (i.e., bought from the manufacturer) multiplied by the market price. In the example, retailer 1 has revenues of 2 × 30 ECU = 60 ECU, while retailer 1 has revenues of 1 × 30 ECU = 30 ECU.

Retailers' stores are run without cost. The profit of a retailer is thus the revenues minus the payment to the manufacturer. Suppose that, in the example, retailer 2 agreed to pay 35 ECU for the 1 unit he received. Then he would actually make a loss of 5 ECU. If he agreed to pay only 5 ECU, a profit of 25 ECU would result.

Also the manufacturer produces without cost. The manufacturer's profit is thus simply the payments of the two retailers.

What kind of information do you get?

(This paragraph only for *PUBLIC*.) Each retailer knows his own offer as well as the offer of the other retailer. The manufacturer and both retailers are told at the end of each period whether or not the retailers accepted the offers. The manufacturer and both retailers are informed at the end of each period about the profit of all three participants involved.

(This paragraph only for *SECRAN*, *SECFIX*.) Each retailer knows only his own offer but not the offer of the other retailer. Each retailer is told his own profit at the end of each period. The manufacturer is told whether or not the retailers accepted the offers at the end of each period. The manufacturer is informed about his own profit and the profit of the two retailers at the end of each period.

☐ **Further instructions for *INTEGR*.**

What is the experiment about? The experiment is about decision making in a market with one manufacturer and one retailer. Some of you will make decisions for a manufacturer, others for a retailer. You will be a manufacturer or a retailer for all 10 periods of the experiment. Consumers in the market are simulated by the computer program. You will be told whether you are a manufacturer or a retailer during the trial period. Currently, you are all reading the same instructions.

Note that in every period the manufacturer-retailer-groups change. More precisely speaking, you will never meet in the same manufacturer-retailer-group twice.

The basic structure of the market is such that both—not only the retailer, but also the manufacturer—own a store. The product is produced only by the manufacturer. The manufacturer may sell the product in his own store as well as in the store of the retailer. In the stores, the product is sold to the consumers.

What are you supposed to do as a manufacturer or retailer? A manufacturer has to decide how many units of the product he wants to sell at which price to the retailer. This decision has the form of an offer to the retailer: the retailer is offered a specified quantity of the product at a specified total price. The manufacturer may also decide not to offer the product to the retailer.

If the retailer receives an offer, he has to decide either to accept the offer or to reject it. If he accepts the offer, he receives the number of units of the product specified in the offer and has to pay the total price. If he rejects the offer, he does not receive the product and does not pay anything to the manufacturer.

Afterwards, the manufacturer gets a message indicating whether or not the retailer accepted the offer. Then the manufacturer decides how much to supply in his own store. If he does not want to supply the product in his own store, the manufacturer enters a zero in the computer.

What price is received for the product in the stores? The market price paid by the consumers is determined by the computer program in the following way. The market price per unit depends on the total quantity supplied together by the store of the retailer and the store of the manufacturer. Here the following relation between the quantity supplied and the market price holds:

Total quantity	Market price
1	60
2	50
3	30
4	18
5	5
6 or more	0

The table reads as follows. In the left column, one finds the total quantity of the product supplied in both stores. For each total quantity there is exactly one market price. Take an example: Suppose retailer 1 supplies 2 units and the manufacturer supplies 1 unit. As the total number of units is 3, the market price per unit is 30 ECU.

Revenues in the stores are the number of units supplied multiplied by the market price. In the example, the store of the retailer has revenues of 2×30 ECU = 60 ECU, while the store of the manufacturer has revenues of 1×30 ECU = 30 ECU.

The stores are run without cost. For the manufacturer also, production is without cost. The manufacturer's profit is thus the sum of the retailer's payment and the revenues of his store.

The profit of the retailer is his revenues minus the payment to the manufacturer. Suppose that, in the example, the retailer agreed to pay 65 ECU for the 2 units he received. Then he would actually make a loss of 5 ECU. If he agreed to pay only 35 ECU, a profit of 25 ECU would result.

What kind of information do you get? As mentioned, the manufacturer gets a message whether or not the retailer accepted his offer before deciding about the number of units to be supplied in his store. The retailer is told his own profit at the end of each period. The manufacturer is informed about his own profit and the profit of the retailer at the end of each period.

References

ABBINK, K. AND SADRIEH, K. "RatImage—Research Assistant Toolbox for Computer-Aided Human Behavior Experiments." Discussion Paper no. B–325, University of Bonn, 1995.

ALLEN, B.T. "Vertical Integration and Market Foreclosure: The Case of Cement and Concrete." *Journal of Law and Economics*, Vol. 14 (1971), pp. 251–274.

BOLTON, P. AND WHINSTON, M.D. "Incomplete Contracts, Vertical Integration, and Supply Assurance." *Review of Economic Studies*, Vol. 60 (1993), pp. 121–148.

BORK, R.H. *The Antitrust Paradox*. New York: Basic Books, 1978.

CARLTON, D.W. "Vertical Integration and Market Foreclosure: Comment." *Brookings Papers on Economic Activity: Microeconomics*, (1990), pp. 277–280.

CHIPTY, T. "Vertical Integration, Market Foreclosure, and Consumer Welfare in the Cable Television Industry." *American Economic Review*, Vol. 91 (2001), pp. 428–453.

CHOI, J.P. AND YI, S.-S. "Vertical Foreclosure with the Choice of Input Specifications." *RAND Journal of Economics*, Vol. 31 (2000), pp. 717-743.

CHURCH, J. AND GANDAL, N. "Systems Competition, Vertical Merger and Foreclosure." *Journal of Economics and Management Strategy*, Vol. 9 (2000), pp. 25–51.

DEGRABA, P. "Most-Favored-Customer Clauses and Multilateral Contracting: When Nondiscrimination Implies Uniformity." *Journal of Economics and Management Strategy*, Vol. 5 (1996), pp. 565–579.

FORSYTHE, R., HOROWITZ, J.L., SAVIN, N.E., AND SEFTON, M. "Fairness in Simple Bargaining Experiments." *Games and Economic Behavior*, Vol. 6 (1994), pp. 347–369.

FRANCIOSI, R., KUJAL, P., MICHELITSCH, R., SMITH, V., AND DENG, G. "Fairness: Effect on Temporary and Equilibrium Prices in Posted-Offer Markets." *Economic Journal*, Vol. 105 (1995), pp. 938–950.

FRIEDMAN, D. AND SUNDER, S. *Experimental Methods—A Primer for Economists.* New York: Cambridge University Press, 1994.

FUDENBERG, D. AND MASKIN, E. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica*, Vol. 54 (1986), pp. 533–554.

GOULD, J.R. "Price Discrimination and Vertical Control: A Note." *Journal of Political Economy*, Vol. 85 (1977), pp. 1063–1071.

GRIMM, C.M., WINSTON, C., AND EVANS, C.A. "Foreclosure of Railroad Markets: A Test of Chicago Leverage Theory." *Journal of Law and Economics*, Vol. 35 (1992), pp. 295–310.

GÜTH, W., MARCHAND, N., AND RULLIERE, J.-L. "On the Reliability of Fairness—An Experimental Study." Discussion Paper no. 102, Humboldt University Berlin, 1997.

———, OCKENFELS, P., AND RITZBERGER, K. "On Durable Goods Monopolies: An Experimental Study of Intrapersonal Price Competition and Price Discrimination Over Time." *Journal of Economic Psychology*, Vol. 16 (1995), pp. 247–274.

———, SCHMITTBERGER, R., AND SCHWARZE, B. "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior and Organization*, Vol. 3 (1982), pp. 367–388.

HART, O. AND TIROLE, J. "Vertical Integration and Market Foreclosure." *Brookings Papers on Economic Activity: Microeconomics*, (1990), pp. 205–276.

HOFFMAN, E., MCCABE, K., SHACHAT, K., AND SMITH, V. "Preferences, Property Rights, and Anonymity in Bargaining Experiments." *Games and Economic Behavior*, Vol. 7 (1994), pp. 346–380.

KLASS, M.W. AND SALINGER, M.A. "Do New Theories of Vertical Foreclosure Provide Sound Guidance for Consent Agreements in Vertical Merger Cases?" *Antitrust Bulletin*, Vol. 40 (1995), pp. 667–698.

KLEIN, B., CRAWFORD, R.C., AND ALCHIAN, A.A. "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process." *Journal of Law and Economics*, Vol. 21 (1978), pp. 297–326.

KWOKA, J.E., JR. AND WHITE, L.J. *The Antitrust Revolution*, 3d ed. New York: Oxford University Press, 1998.

MACAULEY, S. "Noncontractual Relations in Business." *American Sociological Review*, Vol. 28 (1963), pp. 55–70.

MARTIN, S. "Strategic and Welfare Effects of Bundling." *Economics Letters*, Vol. 62 (1999), pp. 371–376.

MAS-COLELL, A., WHINSTON, M.D., AND GREEN, J.R. *Microeconomic Theory*. New York: Oxford University Press, 1995.

MASON, C.F. AND PHILLIPS, O.R. "Vertical Integration and Collusive Incentives: An Experimental Analysis." *International Journal of Industrial Organization*, Vol. 18 (2000), pp. 471–496.

MCAFEE, R.P. AND SCHWARTZ, M. "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity." *American Economic Review*, Vol. 84 (1994), pp. 210–230.

MCBRIDE, M.E. "Spatial Competition and Vertical Integration: Cement and Concrete Revisited." *American Economic Review*, Vol. 73 (1983), pp. 1011–1022.

O'BRIEN, D.P. AND SHAFFER, G. "Vertical Control with Bilateral Contracts." *RAND Journal of Economics*, Vol. 23 (1992), pp. 299–308.

ORDOVER, J.A., SALONER, G., AND SALOP, S.C. "Equilibrium Vertical Foreclosure." *American Economic Review*, Vol. 80 (1990), pp. 127–142.

PASS, C.L AND HAWKINS, K.H. "Exclusive Dealing, Supplier Ownership of Outlets and the Public Interest: The Petrol Case." *Antitrust Bulletin*, Vol. 17 (1972), pp. 567–595.

PERRY, M.K. "Vertical Integration: Determinants and Effects." In R. Schmalensee and R.D. Willig, eds., *Handbook of Industrial Organization*. New York: North-Holland, 1989.

REIFFEN, D. "Equilibrium Vertical Foreclosure: Comment." *American Economic Review*, Vol. 82 (1992), pp. 694–697.

——— AND VITA, M. "Is There New Thinking on Vertical Mergers? A Comment." *Antitrust Law Journal*, Vol. 63 (1995), pp. 917–941.

REY, P. AND TIROLE, J. "A Primer on Foreclosure." In M. Armstrong and R.H. Porter, eds., *Handbook of Industrial Organization*, Vol. 3. New York: North-Holland, forthcoming.

REYNOLDS, S.S. "Durable-Goods Monopoly: Laboratory Market and Bargaining Experiments." *RAND Journal of Economics*, Vol. 31 (2000), pp. 375–394.

RIORDAN, M.H. "Anticompetitive Vertical Integration by a Dominant Firm." *American Economic Review*, Vol. 88 (1998), pp. 1232–1248.

——— AND SALOP, S.C. "Evaluating Vertical Mergers: A Post-Chicago Approach." *Antitrust Law Journal*, Vol. 63 (1995), pp. 513–568.

ROSENGREN, E.S. AND MEEHAN, J.W., JR. "Empirical Evidence on Vertical Foreclosure." *Economic Inquiry*, Vol. 32 (1994), pp. 303–317.

ROTH, A.E. "Bargaining Experiments." In J.H. Kagel and A.E. Roth, eds., *The Handbook of Experimental Economics*. Princeton, N.J.: Princeton University Press, 1995.

SALINGER, M.A. "Vertical Mergers and Market Foreclosure." *Quarterly Journal of Economics*, Vol. 103 (1988), pp. 345–356.

———. "The Meaning of 'Upstream' and 'Downstream' and the Implications for Modeling Vertical Mergers." *Journal of Industrial Economics*, Vol. 37 (1989), pp. 373–387.

SCHMALENSEE, R. "A Note on the Theory of Vertical Integration." *Journal of Political Economy*, Vol. 81 (1973), pp. 442–449.

SEGAL, I. "Contracting with Externalities." *Quarterly Journal of Economics*, Vol. 114 (1999), pp. 337–388.

SELTEN, R. AND STOECKER, R. "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames: A Learning Theory Approach." *Journal of Economic Behavior and Organization*, Vol. 7 (1988), pp. 47–70.

SMITH, V.L. "Microeconomic Systems as an Experimental Science." *American Economic Review*, Vol. 72 (1982), pp. 923–925.

SNYDER, C.M. "Empirical Studies of Vertical Foreclosure." *Industry Economics Conference Papers and Proceedings*, University of Melbourne and Bureau of Industry Economics, Vol. 23 (1995), pp. 98–127.

STATACORP. *Stata Statistical Software: Release 5.0*. College Station, Tex.: Stata Corporation, 1997.

TIROLE, J. *The Theory of Industrial Organization*. Cambridge, Mass.: MIT Press, 1988.

U.S. DEPARTMENT OF JUSTICE, FEDERAL TRADE COMMISSION. *Merger Guidelines, 1969*. Washington, D.C.: U.S. Government Printing Office, 1969.

———. *Merger Guidelines, 1982*. Washington, D.C.: U.S. Government Printing Office, 1982.

———. *Merger Guidelines, 1984*. Washington, D.C.: U.S. Government Printing Office, 1984.

———. *Horizontal Merger Guidelines, 1992*. Washington, D.C.: U.S. Government Printing Office, 1992.

VERNON, J.M. AND GRAHAM, D.A. "Profitability of Monopolization by Vertical Integration." *Journal of Political Economy*, Vol. 79 (1971), pp. 924–925.

WARREN-BOULTON, F.R. "Vertical Control with Variable Proportions." *Journal of Political Economy*, Vol. 82 (1974), pp. 783–802.

WATERMAN, D. AND WEISS, A.A. "The Effects of Vertical Integration Between Cable Television Systems and Pay Cable Networks." *Journal of Econometrics*, Vol. 72 (1996), pp. 357–395.

WHITE, H. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica*, Vol. 8 (1980), pp. 817–838.

WHITE, L.J. "The Revolution in Antitrust Analysis of Vertical Relationships: How Did We Get from There to Here?" In R.J. Larnere and J.W. Meehan, eds., *Economics and Antitrust Policy*. Westpoint, Conn.: Quorum, 1989.

WILLIAMSON, O.E. "The Vertical Integration of Production: Market Failure Considerations." *American Economic Review*, Vol. 61 (1971), pp. 112–123.

———. "Vertical Integration and Market Foreclosure: Comment." *Brookings Papers on Economic Activity: Microeconomics*, (1990), pp. 280–283.