

VERTICAL MERGERS, FORECLOSURE AND RAISING RIVALS' COSTS – EXPERIMENTAL EVIDENCE*

HANSTHEO NORMANN[†]

The hypothesis that vertically integrated firms have an incentive to foreclose the input market because foreclosure raises its downstream rivals' costs is the subject of much controversy in the theoretical industrial organization literature. A powerful argument against this hypothesis is that, absent commitment, such foreclosure cannot occur in Nash equilibrium. The laboratory data reported in this paper provide experimental evidence in favor of the hypothesis. Markets with a vertically integrated firm are significantly less competitive than those where firms are separate. While the experimental results violate the standard equilibrium notion, they are consistent with the quantal-response generalization of Nash equilibrium.

I. INTRODUCTION

THE THEORETICAL INDUSTRIAL ORGANIZATION LITERATURE HAS SUGGESTED THAT 'raising rivals' costs' may be a profitable strategy in oligopoly. Raising-rivals'-costs arguments are based on the simple fact that it is easier to compete with less efficient firms. If a firm's production costs are raised, it will reduce output and increase price, and the other firms in the market will benefit from this as they can increase their market shares and prices. It follows that firms may pursue strategies from which they do not benefit directly (e.g., through production efficiencies) but rather indirectly because a competitor's costs are affected negatively. Cost-raising strategies were first proposed by Salop and Scheffman [1983, 1987] and include boycott and other exclusionary behavior, advertising, R&D, and lobbying for standards and regulation.

Ordoover, Saloner and Salop's [1990], henceforth OSS [1990], raising-rival's-cost paper has received particular attention because it sets out to establish a connection between vertical integration and foreclosure.¹ In OSS

*The author is grateful to two referees, the Editor, Dirk Engelmann, Wieland Müller, Martin Sefton, Georg Weizsaecker, and seminar audiences at ESA Amsterdam, DIW Berlin, TU Berlin, University of Goettingen, Royal Holloway College London, University of Mannheim, and University of Zurich for useful comments. Thanks also to Volker Benndorf, Nikos Nikiforakis, Holger Rau, Silvia Platoni and Brian Wallace for their help on the research. Financial support through Leverhulme grant F/07537/S is gratefully acknowledged.

[†]Authors' affiliation: Düsseldorf Institute for Competition Economics (DICE), University of Düsseldorf, Germany.

email: normann@dice.uni-duesseldorf.de

¹Salinger [1988], Hart and Tirole [1990] and OSS [1990] are the first generation of post-Chicago foreclosure theories. See Rey and Tirole [2007] or Riordan [2008] for recent surveys.

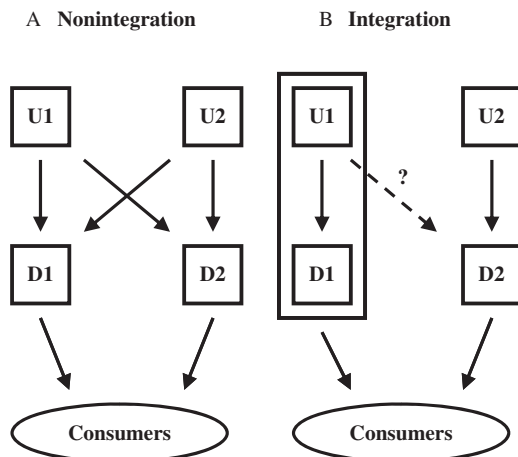


Figure 1
Market Structure

[1990], foreclosure means that a vertically integrated firm withdraws from the input-good market, that is, it stops supplying the input to nonintegrated downstream firms. OSS [1990] argue that firms have an incentive to integrate vertically and engage in such foreclosure because they gain from a raising-rival's-costs effect. The logic is that, when a vertically integrated firm forecloses, competition in the input-good market becomes weaker. The reduction in competition implies higher input costs for the nonintegrated downstream firms. Since the downstream unit of the integrated firm benefits when downstream rivals' costs are raised, the integrated firm is better off with the foreclosure strategy than when it actively competes. In other words, it pays for the integrated firm to forgo business in the input-good market and instead to gain because the downstream rivals become less competitive. Only a vertically integrated firm can pursue such a strategy profitably. It would not make sense for nonintegrated upstream firms to foreclose (they would simply lose money) nor would this strategy be feasible for nonintegrated downstream firms.

Consider OSS' [1990] setup as shown in Figure 1. There are two upstream firms and two downstream firms. In panel A, neither firm is vertically integrated and both upstream firms compete for both downstream firms. For example, if the two upstream firms are Bertrand competitors, the input market will be perfectly competitive. Now suppose $U1$ and $D1$ merge as in panel B of Figure 1. Since $U1$ will supply $D1$ with the input internally, $U2$ can no longer compete for $D1$. More importantly, if $U1$ - $D1$ stops delivering $D2$ (or alternatively if it charges a very high price for the input), $U2$ will increase its price above that before the merger and $U1$ - $D1$ will benefit from this price increase because $D2$'s increased input costs ultimately improve $U1$ - $D1$ profits – this is the raising-rival's-cost effect.

Researchers soon suggested that there was a problem with this argument. Hart and Tirole [1990] and Reiffen [1992] pointed out that, even though foreclosure would be a profitable strategy for the integrated firm, it still has an incentive to compete in the input market. The outcome OSS [1990] derive is therefore not a Nash equilibrium. To understand this argument, note that, given that the integrated firm withdraws from the input-good market and $U2$ becomes a monopolist supplier of the input for $D2$, the integrated firm has an incentive to deviate. Rather than stay out of the input market, it will compete in order to gain the business with $D2$. $U2$ will anticipate this and then the Nash equilibrium is the same both with and without vertical integration.

The problem with OSS [1990] is actually more subtle than the last paragraph (and the discussion in much of the literature) suggests. OSS [1990] assume that a vertical merger enables a firm to commit to not delivering to the downstream rival ($D2$ in Figure 1). When such commitment is available, there is no formal problem with the OSS [1990] argument and, indeed, the integrated firm will foreclose the market. In that case, the raising-rivals'-costs effect will result. However, Hart and Tirole's [1990] point is not whether commitment works but whether it will be available at all. They conclude that 'commitment is unlikely to be believable' [Hart and Tirole, 1990].

The availability of commitment also concerns OSS [1992], which is a reply to Reiffen's [1992] comment. OSS [1992] show that their results hold when upstream firms bid for a nonintegrated downstream firm in a descending-price auction. Yet, to some extent, this merely circumvents the problem that the original analysis posed. In the descending-price auction, the integrated firm will drop out early in the auction, the input market will be monopolized by the non-integrated upstream firm and, hence, the outcome is as in OSS [1990]. Note that a deviation from this equilibrium is prevented by the rules of the auction. By dropping out, the integrated firm commits itself not to compete. Whether the integrated firm can commit is still subject to debate.²

Despite this criticism, OSS' [1990] is a seminal paper in the theoretical Industrial Organization literature. It is the basis of the EU's recent *Non-Horizontal Merger Guidelines*.³ It is featured in various textbooks and it remains a fruitful agenda for theoretical research. For example, Linnemer [2003] simply assumes that a raising-rival's-cost effect of vertical integration

² Similar criticism could also be made of related approaches that attempt to rectify OSS' [1990] conclusion. Choi and Yi [2000] and Church and Gandal [2000] assume that upstream firms can commit to a technology which makes the input incompatible with the technology of nonintegrated downstream firms, in which case an outcome similar to the one proposed by OSS [1990] results. As with OSS [1990, 1992], the central assumption is that some form of commitment is available.

³ See the 'Guidelines on the Assessment of Non-Horizontal Mergers under the Council Regulation on the Control of Concentrations between Undertakings,' *Official Journal of the European Union*, 18.10.2008, C265/6-25.

exists, and he uses this as a base for further theoretical analysis (see also Buehler and Schmutzler [2005]).

The continued interest suggests that there may be more to the OSS' [1990] hypothesis than would come from a model that is plainly wrong. OSS [1992] themselves make such a claim in their reply to Reiffen [1992], and, perhaps somewhat surprisingly for mainstream theorists at the time, they use a behavioral argument when defending their position:

The notion that vertically integrated firms behave differently from nonintegrated ones in supplying inputs to downstream rivals would strike a business person, if not an economist, as common sense (OSS [1992]).

One interpretation of this quotation is that OSS [1992] suggest that their model may have predictive power even though their scenario cannot be supported in a Nash equilibrium. (It should be added that, as noted above, OSS [1992] propose the descending price auction where the foreclosure effect does occur in Nash equilibrium.)

The quotation from OSS [1992] may also suggest that there are actually two interpretations of the foreclosure notion. Foreclosure in a narrow sense can be said to occur when integrated firms refrain completely from supplying the input market. In OSS [1990], the integrated firms charge a price above the monopoly price U_2 would choose – a strategy which is equivalent to the exit of the integrated firm from the input-good market. A broader interpretation of the term would be that integrated firms 'behave differently' from nonintegrated firms, that is, presumably, charge higher prices, but they need not completely foreclose the input market. Broadly speaking, as long as vertical integration causes input prices to be higher than those in markets where firms are separated, foreclosure occurs. Rey and Tirole [2007] define foreclosure in the same broad sense.

This paper reports on an experimental analysis of the OSS [1990] argument. The experiments were designed to investigate whether vertical integration *per se* affects the behavior of integrated firms in the original OSS [1990] setup and does this without requiring any formal commitment. The experimental design allows us to study the effects (in otherwise identical markets) of vertical integration compared to non-integration.⁴ Even if the static Nash equilibrium does not predict an effect of vertical integration, experimental data may reveal whether vertical integration results in foreclosure in the broad sense or the narrow sense.

⁴The effect of vertical integration has been analyzed in experiments before. Mason and Phillips [2000] study a bilateral Cournot duopoly when there is a (large) competitive market that also demands the input from the upstream firms. Durham [2000] and Badasyan *et al.* [2009] compare integrated and nonintegrated monopolies and investigate whether vertical integration mitigates the double marginalization problem. Martin, Normann and Snyder [2001] analyze whether an upstream monopoly loses its monopoly power when selling a good to multiple retailers using two-part tariffs. None of these experiments have investigated the OSS [1990] hypothesis and the design of the experiments would not be suitable for doing so.

Another contribution the paper makes relates to repeated interaction. The arguments in OSS [1990, 1992], and in most of the theoretical literature, are based on the one-shot game. However, interaction in the field is often repeated. Studying a repeated setting seems particularly relevant here, since the commitment problem of the integrated firm may be resolved with repeated interaction. After all, repeated interaction (Macauley [1963]) can serve as an informal commitment device. Here, it may help the integrated firm to establish a reputation for foreclosure. Experiments with repeated firm interaction investigate this hypothesis. They are related to the recent theoretical literature that argues that vertical integration facilitates collusion (Nocke and White [2007], Normann [2009], Riordan and Salop [1995]).⁵

The design of the experiments in this paper follows the distinctions between vertical integration and separation and between the static and the repeated game. The first two treatments allow markets where both firms are vertically separated to be compared to markets where one firm is vertically integrated under a random-matching scheme, such that incentives are as in a one-shot game. Treatments three and four make the same comparison with a fixed-matching scheme in order to allow for repeated-game effects. This yields a two-by-two treatment design with vertical integration and the matching scheme as treatment variables.

The experimental setup is simplified as far as possible with the goal of providing a clean test of the commitment issue, which is at the heart of the debate around the OSS [1990] model. Only upstream behavior (the input market) is part of the experiment, because 'how input prices are set is a crucial determinant of the overall game' OSS [1990, p. 133]. A Bertrand duopoly experiment was designed to address the foreclosure issue. A downside of this simple design is that the experiment cannot fully address issues that may occur in a richer field environment. For example, downstream firms are not represented by subjects in the experiment and this may preclude effects that work in favor or against the foreclosure hypothesis.

There are three main experimental results. All three results hold with both random and fixed matching. Firstly, prices are significantly higher in markets where one firm is vertically integrated compared to those markets where both upstream firms are separated. Second, if one firm is integrated, the integrated firms' pricing behavior is less competitive than that of nonintegrated firms. Third, despite these anti-competitive effects, integrated firms only rarely completely foreclose the input market. That is, on the one hand, there is foreclosure in the broad sense of higher prices for the input,

⁵ It should be added that the experiments do not provide a formal test of these papers as the experimental design differs in various dimensions from the frameworks of the formal models.

but, on the other hand, there is almost no evidence of foreclosure in the narrow sense.

As these results violate the Nash equilibrium prediction, an explanation for the findings is needed. It turns out the results are consistent with a quantal response equilibrium analysis of the game (McKelvey and Palfrey [1995]). Quantal response equilibrium takes decision errors into account, so that players do not choose the best response with probability one but choose better choices more frequently. Quantal response equilibrium captures the fact that the vertical merger affects the payoff structure of the game. While this does not change the standard Nash prediction, it has an impact on the QRE outcomes.

Goeree and Holt [2001] have shown for several games that changes in the payoff structure that do not affect the Nash prediction can nevertheless have effects on the results of experiments.⁶ More closely related to this paper, Capra *et al.* [2002] analyze experiments with price-setting duopolies in which the unique Nash equilibrium is the Bertrand outcome. Competition, however, is not perfect because the market share of the high-price firm is larger than zero.⁷ The experimental results show that price levels are positively correlated with the market share of the high-price firm – which is a violation of the Nash prediction. The results in Capra *et al.* [2002], and in most of Goeree and Holt's [2001] examples, are well explained by quantal response equilibrium.

For the setting of this paper, quantal response equilibrium implies that integrated firms do indeed price less competitively than nonintegrated ones. Integrated firms still compete in the input market (that is, there is no foreclosure in the narrow sense), but the broad implication of the OSS model – that integration raises the price of the input – is consistent with the quantal response equilibrium generalization of Nash equilibrium.

II. EXPERIMENTAL DESIGN

To clarify the scope of the experiment, it seems useful to start with a recapitulation of the moves of OSS's [1990] foreclosure game as it perceived in the literature (refer to Figure 1 for firms' labels again).

1. Firms $U1$ and $D1$ decide whether to integrate.
2. Given the integration decision, $U1$ (or $U1-D1$) and $U2$ simultaneously decide about upstream prices.

⁶Goeree and Holt [2001] analyze ten simple one-shot games where the experimental data support the Nash equilibrium (the 'treasure' treatments). For all ten games, they find an 'intuitive contradiction' which results from a change in the payoff structure that leaves the Nash prediction unchanged but drastically alters the experimental results.

⁷Morgan, Orzen and Sefton [2006] also ran experiments with imperfect Bertrand competition. In their model, a firm that is not charging the lowest price still sells a positive amount due to brand-loyal consumers. Morgan, Orzen and Sefton [2006] analyze how the comparative statics predictions for changes in the degree of consumer loyalty are borne out in the data.

TABLE I
THE PAYOFF TABLE FOR THE TREATMENTS WITHOUT VERTICAL INTEGRATION

Price	1	2	3	4	5	6	7	8	9
Bertrand profit	39	54	69	81	90	99	90	72	51

3. Knowing upstream prices (and therefore knowing their input costs), $D1$ and $D2$ simultaneously set downstream prices.
4. Consumers make purchasing decisions given downstream prices.

The *second* stage of the model is at the heart of the OSS [1990, 1992] papers and the debate around them, and the experiments are about this second stage only. It is at this point where OSS [1990] assume that integration enables a firm to commit to a high price or to stay out of the market whereas the subsequent literature assumes that no commitment is available. The experiment follows the subsequent literature as pricing decisions are simultaneous moves without commitment.

The experiment abstracts from the other three stages in order to be as simple as possible. The first stage is exogenously given in the experiment. There are experimental treatments with and without vertical integration, and integration is not a choice for the subjects. Neither the third nor the fourth stages are present in the experiments, that is, downstream firms and final-good consumers were not represented by participants in the experiments. Instead, they are assumed to play according to Nash equilibrium. Downstream firms' and consumers' equilibrium behavior implies the payoffs tables that were given to the subjects.

How was this decision problem of the second stage implemented in the experiments? In all treatments, two subjects representing the two upstream firms have to make one single choice in every period, they simultaneously set an (upstream) price which has to be an integer between one to nine. The treatments (integration or separation) differ in the payoff table that is given to the subjects. These payoff tables are derived from a parametrized model (see Appendix). The tables are fully consistent with the analysis in OSS [1990] who indeed use the same parametrized model for some of their analysis.

In the treatments with vertical separation, the basis for the decision making is the profit table reproduced in Table I. Here, participants play a Bertrand duopoly experiment (similar to the first price-setting oligopoly experiments, conducted by Fouraker and Siegel [1963]). The firm that charges the strictly lowest price will gain the profit in the 'Bertrand profit' cells.⁸ In the case of a tie, both firms get half that profit. In the instructions,

⁸ In the instructions of the experiments, neutral labels were used instead of the terms 'Bertrand' and 'integrated firm' (instructions are available from the author upon request).

TABLE II
THE PAYOFF TABLE FOR THE TREATMENTS WITH VERTICAL INTEGRATION

Price	1	2	3	4	5	6	7	8	9
Bertrand profit (both firms)	39	54	69	81	90	99	90	72	51
Additional profit (integrated firms only)	66	74	84	96	105	132	159	180	198

this was illustrated with two examples, one of which read 'If you charge a price of 7 and the other firm charges a price of 4, you will get zero and the other firm gets 81 pence'.

In the treatments with vertical integration, the two participants play the same Bertrand game⁹ but the twist is that one subject (firm 1, the integrated firm) now makes an extra profit. In these treatments, the basis of decision making is the profit table reproduced in Table II. In addition to the 'Bertrand profits' (with the rules of the game as above), there is the 'additional profit' in Table II. This extra payoff represents the profit the downstream affiliate of the integrated firm earns. This row of the payoff table thus applies to integrated firms only. Consistent with the theory of OSS [1990], the higher the price in the Bertrand game, the more 'additional profit' the integrated firm earns – this is the raising-rival's-costs effect. Thus, as it was put in the instructions, 'it is the lowest of the two prices that determines the ['additional'] profit, no matter whether firm 1 [the integrated firm] or firm 2 (or both) charged the lowest price'. One of two illustrative examples in the instructions reads: 'If firm 1 charges a price of 7 and firm 2 charges a price of 4, firm 1 only gets 96 pence' additional profit.

Note, once more, the commitment problem. Ideally, the integrated firm would want to commit to a price of 7 or higher, because the nonintegrated firm would then best respond by setting the monopoly price of 6. In that case, profits would be 132 for the integrated firm and 99 for the nonintegrated firm. However, this foreclosure strategy is not feasible without commitment, as the integrated firm can obtain $90 + 105 > 132$ by deviating to a price of 5. Thus, absent commitment, vertical integration may not make any difference at all (Hart and Tirole [1990], Reiffen [1992]).

The experimental markets were designed such that firms still make a positive profit when they both charge 1 (the Nash equilibrium price, derived below). The reason is that subjects might be biased against an action with zero profit (Dufwenberg *et al.* [2007]). Consistent with the underlying

⁹ One general implication of vertical integration is that the downstream unit of the integrated firm no longer buys the input from the market but instead obtains it at marginal cost internally. This implies that the input market has a bigger volume with vertical separation (twice as big) and thus the 'Bertrand profits' in Table I should also be bigger without integration. However, as the experimental design needs to avoid possible wealth effects, the 'Bertrand profits' are kept equal across treatments (see also Appendix).

TABLE III
THE TREATMENTS

		Matching	
		random	fixed
vertical	separation	<i>SEPAR_RAND</i>	<i>SEPAR_FIX</i>
	integration	<i>INTEG_RAND</i>	<i>INTEG_FIX</i>

theoretical model, the ‘additional profit’ is larger than the one in the ‘Bertrand’ row. This means that the integrated firm makes a larger profit than the nonintegrated firm even if it does not get any profit in the Bertrand game.

The two treatment variables are the vertical structure and the matching scheme. The treatments with and without vertical integration are labeled *INTEG_* and *SEPAR_*, respectively. Treatments where participants were randomly rematched in every period have the label *_RAND*, and treatments where subjects repeatedly interacted in pairs of two (fixed matching) are labeled *_FIX*. Table III summarizes the 2×2 treatment design.

The treatments in Table III had a length of 15 periods. As a robustness check, additional sessions with a length of 25 periods were conducted for the treatments with random matching. Below, these treatments are referred to as *SEPAR_RAND25* and *INTEG_RAND25*. In all treatments, subjects knew the number of periods and the end period from the instructions.

III. PREDICTIONS

The subgame perfect Nash equilibrium prediction is the same for all treatments but it seems worthwhile to go through the four variants in detail separately.

In *SEPAR_RAND*, both firms charge the lowest price of 1 in equilibrium (this is the standard Bertrand-Nash equilibrium). Equilibria where firms charge a higher price do not exist because firms have a strict incentive to undercut at any price larger than 1. Both firms earn $39/2 = 19.5$ in equilibrium.

In *INTEG_RAND*, the integrated firm would want to commit to a price larger than 6 but this is not a Nash equilibrium. As emphasized by Hart and Tirole [1990] and Reiffen [1992], the unique Nash equilibrium has both firms choosing 1 also with vertical integration. In the equilibrium of this treatment, the nonintegrated firm earns 19.5 and the integrated firm earns $19.5 + 66 = 85.5$.

In *SEPAR_FIX* there is repeated interaction and it is well known that some collusion may occur. If so, the price of 6 would maximize joint profits. In any event, the subgame perfect Nash equilibrium is for both firms to

charge the price 1 just as in the *RAND* treatments, as follows from backward induction in the finitely repeated game.

In *INTEG_FIX*, firms may collude by charging the same price. In that case, any price between 6 and 9 is Pareto efficient (from the firms' point of view). Vertical integration may also allow for another form of collusion where firms charge different prices and collude by coordinating on foreclosure (in the narrow sense). The integrated firms could set a price larger than 6 and the nonintegrated firm could set a price of, e.g., 6. However, neither way of colluding is a subgame perfect Nash equilibrium, as follows from backward induction, and both firms choosing 1 is the subgame perfect Nash prediction once again.

IV. PROCEDURES

All treatments were run in sessions with 10 participants. Five participants acted as 'firm 1' (the integrated firm in *INTEG* treatments) and the other five participants acted as 'firm 2.' These roles were fixed for the entire course of the experiment. In the *SEPAR* treatments, firms are symmetric but the 'firm 1' – 'firm 2' labels were nevertheless given in order to keep matching scheme and instructions comparable.

There were four sessions (each with ten participants) for treatment *SEPAR_RAND* and also four for *INTEG_RAND*. There was one session each for treatments *SEPAR_FIX* and *INTEG_FIX*. Having more sessions with random matching is motivated by the possibility of group effects within sessions under random matching.

Experiments were computerized, with the programming done in *z-Tree*, developed by Fischbacher [2007]. The treatments listed in Table III were conducted at Royal Holloway College, University of London, in autumn 2004 and spring 2005. The payoffs in Tables 1 and 2 denote cash payments British pence. Subjects' average monetary earnings were £ 12.50, including a flat payment of £ 5 in the London sessions. In total, 140 subjects participated (100 in the main treatments plus 40 in the sessions with 25 periods, reported below in Section 6). Subjects were mainly undergraduate students and a large proportion of them were from faculties other than economics or business studies.

V. RESULTS

Table IV and Figures 2 and 3 summarize the results.¹⁰ The averages in Table IV and most formal tests are based on data from periods 6 to 15. All results reported also hold qualitatively if the analysis is based on all periods or on

¹⁰ Based on the main treatments with 15 periods length. The results from the treatments with 25 periods length are reported below.

TABLE IV
 AVERAGE PRICES (BASED ON SESSION AND GROUP AVERAGES, STANDARD DEVIATION IN PARENTHESIS) AND (ONE-SIDED) p -VALUES OF MANN-WHITNEY U TESTS FOR DIFFERENCES BETWEEN THE PRICE DISTRIBUTIONS OF THE TREATMENTS

		vertical structure		
		<i>SEPAR</i>	<i>INTEG</i>	
matching	<i>RAND</i>	1.81 (0.43)	2.83 (0.77)	$p = 0.021$
	<i>FIX</i>	2.67 (1.15)	4.40 (1.02)	$p = 0.014$
		$p = 0.115$	$p = 0.033$	

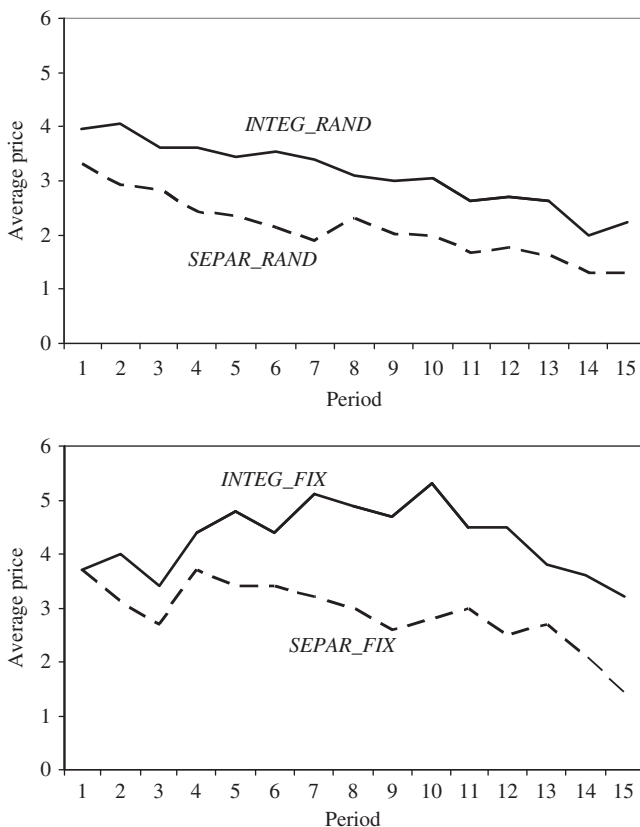


Figure 2

Average Prices in *SEPAR* (dashed lines) and *INTEG* (solid lines) for Random Matching (top panel) and Fixed Matching (bottom panel)

periods 11 to 15. There are four entirely independent observations for the *RAND* treatments and five independent observations for the *FIX* treatments. The non-parametric tests applied here (in this case,

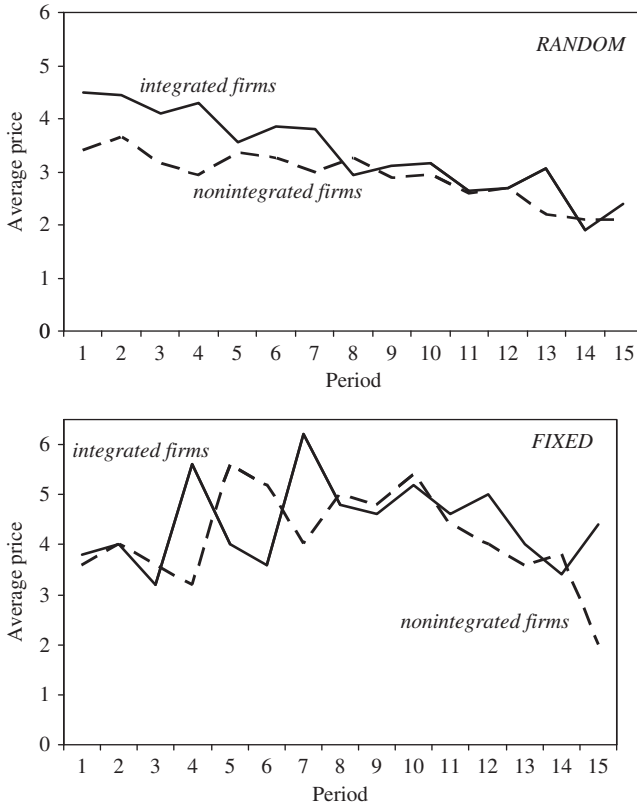


Figure 3
Average Prices of Integrated Firms (solid lines) and Nonintegrated Firms (dashed lines) in the INTEG Treatments, Random Matching (top panel) and Fixed Matching (bottom panel)

Mann-Whitney *U* tests) conservatively only use these four and five observations, respectively.¹¹

Prices in *SEPAR_RANDOM* are lower than those in *INTEG_RANDOM*. The top panel of Figure 2 confirms this for the average prices in *SEPAR_RANDOM* and *INTEG_RANDOM* across the 15 periods and Table IV shows that prices are about 36% higher with vertical integration. The significance of this result follows from a Mann-Whitney *U* test ($p = 0.021$, see also Table IV).

The top panel of Figure 3 indicates that there are some differences between the prices of integrated and nonintegrated firms in *INTEG_RANDOM*.

¹¹ Nonparametric tests are distribution-free tests that do not rely on assumptions regarding the distribution the data are drawn from (e.g., normal distribution). The tests work with cardinal rankings of the observations rather than ordinally measured data. See, for example, Hollander and Wolfe [1999].

Averaging across periods 6 to 15, integrated firms' prices are about 8% higher with random matching. These differences are significant (one-sided matched-pairs Wilcoxon, $p = 0.034$) although quantitatively perhaps not particularly big.

Essentially the same results also hold in the *FIX* treatments. Comparing *SEPAR_FIX* and *INTEG_FIX*, Figure 2 and Table IV indicate differences due to integration which are quantitatively bigger than with random matching. The relative increase is roughly the same, however, as prices are about 39% higher with integration in *FIX*. These differences are significant ($p = 0.014$). As with random matching, integrated firms in *INTEG_FIX* charge 8% higher prices than nonintegrated firms (significant according to a matched-pairs Wilcoxon, $p = 0.039$). (See below for an analysis of whether this results holds with a longer horizon.)

Result 1: There is evidence of foreclosure broadly defined. Markets with a vertically integrated firm have significantly higher prices than markets where the two firms are separated. In markets with integration, the vertically integrated firms charge significantly higher prices than the nonintegrated firms.

How about foreclosure in the narrow sense then? Strong evidence in favor of that would be if integrated firms charged a price higher than 6, as this would imply a complete withdrawal from the input market.

It is already clear from Figures 2 and 3 and the averages in Table IV that there is only little evidence of such behavior, and a concrete search for these foreclosure outcomes confirms that they are rare. In *INTEG_RAND*, only one of 20 subjects representing an integrated firm charged prices which deviated from the general pattern visible in Figure 2. This subject charged prices of 7 and 8 from period 2 to 13 and clearly did not compete in the input market except for the last two periods. This can be interpreted as foreclosure behavior, narrowly defined. In total, however, only 29 of 300 observations (data from all periods) include prices of 7 or higher, and 12 of these cases are accounted for by the subject just mentioned. For comparison, nonintegrated firms in *INTEG_RAND* charged a price of 7 or higher in 4 (of 300) cases, and in *SEPAR_RAND* there were 11 (of 600) such observations. Hence, whereas these shares are somewhat lower than those of the integrated firms in *INTEG_RAND*, too few observations in *INTEG_RAND* are consistent with (narrow) foreclosure to suggest that it is important in the data.

In treatment *INTEG_FIX*, integrated firms charged prices of 7 or higher in 5 of 75 cases. Compared to this, nonintegrated firms did so in 2 of 75 cases, and in *SEPAR_FIX* there are 5 (of 150) such cases (data from all periods). Hence, *INTEG_FIX* does not contain more evidence of foreclosure in the narrow sense than *INTEG_RAND* does.

TABLE V
 NUMBER OF OBSERVATIONS WHEN THE INTEGRATED FIRM OR THE NONINTEGRATED FIRM CHARGED THE LOWEST PRICE, AND NUMBER OF TIES

Low-price firm	Treatment	
	<i>INTEG_RAND</i>	<i>INTEG_FIX</i>
integrated	32%	22%
nonintegrated	39%	40%
Ties	29%	38%
# observations	200	50

Looking at the five individual duopoly pairs in *INTEG_FIX* yields further insights. Duopoly #1 had both firms charging a collusive price of 6 in all periods except for the first two and the last three. Duopoly #2 priced competitively in the first and last third of the experiment and only in two outcomes in the middle of the experiment did the integrated firm charge high prices. Duopoly #3 colluded non-systematically. Sometimes, there was symmetric collusion, sometimes there were apparently competitive outcomes with either firm being the low-price firm. In duopoly #4, the integrated firm never charged a price lower than the rival, possibly suggesting a foreclosure strategy – but then, why did this firm not go all the way and set a price above 6? Finally, duopoly #5 started competitively and then colluded symmetrically at a price of 6. To summarize, there is not more evidence of foreclosure with fixed matching either. If firms collude successfully at all, they both tend to choose a price of 6 rather than have the integrated firm foreclose the input market.

Even if integrated firms do not charge prices higher than 6, it could still be that they do not compete and that they only rarely charge a lower price than nonintegrated firms as a result. Table V shows data on which type of firm turns out to be the low-price firm in the *INTEG* treatments (in periods 6–15). The table also lists the number of ties. Integrated firms charge the lowest price less frequently than nonintegrated firms both in *INTEG_RAND* and *INTEG_FIX*. These differences are significant according to binomial tests in *INTEG_RAND* ($p = 0.0625$, one sided) and *INTEG_FIX* ($p = 0.03125$, one sided).¹² However, while consistent with Result 1, they are quantitatively too minor to support the narrow foreclosure hypothesis.

Result 2: There is little evidence of foreclosure narrowly defined. Even though integrated firms are the low-price firm significantly less often, they still compete actively in the input-good market.

¹² Integrated firms charge the lowest price less frequently than nonintegrated firms in all four (five) groups of *INTEG_RAND* (*INTEG_FIX*). Accordingly, a binomial test rejects the null hypotheses of equal likelihood at a significance level of $p = (\frac{1}{2})^n$, where n is the number of observations.

VI. ROBUSTNESS CHECK: GAMES WITH 25 PERIODS

Both Figure 2 and Figure 3 indicate a negative time trend in the data. In *INTEG_FIX*, prices are stable except for an end-game effect in the last three periods. Regarding *INTEG_RAND* and *SEPAR_RAND*, however, it could be that both treatments converge to the Nash equilibrium price of 1 only at different rates. In that case, Result 1 would not be robust.

To check for the impact of the length of the horizon, additional sessions with a length of 25 periods were carried out in the spring of 2010 at the University of Düsseldorf. Specifically, there were 40 subjects participating in two sessions each for treatments *SEPAR_RAND25* and *INTEG_RAND25*.

The data from treatments *INTEG_RAND25* and *SEPAR_RAND25* indicate that with a longer horizon the above results remain.¹³ In periods 16 to 25, average prices in *INTEG_RAND25* are 2.51, and they are 1.46 in *SEPAR_RAND*. Importantly, there are no significant time trends in this phase of the experiments any more with Spearman correlations coefficients (of prices and time) being $\rho = -0.105$ ($p > 0.1$) in *INTEG_RAND25* and $\rho = 0.085$ ($p > 0.1$) in *SEPAR_RAND25*. (Average prices in periods 6 to 15 suggest more pronounced differences than in the main treatments with average prices of 1.43 in *SEPAR_RAND25* and 3.06 in *INTEG_RAND25*). Thus, these averages suggest that the differences persist.¹⁴ On the other hand, the differences are arguably quantitatively weak, and it is not clear how the effect of a vertical merger would carry over to a more complex setting in which other effects may confound the foreclosure effect.

VII. QUANTAL RESPONSE EQUILIBRIUM ANALYSIS

The results suggest that vertical integration has an impact. We saw that integrated firms charge significantly higher prices. This effect was sufficient to render less competitive the markets where vertical integration is present. While this confirms the OSS [1990] hypothesis in a broad sense, the lack of evidence for foreclosure rejects the narrow OSS prediction. How can these results be accounted for?

¹³Note that extending the length of the experiment is not without cost. The longer the horizon, the more the game may have aspects of a repeated game, despite the random matching scheme. (Subjects may simply recognize that they interact 'many' times with the individual members of the group, even if each interaction is randomly determined). There is a tradeoff between the goal of a longer horizon and avoiding repeated-game effects.

¹⁴Because only two (randomly rematched) sessions were conducted for each treatment of this robustness check, one cannot conduct the non-parametric tests applied above. However, the significance of the result in periods 16 to 25 can be explored by running *t*-tests on average prices. Counting each individual as one observation, one can establish a lower bound for the significance value at $p < 0.001$ ($t = -4.40$). With just one average price for each session, the result is weakly significant at $p = 0.051$ (one sided, $t = -1.85$).

In this section, it will be argued that the above findings are consistent with the quantal-response equilibrium (QRE) analysis (McKelvey and Palfrey [1995]) of the game. QRE is a generalization of Nash equilibrium that takes decision errors into account. Players do not always choose the best response with probability one but they do choose better choices more frequently. Because of this, changes in the payoff structure that do not affect the standard Nash prediction can still have an impact on the QRE outcome(s). In the model of this paper, vertical integration (compared to nonintegration) has exactly this impact. Therefore, QRE is a good candidate for explaining the results.

Consider the logit equilibrium variant of QRE. Firm i , $i = 1, 2$, believes that the other firm will choose price p_k , $k \in \{1, 2, \dots, 9\}$, with probability ρ_i^k . Accordingly, firm i 's expected profit from choosing price j is

$$\Pi_i^j = \sum_{k=1}^9 \rho_i^k \pi_i(p_j, p_k), j = 1, \dots, 9.$$

where $\pi_i(p_j, p_k)$ are the profits as in the Bertrand game of Table I (note that profit functions are not symmetric in the *INTEG* treatments). As mentioned above, firms choose better choices more frequently. In particular, choice probabilities, σ_i^j , are specified to be ratios of exponential functions

$$\sigma_i^j = \frac{e^{\lambda \Pi_i^j}}{\sum_{k=1}^9 e^{\lambda \Pi_i^k}}, j = 1, \dots, 9.$$

λ is the error parameter. If $\lambda = 0$ behavior is completely noisy and all prices are equally likely regardless of their expected profit. As $\lambda \rightarrow \infty$, firms choose the best response with probability one. In the logit equilibrium, beliefs and choice probabilities have to be correct, that is, $\rho_1^j = \sigma_2^j$ and $\rho_2^j = \sigma_1^j$, $j = 1, \dots, 9$.

Using *Gambit* (McKelvey *et al.* [2005]), a unique equilibrium (given λ) is found, as illustrated in Figure 4. The figure shows the relative frequency of the price of 1 in QRE conditional on the error parameter, λ . The impact of λ is intuitive. If $\lambda = 0$, the price of 1 (like any other price) is chosen with probability 1/9 and, as $\lambda \rightarrow \infty$, it is chosen with probability one, as in the standard Nash equilibrium. When $\lambda \in (0, \infty)$, vertical integration implies (among other things) that the frequency of the price 1 is lower than in markets with separation. Only when $\lambda = 0$ and $\lambda \rightarrow \infty$, does vertical integration not have any impact. The figure illustrates these findings for a relevant range of lambda. It also shows that the integrated firms in *INTEG* set the price of 1 less often than nonintegrated firms.¹⁵

¹⁵ This can be generalized. It turns out that, for any $\lambda \in (0, \infty)$, the distribution of the prices of the nonintegrated firm first-order stochastically dominates that of integrated firms in the *INTEG* setup. This supports the claim that integrated firms are less competitive than their

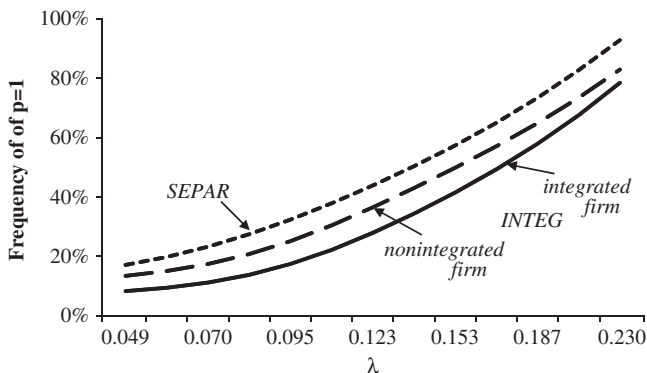


Figure 4

Quantal Response Equilibrium Simulations for the Frequency of the Price of 1

To summarize up to this point, whereas neither the static Nash equilibrium nor the foreclosure outcome organize the data in the *RANDOM* treatments well, QRE does. The qualitative predictions of QRE are confirmed. In particular, the QRE analysis is consistent with the finding that, although integrated firms charge somewhat higher prices, they do not completely refrain from competing in the input market.

What is the intuition behind the mechanics of the QRE analysis, and why can it explain treatment differences? As mentioned above, players do not play a best reply with certainty in QRE and prices higher than the Nash equilibrium price are played with positive probability – however, this applies to both treatments with and without integration. The key difference vertical integration makes is that the integrated firm loses less profit when the rival undercuts its price. Therefore, integrated firms charge higher prices (in a probabilistic sense) and this also pushes up the prices of non-integrated firms in equilibrium (where expectations are correct), rendering the *INTEG* treatments less competitive. Put in terms of the foreclosure story, an integrated firm still has an incentive to compete (which confirms Hart and Tirole [1990] and Reiffen [1992]) but that this incentive is weaker than for a nonintegrated firm (which confirms OSS' broad foreclosure interpretation).

The next step is to estimate λ Using data from *INTEG_RANDOM* and *SEPAR_RANDOM* jointly,¹⁶ Maximum-Likelihood estimates of the error term

nonintegrated counterparts. However, it does not generally follow that prices will be lower in *SEPAR* compared to *INTEG* (in the sense of first-order stochastic dominance).

¹⁶Note this yields one Maximum Likelihood estimates of λ for two treatments (*INTEG_RANDOM* and *SEPAR_RANDOM*). Recently, Haile *et al.* [2008, p. 188] have criticized that '[a]lthough many papers have examined the fit of the logit QRE in different treatments (varying payoffs), typically a new value of the logit parameter is estimated each time.' Here, a single estimate is conducted only, and it can rationalize the observed difference between the two treatments. Haile *et al.* [2008, p. 188] explicitly do not criticize this. Note also that QRE

are as follows. For periods 6 to 15 (which the results in the above section were based on), $\lambda = 0.134$ results with a standard error of 0.006. Intuitively, the λ estimates increase over time. In periods 1 to 5, $\lambda = 0.051$ (0.012) results; whereas, in periods 11 to 15, $\lambda = 0.182$ (0.009) is the estimate. This is consistent with the decline in prices observed in the *RAND* treatments and the notion that subjects learn over time and become 'more rational'.

How well do QRE and the actual estimate of λ fit with the differences observed between the two treatments (*INTEG_RAND* and *SEPAR_RAND*)?

- The predicted (expected) average QRE price in *SEPAR_RAND* is 2.13, and the actual average turns out to be 1.81. In *INTEG_RAND*, predicted average prices are 2.89 and 2.23 for the integrated and nonintegrated firm, respectively. Observed average prices of 2.96 and 2.71.
- The expected winning price (the minimum of the two prices) is predicted to be 1.30 in treatment *SEPAR_RAND* and the actual average is exactly 1.30. In *INTEG_RAND* the prediction is 1.57 and the average is 2.17.

Whereas the broad magnitude of expected values corresponds to the actual values, it appears that the QRE prediction (given $\lambda = 0.134$) somewhat overestimates the differences between integrated and nonintegrated firms and underestimates the differences between the treatments' averages.

Note that QRE implies a distribution for the prices firms charge that differs between the treatments, given the estimate $\lambda = 0.134$. (By contrast, if firms merely made decision errors, a uniform distribution across prices would result in both treatments.) In *SEPAR_RAND*, the QRE predicted and observed frequency of prices are as follows. Price of 1: 50% (QRE) and 55% (data); price of 2: 20% and 29%; prices larger than 2: 16% and 30%. The same numbers for *INTEG_RAND* are like this. Price of 1: 32% (QRE) and 24% (data); price of 2: 28% and 27%; prices larger than 2: 40% and 50%. It appears that the QRE predicted and the observed distribution do not differ substantially. But, more importantly QRE, seems to capture the treatment differences rather well on basis of the same estimate for λ .

VIII. CONCLUSION

This paper contributes to the literature on vertical integration and raising rivals' cost with the use of a laboratory experiment. The experiments were designed to analyze the raising-rivals'-costs argument of Ordober, Saloner and Salop [1990]. In simple duopoly treatments (with random and fixed matching), the data show how the presence of an integrated firm affects market outcomes.

arguments generally have less bite in repeated-game settings because there is less uncertainty about the action of the other player. Therefore, the estimations are based in the data from the *_RAND* treatments.

The experimental results support the hypothesis of Ordober, Saloner and Salop [1990] in that overall competition is reduced when one firm vertically integrates and, in markets where an integrated firm is present, it charges higher prices compared to nonintegrated firms. While the effects are quantitatively small, they are statistically significant. On the other hand, there is very little evidence of foreclosure in the sense that virtually no integrated firm completely refrains from competing in the input market. Whereas these results are inconsistent with the standard notion of Nash equilibrium, these results are consistent with the quantal response equilibrium (McKelvey and Palfrey [1995]) generalization of Nash equilibrium. The results are also consistent with Ordober, Saloner and Salop's [1992] broad notion of foreclosure which says that vertical integration generally causes an anticompetitive effect even if no refusal to supply the input market is observed.

The lack of evidence for foreclosure (narrowly defined) suggests that the commitment problem of the integrated firm pointed out by Hart and Tirole [1990] and Reiffen [1992] is significant. In experiments, participants do generally not manage to resolve commitment problems simply with mere intentions. This has been found in Huck and Müller [2000], Reynolds [2000], Cason and Sharma [2001] and Martin, Normann and Snyder [2001].¹⁷ In these experiments, subjects failed to achieve desirable outcomes when there was no formal commitment mechanism, and the same appears to be the case in this study.

Further investigating the commitment issue also seems promising for future research. For example, will firms commit if they are given the opportunity to do so? Likewise, will firms learn to commit if they are forced to do so over a transitory period? Further, one could imagine effects arising in a more complex environment that seem worth investigating. When both upstream and downstream competition are part of the experiment, how will this affect the firm that explicitly makes both upstream and downstream decisions, and how will integrating previously separate firms influence pricing behavior?

APPENDIX THE MODEL

This appendix presents the model underlying the payoff table of the experiment. The model has two upstream firms ($U1$ and $U2$) which are Bertrand competitors and two

¹⁷Huck and Müller's [2000] experiments show that a Stackelberg leader has serious difficulties exploiting the first-mover advantage when second movers obtain a noisy signal of its action. Reynolds [2000] and Cason and Sharma [2001] show that monopolies producing durable goods often fail to achieve full monopoly profits. Similarly, in the experiments of Martin, Normann and Snyder [2001], a firm loses its monopoly power when selling its product through multiple retailers.

downstream firms ($D1$ and $D2$) which transform the input into differentiated final goods.

We begin at the downstream level. Downstream firm D_i 's demand is

$$(1) \quad q_i(p_i, p_j) = a - bp_i + dp_j; i, j = 1, 2; i \neq j,$$

where p_i and p_j are the prices the downstream firms i and j set ($i, j = 1, 2; i \neq j$). Suppose downstream firm i purchases the input good at a linear price of c_i per unit. As the D firms incur no other costs, they operate at constant marginal costs of c_1 and c_2 respectively. Thus, at the downstream level, this is a standard Bertrand duopoly model with product differentiation and asymmetric cost. It is straightforward to solve for downstream Nash equilibrium prices

$$(2) \quad p_i^*(c_i, c_j) = \frac{(2b + d)a + 2b^2c_i + bdc_j}{4b^2 - d^2},$$

outputs

$$(3) \quad q_i^*(c_i, c_j) = b \frac{(2b + d)a - (2b^2 - d^2)c_i + bdc_j}{4b^2 - d^2},$$

and profits $\pi_{Di}^* = (q_i^*(c_i, c_j))^2/b$.

Upstream firms have constant marginal costs which are assumed to be zero for simplicity. The upstream firms compete for each of the two downstream markets in a Bertrand fashion. Specifically, upstream firm k sets two prices, c_1^{Uk} and c_2^{Uk} , for downstream firms 1 and 2, respectively. The Bertrand logic implies that the downstream firm i buys from the upstream firm with the lowest price, formally $c_i = \min\{c_i^{U1}, c_i^{U2}\}, i = 1, 2$. Put it another way, an upstream firm will sell a positive amount to Di only if it charges the lowest price. Formally, when upstream firm k bids c_i^{Uk} to downstream firm I , it will make the following profit with Di

$$(4) \quad \pi_i^{Uk}(c_i^{Uk}, c_i^{Ul}) = \begin{cases} c_i^{Uk} q_i^*(c_i^{Uk}, c_i^{Uk}) & \text{if } c_i^{Uk} < c_i^{Ul} \\ c_i^{Uk} q_i^*(c_i^{Uk}, c_i^{Uk})/2 & \text{if } c_i^{Uk} = c_i^{Ul} \\ 0 & \text{if } c_i^{Uk} > c_i^{Ul} \end{cases}$$

where $k, l = 1, 2; k \neq l; i = 1, 2$.

When neither firm is integrated, in the general model, $U1$ and $U2$ set prices (c_1^{U1}, c_2^{U1}) and (c_1^{U2}, c_2^{U2}), respectively. For the derivation of the payoff tables, c_1 is set equal to zero. The reason is that $c_1 = 0$ with vertical integration. Thus, in order to keep treatments comparable and avoid wealth effects, one also needs $c_1 = 0$ without integration. Essentially, this implies that firms only compete for $D2$ also absent integration. This is without loss of generality as the qualitative features of Bertrand competition are unaffected by this. $D2$ buys at the lower of the two prices such that $c_2 = \min\{c_2^{U1}, c_2^{U2}\}$. Next, given $c_1 = 0$ and c_2 , $D1$ and $D2$ set the final good prices. In equilibrium, $D1$ and $D2$ charge $p_1^*(0, c_2)$ and $p_2^*(c_2, 0)$, respectively. Downstream profits are $\pi_{D1}^* = (q_1^*(0, c_2))^2/b$ and $\pi_{D2}^* = (q_2^*(c_2, 0))^2/b$, and upstream profits are π_2^{U1}, π_2^{U2} and $\pi_1^{Uk} = 0$.

A vertical merger of $U1$ and $D1$ implies that the integrated firm's true input price is $U1$'s marginal cost (Bonanno and Vickers [1988]). Thus, $D1$ will be delivered efficiently at $c_1 = 0$ and $U2$ cannot compete for the $D1$ business any more. For both upstream firms, only the $D2$ market remains a source for potential business. Profits are as follows.

$D2$ earns $\pi_{D2}^* = (q_2^*(c_2, 0))^2/b$, $U2$ earns π_2^{U2} , and the integrated firm $U1 - D1$ makes a profit of $\pi_2^{U1} + \pi_{D1}^* = \pi_2^{U1} + (q_1^*(0, c_2))^2/b$.

Tables 2 and 3 can be derived from these closed-form solutions for the parameters $a = 35/2$, $b = 4$, $d = 2$. The actual price parameters used to derive the profits in the table differ from the prices labels '1' to '9'. In particular, profits around the joint-profit maximizing prices are quite flat. Hence, prices were increased in steps larger than one in this range to avoid the 'flat-maximum' critique (Harrison [1989]). The actual price parameters underlying the values in the table are {1.1, 1.6, 2.2, 2.9, 3.5, 5.0, 6.5, 7.6, 8.5}. Additionally, profits were multiplied by three and rounded to yield the payoff in real currency subjects received.

REFERENCES

- Badasyan, N.; Goeree, J. K.; Hartmann, M.; Morgan, J.; Rosenblat, T.; Servatka, M. and Yandell, D., 2009, 'Vertical Integration of Successive Monopolists: A Classroom Experiment,' *Perspectives on Economic Education Research*, 5(1).
- Bonanno, G. and Vickers, J., 1988, 'Vertical Separation,' *Journal of Industrial Economics*, 36(3), pp. 257–265.
- Buehler, S. and Schmutzler, A., 2005, 'Asymmetric Vertical Integration,' *Advances in Theoretical Economics*, 5, Article 1.
- Capra, M.; Gomez, R.; Goeree, J. and Holt, C. H., 2002, 'Learning and Noisy Equilibrium Behavior in an Experimental Study of Imperfect Price Competition,' *International Economic Review*, 43, pp. 613–636.
- Cason, T. and Sharma, T., 2001, 'Durable Goods, Coasian Dynamics and Uncertainty: Theory and Experiments,' *Journal of Political Economy*, 109, pp. 1311–1354.
- Choi, J. P. and Yi, S.-S., 2000, 'Vertical Foreclosure and the Choice of Input Specifications,' *RAND Journal of Economics*, 31, pp. 717–743.
- Church, J. and Gandal, N., 2000, 'Systems Competition, Vertical Merger and Foreclosure,' *Journal of Economics and Management Strategy*, 9, pp. 25–51.
- Dufwenberg, M.; Goeree, J.; Gneezy, U. and Nagel, R., 2007, 'Price Floors and Competition,' *Economic Theory*, 33(1), pp. 211–224.
- Durham, Y., 2000, 'An Experimental Examination of Double Marginalization and Vertical Relationships,' *Journal of Economic Behavior and Organization*, 42, pp. 207–229.
- Fischbacher, U., 2007, 'Z-Tree: Zurich Toolbox for Readymade Economic Experiments,' *Experimental Economics*, 10(2), pp. 171–178.
- Fouraker, L. and Siegel, S., 1963, *Bargaining Behavior* (McGraw-Hill, New York, New York, U.S.A.).
- Goeree, J. and Holt, C. H., 2001, 'Ten Little Treasures of Game Theory and Ten Intuitive Contradictions,' *American Economic Review*, 91, pp. 1402–1422.
- Haile, P. A.; Hortaçsu, A. and Kosenok, G., 2008, 'On the Empirical Content of Quantal Response Equilibrium,' *American Economic Review*, 98(1), 180–200.
- Harrison, G., 1989, 'Theory and Misbehavior of First-Price Auctions,' *American Economic Review*, 79, pp. 749–762.
- Hart, O. and Tirole, J., 1990, 'Vertical Integration and Market Foreclosure,' *Brookings Papers on Economic Activity: Microeconomics*, pp. 205–276.
- Hollander, M. and Wolfe, D. A., 1999, *Nonparametric Statistical Methods* (John Wiley and Sons, New York, New York, U.S.A.).
- Huck, S. and Müller, W., 2000, 'Perfect vs. Imperfect Observability: An Experimental Test of Bagwell's Result,' *Games and Economic Behavior*, 31, pp. 174–190.

- Linnemer, L., 2003, 'Backward Integration by a Dominant Firm,' *Journal of Economics and Management Strategy*, 12, pp. 231–259.
- Martin, S.; Normann, H. T. and Snyder, C., 2001, 'Vertical Foreclosure in Experimental Markets,' *RAND Journal of Economics*, 32, pp. 667–685.
- Mason, C. F. and Phillips, O. R., 2000, 'Vertical Integration and Collusive Incentives,' *International Journal of Industrial Organization*, 18, pp. 471–496.
- Macauley, S., 1963, 'Non-Contractual Relations in Business,' *American Sociological Review*, 28, pp. 55–70.
- McKelvey, R. D. and Palfrey, T., 1995, 'Quantal Response Equilibria in Normal Form Games,' *Games and Economic Behavior*, 10, pp. 6–38.
- McKelvey, R. D.; McLennan, A. M. and Turocy, T. L., 2005, 'Gambit: Software Tools for Game Theory', Version 0.2005.06.13, <http://econweb.tamu.edu/gambit>.
- Morgan, J.; Orzen, H. and Sefton, M., 2006, 'An Experimental Study of Price Dispersion,' *Games and Economic Behavior*, 54(1), pp. 134–158.
- Nocke, V. and White, L., 2007, 'Do Vertical Mergers Facilitate Upstream Collusion?,' *American Economic Review*, 97, pp. 1321–1339.
- Normann, H. T., 2009, 'Vertical Integration, Raising Rival's Cost and Upstream Collusion,' *European Economic Review*, 53(4), pp. 461–480.
- Ordover, J. A.; Saloner, G. and Salop, S. C., 1990, 'Equilibrium Vertical Foreclosure,' *American Economic Review*, 80, pp. 127–142.
- Ordover, J. A.; Saloner, G. and Salop, S. C., 1992, 'Equilibrium Vertical Foreclosure: Reply,' *American Economic Review*, 82, pp. 698–703.
- Reiffen, D., 1992, 'Equilibrium Vertical Foreclosure: Comment,' *American Economic Review*, 82, pp. 694–697.
- Rey, P. and Tirole, J., 2007, 'A Primer on Foreclosure,' in M. Armstrong and Porter, R. (eds.), *Handbook of Industrial Organization*, vol. III, pp. 2145–2220 (North Holland Press, Elsevier, Amsterdam, The Netherlands).
- Reynolds, S., 2000, 'Durable-Goods Monopoly: Laboratory Market and Bargaining Experiments,' *RAND Journal of Economics*, 31, pp. 375–394.
- Riordan, M., 2008, 'Competitive Effects of Vertical Integration,' Forthcoming in: *Handbook of Antitrust Economics*, Buccirosi Paolo (ed.), pp. 145–182 (MIT Press, Cambridge, Massachusetts, U.S.A.).
- Riordan, M. and Salop, S., 1995, 'Evaluating Vertical Mergers – A Post Chicago Approach,' *Antitrust Law Journal*, 63, pp. 513–568.
- Salinger, M., 1988, 'Vertical Mergers and Market Foreclosure,' *Quarterly Journal of Economics*, 103, pp. 345–356.
- Salop, S. and Scheffman, D., 1983, 'Raising Rival's Cost,' *American Economic Review*, 73, pp. 267–271.
- Salop, S. and Scheffman, D., 1987, 'Cost Raising Strategies,' *Journal of Industrial Economics*, 36, pp. 19–34.