Communication in Vertical Markets: Experimental Evidence

Claudia Möllers

Duesseldorf Institute for Competition Economics (DICE), 40225 Duesseldorf, Germany, tel: +49 211-81-10249, email: moellers@dice.hhu.de

Hans-Theo Normann

Duesseldorf Institute for Competition Economics (DICE), 40225 Duesseldorf, Germany, tel: +49 211-81-15297, email: normann@dice.hhu.de

Christopher M. Snyder

Dartmouth College, Department of Economics, 301 Rockefeller Hall, Hanover, NH 03755, tel: 603-646-0642, email: chris.snyder@dartmouth.edu

October 2014

Abstract: [TBD] When an upstream monopolist supplies several competing downstream firms, it may fail to monopolize the market because of opportunistic behavior towards the downstream firms. We analyze this well-known commitment problem in an experiment where we extend previous research by allowing for communication. In one treatment, the upstream firm can bilaterally talk to either of two downstream firms. In a second treatment, all three firms talk together. We find that the treatment with bilateral communication leads to fewer rejections of offers and higher joint profits than a baseline treatment without communication, but output is still above the monopoly benchmark. Only the treatment where all three firms can communicate leads to complete monopolization. Such communication effectively works as a vertical restraint and should be regarding as potentially anticompetitive.

JEL classification: L42, C90, C70

Keywords: commitment, communication, experiments, vertical restraints

Acknowledgments: The authors are grateful to Patrick Rey and conference participants at CISS Turunç, DFG-ANR Rennes and EARIE Milan for helpful comments.

1. Introduction

[TBD] When an upstream monopolist supplies several competing downstream firms, it may fail to monopolize the market because the upstream firm will behave opportunistically towards the downstream firms. Ideally, the upstream firm wants to restrict output at the monopoly level, but the contracts the monopolist offers suffer from a commitment problem[1]: when dealing with a downstream firm, the monopolist has an incentive to opportunistically negotiate with another downstream firm. These negotiations will disadvantage the first competitor but improve joint profits of the monopolist and the second downstream firm. If this kind of opportunistic behavior is anticipated by the downstream firms, the market will not be monopolized, despite an uncontested upstream monopoly.

Whereas the commitment problem has significant policy implications (Rey and Tirole 2007), caveat is that its relevance depends on players' beliefs as different beliefs suggest different conclusions (see McAfee and Schwartz 1994, and Rey and Vergé 2004). The commitment problem occurs when contracting between firms is secret. So there is imperfect information and downstream firms need to form beliefs about the contracts their rival will be offered. For some beliefs (passive and wary[2]), downstream opportunism arises while for others (symmetric beliefs) it does not. In other words, depending on beliefs, the same model predicts monopolization of the market in the second case but not in the first. Hence, it is an open question whether the commitment problem is relevant from a theoretical perspective.

Experiments seem well suited to investigate the commitment problem. Experiments are helpful here because they can (behaviorally) select between the different solutions the theory offers equilibrium. As the theory is bland (beliefs cannot be rationalized), experiments can help to assess the significance of the commitment problem.

Existing experimental evidence suggests that the problem of downstream opportunism is severe. Martin, Normann and Snyder (2001) (henceforth MNS) analyze markets with one upstream firm who offered non-linear tariffs to two downstream firms. When contracts were secret, markets were rarely monopolized and firms (jointly) earned only about 70 percent of maximum industry

[1] Hart and Tirole (1990), O'Brien and Shaffer (1992) and McAfee and Schwartz (1994) were among the first to highlight this commitment problem.

[2] See the discussion below. In our context, passive and wary beliefs lead to the same outcomes.

profits. By contrast, markets were regularly monopolized when the upstream monopoly was vertically integrated with a downstream firm or when contracts were public, as predicted by theory. The experiments thus support the theory in that the commitment problem is genuine.

Having said that, two aspects of the MNS data are difficult to reconcile with existing theory. First, the markets with secret contracts did not converge to monopoly, but they did not fully approach the duopoly solution predicted by passive or wary beliefs either. MNS extend the theory to allow downstream firms to have heterogeneous (rather than purely passive or symmetric) out-of-equilibrium beliefs. This extension can largely rationalize the data in MNS, among other things because rejections occur in equilibrium (see below). A second unpredicted aspect of the data are bargaining frictions: downstream firms sometimes rejected profitable offers. Such rejections are well known in bargaining experiments. In MNS's treatment with public contracts, this caused efficiency losses such that joint profits were lower than with vertical integration. In any event, the two discrepancies unpredicted by the (standard) theory suggest further investigation.

In this paper, we extend the experimental work of MNS by allowing for communication between players. Our experiments involve one upstream firm, two downstream firms and secret contracts, as in MNS' *SECRAN* treatment. Interactions are essentially one-shot in the experiments due to a random matching scheme, so both the theory with passive/wary beliefs and the previous experimental results suggest the commitment problem will occur. We employ two treatments with communication, in addition to a baseline treatment without communication. In the communication treatments, players can engage in unconstrained chat using a messenger-like tool. In the first treatment, the upstream firm can bilaterally talk to either downstream firm, but the downstream firms cannot observe what the upstream firm and the rival communicate. In the second treatment, all three firms talk together: there is no secrecy, whatever a firm posts will be visible to the other two players in the market.

While communication is cheap talk in our experiments, we expect it to have an impact exactly, specifically with respect to the above-mentioned two aspects underemphasized by the theory. First, communication may clarify which beliefs downstream firms hold. We will argue that, with bilateral communication, downstream firms have little incentive to lie about their beliefs. As a result, we expect better informed upstream firms to offer contracts that are more likely to be accepted by the downstream firms. Put differently, we hypothesize that there will still be heterogeneous beliefs but the upstream firm faces less uncertainty about them with bilateral chat. When all three players can talk together, communication can possibly affect beliefs. There may be downstream firms for whom the thought of beliefs mentioned by others never occurred to them. Likewise, an upstream firm may attempt to talk downstream firms into maintaining beliefs suitable for joint monopolization. If so, markets may converge to monopoly. For both treatments, we expect the data to reject the cheap talk null hypothesis.

A third hypothesis is that communication is likely to reduce bargaining frictions. In previous bargaining experiments, communication often leads to high frequencies of agreement whereas bargaining frictions (rejections) commonly occur in treatments without communication (see Roth 1995). In our case, the downstream firms can inform the monopolist about the share of the surplus they expect. As a result, unexpected rejections of offers would be reduced.

We believe introducing communication to this class of games is interesting from a policy perspective. Communication between vertically related firms is presumably the rule rather than the exception in the field. So it appears relevant to analyze this possibility in the lab. At least bilateral talk between upstream and downstream firms should not violate antitrust laws, however, when the communication not only involves vertically but also horizontally related firms, this would raise antitrust suspicion. Whether such communication has the potential to restrain competition has so far not been investigated.

We also trust that allowing for communication in a vertical structure is a novel and interesting contribution to the experimental literature. The number of research papers on communication in experimental economics is rising (see the literature survey in the next section), but vertical structures have largely not been investigated. While bargaining experiments with communication have been analyzed before, we are not aware of such an experiment where play is subject to a bargaining externality.

Our results are as follows. First, our treatment with bilateral communication between the upstream and downstream firms leads to higher accept rates and higher joint profits, but the average output on the market (the accepted quantities) are just as high as in the baseline treatment without communication, that is, above the monopoly benchmark. This result suggests that the commitment problem is substantial and cannot be overcome by this form of communication. Second, the treatment in which all three firms can communicate does indeed lead to full monopolization. Bargaining frictions are much reduced (as in the first treatment with bilateral communication) but downstream firms obtain a higher share of the joint profits. From a policy perspective, such communication effectively works as a vertical restraint and should be regarding as potentially anticompetitive.

2. Related Literature

[TBD] We review the mostly closely related experimental papers. The main contributions to the theory (Hart and Tirole 1990, O'Brien and Shaffer 1992, McAfee and Schwartz 1994, Rey and Vergé 2004, Rey and Tirole 2004) have been highlighted already in the introduction. For more recent theory papers on the commitment problem, see Avenel (2012) or Rey and Caprice (2012).

2.1. Experiments with vertically related markets

Vertically related markets have been studied in experiments in a few research papers. Mason and Phillips (2000) investigate 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 vertically integrated and nonintegrated monopolies and analyze if the vertical merger mitigates the double marginalization problem. In Normann (2012), the market structure is a bilateral duopoly. He investigates whether a vertical merger have a "raising rivals' cost" effect. MNS (described in detail in this paper) analyze the commitment problem arising in vertically related markets. In this paper, we investigate whether the commitment problem is mitigated by communication.

2.2. Experiments with communication

The experimental literature on the effects of communication is large and we confine ourselves to a few relevant areas. Generally communication is many experiments is cheap talk because it has no binding effect. The effects of such cheap talk communication differ depending on the type of game which is played.

Farrell and Rabin (1996) compare the effects of cheap talk on coordination games with games in which there is a conflict of interests. In coordination games, there is a great benefit of cheap talk because it may lead to efficient outcomes. Believing in the messages from one player to another is an equilibrium in this case (however, it is not unique). Charness (2000) finds evidence in an experiment in favor of this theory: "results show impressive coordination". See also Crawford (1998) for a survey. There are some similarities between coordination games and our vertically related markets: there can be coordination failures when offers made by the upstream firms do not match the beliefs of the downstream firms.

Whereas coordination games seem to be a good example where cheap talk leads to higher efficiency, the case of games of conflicting interests with communication are more complicated. Rational players will generally ignore any agreements to cooperate since it is their best response. However, there is experimental evidence for an improvement of cooperation in dilemma games when communication is possible (see Dawes et al. 1977, Isaac et al. 1984, Isaac and Walker 1988, Balliet 2010). Whereas the firms in our game could be said to be in a dilemma situation (best responding prevents monopolization for certain beliefs), decision making is more complex: no firm can unilaterally "defect" (or "cooperate", for that matter); instead, coordination on certain outputs (corresponding to a "defect" choice) between upstream and downstream firms is required.

Communication can also help in achieving superior outcomes in trust games (Charness and Dufwenberg 2006, Charness and Dufwenberg 2010). Our vertically related markets also involve a situation of trust: a downstream firm need to trust any promises about joint monopolization the upstream firm may make. Whereas in a typical two-player trust game the first mover needs to trust the second mover (for example to return money), it is the second movers (a downstream firm) who needs to trust the first movers (the upstream firm) in our case. Any promises made by the upstream monopolist are not about own future actions but regard a simultaneously made offer to the other downstream firm (who is a trustor at the same time). Even though we have an element of a trustee-tustor relationship in our game, the parallels to standard trust games are perhaps not overly strong.

Perhaps foremost, our experiment is a bargaining game with externalities (and with communication). While the upstream firms make take-it-or-leave it offers, these offers can be intensively discussed upfront. Roth (1995) reports results from such bargaining experiments. With face-to-face communication, he reports near-perfect efficiency rates. One study Roth summarizes is by Radner and Schotter (1989) who run various experiments with a sealed-bid bargaining mechanism. Buyer values and seller costs are private knowledge in their experiment. They report that face-to-face communication leads to 99% efficient trades. Roth also reports on an experiment where a fixed sum of money is to be divided and subjects only very rarely fail to agree how to split the money with face-to-face communication. In experiments with unrestricted types messages (Roth and Murningham 1982), lower levels of efficiency occur but these levels are still higher than in standard (no communication) variants. Our experiments will involve unrestricted typed messages but not face-to-face communication. We hence expect improved but imperfect efficiency rates due to communication. More recent bargaining experiments with similar results include Brosig et al. (2004) and Zultan (2012).

2.3. Exclusive dealing

The literature on exclusive dealing is related because it concerns a similar case of sequential bargaining with externalities between players. Contracts may be secret there, too. A typical setting involves an incumbent seller, a (more efficient) entrant and, say, two buyers with independent demand. Because of fixed costs, the entrant needs to have both buyers signing up with him for entry to be profitable. The question is whether buyers manage to coordinate such that entry occurs. Rasmusen, Ramseyer and Wiley (1991) and Segal and Whinston (2000) show that the incumbent may need only one buyer to sign an exclusive contract to deter entry. Landeo and Spier (2009) report experimental evidence on this issue and employ treatments with and without communication between buyers. They find that, with no-discrimination clauses, communication between buyers reduces the likelihood of exclusion substantially. This is not the case when discrimination is possible. Landeo and Spier (2009) mainly focus on public offers but, for a subset of treatments, they also make the comparison to secret offers. They find that they significantly reduce the likelihood of exclusion. In Landeo and Spier (2009), the coordination problem is between the buyers. In our game, upstream and downstream firm need to coordinate on offers that match beliefs.[3, 4]

[3]In a related paper, Smith (2010) finds that exclusion is more difficult for the incumbent when he needs a larger fraction of buyers to accept a contract and when buyers are allowed to communicate with each other. Another related experiment is by Boone, Müller and Suetens (2014). They find that exclusion rates in the sequential regimes are significantly higher than in the non-discriminatory regime and the simultaneous discriminatory regimes.

[4]Communication has recently found further interesting applications in industrial organization. Anderson and Wengström (2007) analyze costly communication in Bertrand duopoly. They find that prices are higher and collusion is more stable when communication is costly. Fonseca and Normann (2012) investigate Bertrand oligopolies with and without communication. Specifically, they analyze how the gain from communication is affected by the number of firms (ranging from two to eight). Cooper and Kühn (2013) study conditional cooperation: a simple cooperation game is followed by a coordination game, so the threat of coordinating on a payoff-inferior equilibrium in stage two

3. Theoretical Framework

3.1. Model

Consider a simplified version of the model due to Rey and Tirole (2007).[5] The market has a vertical structure shown in Figure 1, with a monopoly upstream firm, U, and two downstream firms, D_i, $i = 1,2$. The upstream firm produces an intermediate product at zero cost. The downstream firms transform this product on a one-for-one basis, also at zero cost, into a final good sold to consumers. Consumers have inverse demand $P(Q)$ for this homogeneous final good.[6]

The timing is as follows. First, U offers contracts (x_i, T_i) to each D_i specifying a quantity x_i and fixed tariff T_i. Second, the D_i simultaneously decide whether to accept $(a_i = 1)$ or reject $(a_i = 0)$ the their contract offers. The rest of the game proceeds deterministically from those decisions. Each D_i produces $q_i = a_i x_i$ resulting in total output $Q = q_1 + q_2$. Profits are $a_1 T_1 + a_2 T_2$ for U and $P(Q)q_i - a_i T_i$ for D_i.

To set some benchmarks, let $Q^m = \text{argmax}_Q P(Q)Q$ be the monopoly quantity for this market and $\Pi^m = P(Q^m)Q^m$ be monopoly profit. Let q^c be a firm's equilibrium quantity from Cournot competition between two firms in a market in which the vertical structure from Figure 1 were compressed into a single level. That is, defining the best-response function

$$BR(q) = \underset{\tilde{q}}{\text{argmax}}\, P(\tilde{q} + q)\tilde{q},$$

q^c is the fixed point $q^c = BR(q^c)$. Let $\pi^c = P(2q^c)q^c$ be a firm's Cournot profit.

3.2. Commitment problem

Suppose first that contracts are public, meaning that each D_i can see the contract offered to its rival. In this case, U can extract the whole monopoly profit in equilibrium. For example, there exists a

is credible. They analyze what type of communication is most effective in achieving cooperation in this setup.

[5] Rey and Tirole (2007) is itself a simplified version of a number of earlier papers including Hart and Tirole (1990) and McAfee and Schwartz (1994). We modify Rey and Tirole (2007) in three ways. First, contracts here specify a single bundle at a fixed tariff rather than a tariff function. Second, downstream firms here make a simple accept/reject decision rather than choosing some continuous quantity. Third, upstream marginal cost is set to $c = 0$ to simplify the analysis and reflect experimental conditions to follow.

[6] Assume $P(Q)$ has properties ensuring that the Cournot game formed by compressing the vertical structure in Figure 1 into a single level is well behaved. In particular, the resulting profit functions are strictly quasiconcave and actions are strategic substitutes. A sufficient condition is $P'(Q) + P''(Q)Q < 0$ for all Q.

symmetric equilibrium in which U offers the contract $(Q^m/2, \Pi^m/2)$ to each D_i, and both accept. To see this is an equilibrium, note U has no incentive to deviate since it cannot earn more than the monopoly profit. The D_i earn zero profit whether or not they accept so have no strict incentive to deviate.

Now suppose that contracts are secret, meaning that D_i cannot see the contract offered to its rival. The outcome in which U offers $(Q^m/2, \Pi^m/2)$ to each D_i may no longer be an equilibrium. U may not be able to commit to supplying just half the monopoly output to each D_i. Instead, it may prefer to deviate to a higher output, for example, the best response to half the monopoly output $BR(Q^m/2)$, which could increase the joint profits of the contracting pair at the expense of the other downstream firm. Whether D_i would accept this deviating offer depends on its posterior beliefs about its rival's contract offer following a deviation, the issue discussed next.

3.3. Equilibria with secret contracts

Secret contracts transforms the model into a dynamic game of imperfect information. The relevant solution concept in such games perfect Bayesian equilibrium, requiring strategies to be best responses at every node given posterior beliefs and requiring posterior beliefs to be formed using Bayes rule along the equilibrium path. Bayes rule does not pin down beliefs off the equilibrium path, in particular D_i's updated beliefs about its rival's secret contract after D_i receives a deviating offer. Different assumptions about out-of-equilibrium beliefs give rise to different perfect Bayesian equilibria.

One assumption, called *symmetric beliefs*, is that D_i believes its rival receives the same deviating contract. Under such beliefs, U can obtain the same monopoly outcome as it did with public contracts, i.e., having both D_i accept contract offers $(Q^m/2, \Pi^m/2)$. To see that this is an equilibrium, note that if U deviates to some quantity x^d in its contract offer, D_i would be unwilling to pay a fixed tariff greater than $P(2x^d)x^d$,[7] which is obviously no greater than the fixed fee $\Pi^m/2$ that U charged in the equilibrium contract. By rendering deviation unprofitable for U, symmetric beliefs effectively preserve U's ability to commit to the monopoly outcome.

Another assumption, called *passive beliefs*, is that after receiving a deviating offer, D_i continues

[7] This assumes D_i also believes its rival accepts the deviating contract. If one or both downstream firms rejects the deviating contract, deviation would be certainly less profitable than the equilibrium $(Q^m/2, \Pi^m/2)$ contracts to each.

to believe its rival receives the equilibrium contract. In this case, there will always exist a strictly profitable deviation unless equilibrium firm quantity q^* is best response to itself, i.e., $q^* = BR(q^*)$. But as we saw above, the Cournot output q^c is the unique quantity satisfying this equation. U can extract at most the Cournot duopoly profit from each firm with fixed tariff π^c. Hence the equilibrium contract offer is (q^c, π^c), which both D_i accept.

Another assumption, introduced by McAfee and Schwartz (1994), called *wary beliefs*, is that after receiving a deviating offer D_i believes its rival receives and accepts a contract that is the best response to this deviation. In the present context in which downstream firms essentially engage in Cournot competition, wary beliefs turn out to select the same perfect Bayesian equilibrium as passive beliefs.[8] In most of the rest of the paper, for brevity, statements that apply equally to wary and passive beliefs will just mention passive beliefs.

Because neither the monopoly outcome predicted when all downstream firms have symmetric beliefs nor the Cournot outcome predicted when they all have passive beliefs fit their experimental results well, Martin, Normann, and Snyder (2001) proposed a model of heterogeneous beliefs. They assume that each D_i is an independent draw from a population in which a fraction $s \in [0,1]$ hold symmetric beliefs and $1-s$ hold passive beliefs. The population distribution is common knowledge, but the realized beliefs for the D_i are private information. The authors show that there exists a threshold $\hat{s}$, the value of which depends on the experimental parameters, such that for $s \in (0, \hat{s})$ the extremal perfect Bayesian equilibrium (i.e., the one yielding the highest profit to U) involves U offering the Cournot duopoly output, q^c, as with passive beliefs. However, the fixed tariff is higher, $T_i > \pi^c$, inducing D_i to respond with an acceptance probability strictly less than one. The heterogeneous-beliefs model could rationalize the modal contract offers observed in the experiment, of the form (q^c, T_i) with $T_i > \pi^C$, as well as the observed acceptance rates.

[INSERT PASSAGE ABOUT BARGAINING AND/OR COMMUNICATION]

4. Experimental Design

We build on the experimental design of Martin, Normann, and Snyder (2001), which is essentially an experimental implementation of the model of Rey and Tirole (2007) discussed in the previous

[8] Rey and Vergé (2004) show that wary and passive beliefs lead to different equilibrium outcomes if downstream firms engage in Bertrand competition.

section. We will maintain their baseline treatment—called *SECRAN* because it involves secret contracts with randomly re-matched players—as our baseline treatment with no communication here. We will then introduce treatments allowing for different forms of communication.

The market, shown in Figure 1, involves three subjects, one playing the role of the upstream firm (called a manufacturer in the experiment) and two playing the role of downstream firms (called retailers in the experiment). The upstream player moves first, making a take-it-or-leave-it offer (x_i, T_i) to each D_i, where x_i had to be an integer in $[0, 10]$ and T_i had to be an integer in $[0, 120]$. After observing its own contract only, D_i chooses whether to accept ($a_i = 1$) or reject it ($a_i = 0$). These decisions result in each D_i supplying $q_i = a_i x_i$ to the final-good market, for a total supply of $Q = q_1 + q_2$. Market price $P(Q)$ is calculated from the discrete demand function in Figure 2A. All firms produce at zero cost. Thus profits are $\pi_U = a_1 T_1 + a_2 T_2$ for U and $\pi_{Di} = P(Q) q_i - a_i T_i$ for D_i. Let $\pi_D = \pi_{D1} + \pi_{D2}$ denote total downstream profit and $\Pi = \pi_U + \pi_D$ denote market profit. Figure 2B graphs the profit function in the experiment; it is concave, achieving a maximum of $\Pi^m = 100$ at an output of $Q^m = 2$.

Participants were randomly assigned to their roles (U or D_i), which they played each round for the entire course of the session. We recruited 15–21 subjects for each session, allowing us to form 5–7 markets. Each session consisted of 15 rounds of game play. As in *SECRAN*, the three subjects constituting a market were randomly re-matched before every round to minimize effects of repeated interaction. (Experimenter effects aside, observations may be dependent within sessions but should be independent across sessions because new subjects were recruited for each session.) After each round, each D_i was told the profit it earned that period; U was told his profit and the profits of the two downstream firms with whom he was matched in that round. All these design features were explained to subjects in the instructions.

We conducted three different treatments, summarized in Table 1. Our baseline treatment replicates the *SECRAN* treatment from Martin, Normann, and Snyder (2001). To compare the communication element with other treatments, in particular that there is no communication involved, we relabel this treatment *No Chat*. The other two treatments, *Bilateral Chat* and *All Chat*, introduced the possibility of communication using an instant-messaging technology via a chat window. In *Bilateral Chat*, U could engage in bilateral communication with each D_i. D_1 and D_2 could not communicate with each other, and D_i could not observe U's communications with his competitor.

U had separate chat windows for each D_i on its screen; each D_i had only one chat window on its screen through which it communicated to U. In *All Chat*, U, D_1, and D_2 could freely communicate with each other. Whatever a player posted into its chat window is displayed to all three players in the market (including the player himself). It was not possible to exclude one of the players and engage in bilateral chat.

Every round of *Bilateral Chat* and *All Chat* included a communication stage prior to U's making its contract offer. Except for threats outside the lab or messages containing information that could be used identify subjects, which were forbidden, the content of the chat was unrestricted. In *Bilateral Chat*, subjects had 90 seconds to chat during the first five periods, reduced to one minute for the last ten periods. In *All Chat*, the communication stage lasted 60 seconds in all 15 periods. Subjects could not leave the chat stage before the time expired. The design of *Bilateral Chat* and *No Chat* were otherwise identical to *No Chat*.

Subjects were invited using the ORSEE system (Greiner 2004). Upon arrival in the lab, each was assigned to a cubicle. After all subjects had arrived, we provided them with instructions, reproduced in the appendix. The instructions were the same in all treatments except for a short section about the chat stage added in the communication treatments. After reading the instructions, subjects were allowed to ask questions privately in their cubicles. Subjects were then informed about their role in the experiment (U or D) and the experiment proceeded.

It is possible for downstream firms to earn a negative payoff in a round of the experiment. To offset this possibility as well to provide a payment for showing up, subjects playing the D role received an initial endowment of 200 ECU (experimental currency units). Subjects playing the U role received an initial endowment of 60 ECU. At the end of the experiment, participants were paid in euros, exchanged at a rate of one euro for each 40 ECU. Participants earned an average of about 14 euros each.

All sessions were run at DICElab of the University of Duesseldorf between November 2013 and May 2014. We ran 12 total sessions, four for each of the three treatments. Each session lasted for about one hour. In total, 216 subjects participated.

5. Hypotheses

[CONSIDER DIFFERENT APPROACH WHERE TAKE WEAKER STANCE TOWARD AT LEAST SOME OF THIS. COULD HAVE BABBLE? BILATERAL CHAT COMPLICATED]

[CONSIDER DIFFERENT TERM, CONJECTURE?, BECAUSE DON'T HAVE A THEORY OF COMMUNICATION GENERATING REAL HYPOTHESES.]

We start with two market benchmarks, monopoly and Cournot duopoly. The joint-profit maximizing solution is a quantity of two, yielding industry profits of 100. With symmetric beliefs the monopoly outcome is expected to emerge, with each D_i supplying one unit. The Cournot outcome predicted under the assumptions of passive/wary beliefs has both downstream firms accepting two units. Hence, there is a total quantity of four and a joint profit of 72 for the Cournot outcome.

As for our *No Chat* treatment, we expect the same results as those in MNS's *SECRAN* treatment. That is, the market will rarely be monopolized, quantities offered will be in excess of the monopoly output, and offers are regularly rejected—as suggested by the model with heterogeneous beliefs.

How will communication affect outcomes in the experiments? As players might simply regard the chat as cheap talk, they might disregard it entirely. The PBE would be unaffected by cheap talk. Hence, the null hypothesis is that the possibility to communicate in *Bilateral Chat* and *All Chat* has no impact on market outcomes.

Moving away from the null, we do expect communication to change results. As noted in our literature survey, efficiency rates in bargaining games with communication can be rather high (Roth, 1995). Assuming this holds, for our setup, two questions arise: how do firms manage to bargain efficiently here? And second: can they manage to overcome the commitment problem through communication?

In *Bilateral Chat* we hypothesize that D_i may honestly indicate to U what kind of belief they hold. In fact, they will have an incentive to tell the truth about their beliefs. Suppose D_1 actually has passive beliefs and would thus accept $(2,36)$, expecting a non-negative profit from this contract. Claiming to have symmetric beliefs (and given U believes the claim) would induce U to offer $(1,50)$ to D_1. Since D_1's true belief is passive, it believes D_2 will be offered and accept $x_2 = 2$, so D_1 would make a loss by accepting $(1,50)$. Thus, a passive D_1 will not pretend to have symmetric beliefs. Now consider the case when D_1 actually has symmetric beliefs. Whatever the (x_1, T_1) contract D_1 is being offered, it believes that D_2 gets the same offer. D_1 expects a non-negative profit from the contract $(1,50)$ and would thus accept it. Pretending to have passive beliefs would U induce to make the offer $(2,36)$. The expected profit is also zero (because of symmetric beliefs) here for D_1 but the surplus (which can possibly be divided) is even smaller. Hence, a D_i with symmetric beliefs would not lie either.[9]

Will the U firms not try to convince the D firms to accept $(1,50)$ offers in *Bilateral Chat*? They might attempt to do so, but because communication is bilateral only, this is unlikely to work out. D_i will distrust chat content suggesting joint monopolization because it realizes that once it agrees to accept (1, 50) in the chat, U will agree on $(2, T_j)$ with D_j. This is the heart of the commitment problem.

What we hence expect in *Bilateral Chat* is improved coordination but no monopolization: given the heterogeneity of beliefs, U now better knows what the D firms believe and can design appropriate offers. While the above model with heterogeneous beliefs involved rejections of contracts in equilibrium, rejection rates should be lower with communication, leading to higher industry profits. However, we expect the heterogeneity of beliefs persists, including players with passive beliefs. Thus, we do not expect commitment problem itself to be resolved and we maintain no hypothesis as to how *Bilateral Chat* should affect outputs.

Hypothesis 1. Acceptance rates and industry profits will be higher *Bilateral Chat* than in *No Chat*.

In the *All Chat* treatment, a similarly improved coordination due to the revelation of beliefs might occur but, beyond that, a change of beliefs now seems possible. Firstly, there may be some downstream firms for whom the thought of the other set of beliefs never occurred to them. So some learning or revising of beliefs may kick in when D_i learns about D_j's beliefs. In addition, U firms may now attempt to convince downstream firms about symmetric beliefs and joint monopolization

[9] If, contrary to the above assumption, a symmetric believer (D_1, say) thinks that the other retailer (D_2) will receive (and accept) the offer $(1,50)$ even if D_1 pretends to have passive beliefs, it looks at first sight that D_1 might gain from communicating false beliefs. If these offers were accepted, $Q = 3$ units would be on the market, D_1 would earn $2 \cdot 30 - 36 = 24$ which is strictly better than the (zero) payoff from accepting $(1,50)$ (D_2 would make a loss of $30 - 50 = -20$ in this case). Such reasoning would, however, be naive as D_2 either holds symmetric beliefs and reasons (and communicates) the same way, with the result that both retailers accept $(2,36)$ and make zero profits. Alternatively, D_2 is actually a passive believer in the first place in which case he will also accept $(2,36)$ (and not lie about his beliefs).

of the market. A chief reason for this is that the channel for secret communication with one D_i only is closed, as opposed to *Bilateral Chat*. In the chat, promises about offers will inevitably be symmetric and therefore conducive to monopolization.

Hypothesis 2. Market output will be lower, and acceptance rates and industry profits will be higher in *All Chat* than in *Bilateral Chat* and *No Chat*.

Finally, we expect a bargaining effect in *All Chat* and *Bilateral Chat*. The *D* firms may not only communicate the expectation about the quantity offered to the rival, they may also indicate a maximum willingness to pay for a certain bundle of units. Such demands may be credible when subjects have other-regarding preferences, but they may also be exaggerated for strategic reasons and hence not be credible from a theory point of view. In any event, we know from previous experiments (Roth, 1995; Zultan, 2012) that they can have an effect. Since players will know (or learn) the size of pie they bargain about, our hypothesis is that *D* firms can obtain a larger *share* of the pie.

Hypothesis 3. The share of industry profits obtained by the downstream firm will be higher *All Chat* and *Bilateral Chat* compared to *No Chat*.

6. Results

6.1. Preliminaries

This section reports our main experimental findings. Though each session ran for 15 rounds, to reduce noise from play by subjects still unfamiliar with the experiment, for most of the analysis we restrict attention to the last 10 rounds. The results are indeed stronger using just the last ten rounds but do not qualitatively change if we use all 15 rounds. Comparing early- to late-round play does yield some interesting insights, but we will defer that comparison until Section 6.7.

Table 2 provides summary statistics for the main experimental variables. The first row can be used to compare our baseline treatment here, *No Chat*, to the *SECRAN* baseline from Martin, Normann, and Snyder (2001). Comparing these treatments provides a consistency check because their designs are identical; they only differ in being run a decade apart with different subjects. Total offered quantity, $X = x_1 + x_2$, averaged 3.44 in *SECRAN*, similar to the 3.63 in *No Chat*. The averages for total accepted quantity $Q = q_1 + q_2$ are also quite similar—2.47 in *SECRAN* versus 2.54 in *No Chat*—as are the averages for industry profit Π—69.9 in *SECRAN* versus 70.2 in *No Chat*. Upstream firms earned somewhat higher profit π_U in *SECRAN* (mean 53.0) compared to *No Chat* (mean 45.3). Besides providing additional confidence in the stability of Martin, Normann, and Snyder's (2001) main results, the remarkable consistency between *SECRAN* and *No Chat* suggests that *No Chat* is a good baseline for comparing treatments with communication.

6.2. Offered output

We begin by analyzing the contract offers observed in the experiment, starting with total offered output, $X = x_1 + x_2$. This single variable captures whether U is able to solve the commitment problem, allowing it to restrict offered output to the monopoly level, $X = 2$, or "burns" profit by supplying more, for example the Cournot output $X = 4$. Table 2 shows that the mean of X is highest in *No Chat*, 3.63, falling to 2.80 in *Bilateral Chat*, falling further to 2.18 in *All Chat*, close to the monopoly output. These results are consistent with more communication facilitating commitment and monopolization. Figure 3 provides a more disaggregated picture, showing the means of the four individual sessions run for each treatment. The filled circles in Panel A represent session means for X. The group of means fall together from *No Chat* to *Bilateral Chat* to *All Chat*, providing confidence that decline toward the monopoly level is not due to an outlying session.

Table 3 provides formal statistical tests of the differences among the means across sessions. The first column regresses X on an exhaustive set of treatment indicators, suppressing the constant. This specification allows us to recover the means from Table 2 as the coefficients on the indicators; the advantage of the regression is that the supplied standard errors allow statistical tests of the differences between the means. We compute White (1980) heteroskedasticity-robust standard errors that are also clustered by session, allowing for dependence among observations arising from the same set of interacting subjects. The bottom part of the table reports differences between all combinations of coefficients on the treatment indicators, providing the appropriate standard errors for these differences. The differences are statistically significant at better than the 1% level as well as being economically substantial. The fall in the mean of X from *No Chat* to *Bilateral Chat* of 0.81 represents 40% of the gap between the monopoly and Cournot market output. The fall from *Bilateral Chat* to *All Chat* of 0.63 brings the offered quantity very close to the monopoly level of $X = 2$.

Figure 4 provides a histogram for X for the various treatments in Panel A. The white bars for *No Chat* show a mode at $X = 4$ and considerable additional mass on yet higher offers. Moving from the white to the grey bars, representing *Bilateral Chat* observations shifts the mass of the distribution from these high levels to the lower levels $X = 2$ and $X = 3$, and $X = 2$ becomes the mode. Moving to the black bars for *All Chat* piles almost all the mass in the monopoly ($X = 2$) bin.

Table 4 can be used to test for the statistical significance of these shifts in the histogram. The first column is a linear probability model regressing a 0–1 indicator for whether $X = 2$ on an exhaustive set of treatment indicators, again suppressing the constant. This specification allows us to recover the relative frequency of the monopoly outcome (graphically, the height of the bars in Figure 4A in the $X = 2$ bin) directly from the coefficients on the treatment indicators. The supplied standard errors, heteroskedasticity robust and clustered at the session level, allow statistical tests of the difference across treatments, which are reported in the lower part of the table. *All Chat* is 34 percentage points more likely to generate monopoly offers than *Bilateral Chat*, a difference significant at the 1% level. *Bilateral Chat* is 23 percentage points more likely to generate monopoly offers than *No Chat*, a difference significant at the 10% level. The next column regresses an indicator for the event $X \geq 4$, i.e., that the offers total to at least the Cournot output. *All Chat* is 19 percentage points less likely than *Bilateral Chat* to have offers this high, and *Bilateral Chat* is 31 percentage points less likely than *No Chat* to have offers this high, both differences significant at the 5% level. We conclude that reducing the intensity of communication entailed by moving from *All Chat* to *Bilateral Chat*, and then from *Bilateral Chat* to *No Chat*, results in each instance in a statistically significant shift in the mass from the monopoly bin to the right.

6.3. Tariffs

We next turn to the other variable in the contract, the fixed tariff T_i. Because it is a pure transfer between parties, this variable provides a clean picture of how communication affects the division of surplus in the experiment. The mean reported in Table 2 is similar in *No Chat* and *Bilateral Chat*, slightly above 30 in each case, but much lower in *All Chat*, 22.5. Table 3 shows that the means in *No Chat* and *Bilateral Chat* are not significantly different from each other, but the mean in *All Chat* is significantly lower than the others at the 1% level. The scatter plots of session-level means in Figure 3B tells a similar story. These results hint that the introduction of bilateral communication does not affect how negotiating parties divide surplus but multilateral communication among all parties does, allowing the downstream firms to extract a greater share.

We cannot make definitive inferences from the raw means of T_i, however, because x_i varies systematically across treatments as well, as we saw, and this variation in x_i can pollute the inference. To understand why, consider the contracts $(1,30)$ and $(2,30)$. While they specify the same fixed tariff of 30, if D_i has symmetric beliefs, the first contract is more generous, providing him with a profit of 20 compared to 6 for the second contract.[10] To purge these quantity effects, the third column of Table 3 includes fixed effects for each different integer value of x_i in the regression of T_i on the treatment indicators. The conclusions from the raw means remain unchanged: *No Chat* and *Bilateral Chat* involve similar tariffs, but the tariff in *All Chat* is significantly lower.

The result that *Bilateral Chat* does not result in tariff reduction relative to *No Chat* seems at odds with previous experimental work showing that introducing communication in the ultimatum game led to more generous splits for the responder. Here the responder (downstream firm) does benefit from communication, not from a bigger slice of a fixed pie, but from a fixed slice of a larger pie due to monopolization of the market. This allows U to be more generous to D_i without having to offer a lower T_i.

6.4. Acceptance behavior

Having analyzed upstream behavior, we next turn to downstream behavior, embodied in the acceptance decision a_i. Based on results from previous experiments we expect communication to improves acceptance rates. The means of a_i in Table 2 bear these expectations out. The acceptance rate rises from 72% in *No Chat* to 85% in *Bilateral Chat* to 92% in *All Chat*, remarkably high in an experimental setting. Table 3 shows that these increases are significant at the 10% level or better in both cases. Figure 3C shows this same pattern holds in the session-level means.

The raw means of a_i provide evidence on how equilibrium acceptance rates vary communication conditional on equilibrium contracts. The fifth column of Table 3 sheds light on how accep-

[10] With passive beliefs, the computation is less clear because the generosity of a contract depends on whether it is an equilibrium or out-of-equilibrium offer.

tance rates vary with communication holding constant the contract offer. This column regresses a_i on the treatment indicators controlling for the contract's terms in a semi-parametric way by including fixed effects for each integer value of x_i, a linear and squared term in T_i, and the interaction between T_i and the fixed effects for x_i. Given that we omitted the fixed effect for $x_i = 1$ to avoid multicolinearity, the coefficients on the treatment indicators can be interpreted as the acceptance rates of a contract offering one unit for free; note that all three are not statistically significantly different from 1 for this.

Controlling for contract offer reduces the gap between the *No Chat* and *Bilateral Chat* acceptance rates and reverses the sign on the difference between *Bilateral Chat* and *All Chat*. The reason the acceptance rates are so high in *All Chat* is not that open communication somehow makes the D_i more receptive to offers but because U offers more generous contracts, involving more profitable output levels and lower tariffs.

[ARE THE HYPOTHESES FOR REDUCED FORM OR STRUCTURAL ACCEPTANCE LEVELS? I'M NOT SURE.]

6.5. Market Output

The rest of the variables for which provide summary statistics in Table 2 are deterministic functions of subjects' actions in the experiment. Still they deserve some study because these would be the observables in a non-experimental market. The means for market output $Q = q_1 + q + 2$ show the same pattern in Table 2 as X, falling from 2.54 in *No Chat* to 2.37 in *Bilateral Chat* to 1.98 in *All Chat*. The statistical tests in Table 3 show that the 0.19 drop from *Bilateral Chat* to *All Chat* is statistically significant, but the 0.16 drop from *No Chat* to *Bilateral Chat* is not.

These relatively small changes in the mean of Q belie much more dramatic changes to the overall distribution, shown in Figure 4B. Moving from *No Chat* to *Bilateral Chat* to *All Chat*, the distribution is increasingly concentrated from above and below on the mode at the monopoly outcome. The concentration from above is inherited from the effect of communication on the concentration of offers on the monopoly ($X = 2$). The concentration from below is inherited from the increase in the raw acceptance rate with better communication, reducing the mass in the $Q = 0$ and $Q = 1$ bins, which, except for one case out of 720, never arise unless there has been a rejection. Looking at the coefficient differences in the sixth column of Table 4, the monopoly outcome ($Q =$

2) is 18 percentage points more likely in *Bilateral Chat* than *No Chat* and 31 percentage points more likely in *All Chat* than *Bilateral Chat*, both significant at at least the 5% level.

Thus, more communication leads to more monopolization. *All Chat* is conducive to monopolization not just in a relative sense (i.e., relative to the other treatments) but in an absolute sense, attaining the monopoly outcome in a remarkable 81% of the observations. Thus, free communication facilitates nearly complete monopolization whether measured in terms of offered or actual quantity.

6.6. Profits

Although profits are determined by the variables already studies, further analysis of profits will let us put a monetary value on the differences across treatments uncovered so far. First consider industry profit, Π. Table 2 shows that the mean rises from 70.3 to 84.3 to 92.4 ECU. Table 3 shows that the increase in the mean of Π from *No Chat* to *Bilateral Chat* of 8.1 ECU and from *Bilateral Chat* to *All Chat* of 14.1 ECU are statistically significant at the 1% level. These profit increases are the direct consequence of the concentration of the distribution of Q on the bin ($Q = 2$) that maximizes industry profits. Mean profit in *All Chat*, 92.4 ECU, is remarkably close to the monopoly profit of 100.

Moving to the allocation of profit across industry levels, we have seen that the move from *No Chat* to *Bilateral Chat* resulted in a slight fall in T_i and a moderate increase in the acceptance rate. These two forces combine to increase the mean of π_U slightly from 45.3 to 49.9 ECU, although Table 3 shows this 4.6 ECU increase is not significant. The large fall in T_i from the other treatments to *All Chat* does not offset the increase in acceptance rate for U. The mean of π_U is significantly lower in *All Chat* than either *No Chat* or *Bilateral Chat*. Therefore, while more communication benefits the industry as a whole by allowing it to increasingly monopolize the market, the firm that one might think would be in the best position (because it has no competitors and makes contract offers) to appropriate this benefit does not. Introducing communication has at best a small effect but at worst a measurably harmful effect on U. Apparently, U's take-it-or-leave-it bargaining power is eroded faster than its commitment power is increased. Thus, while a meeting in the proverbial "smoke-filled room" would benefit the industry as a whole, U would veto such a meeting.

Downstream firms obviously benefit from more communication because this increases industry profits but reduces their tariffs. Table 2 confirms this intuition, showing an increase in the mean of π_D from 25.0 in *No Chat* to 34.5 ECU in *Bilateral Chat*, and a further increase to 51.6 ECU in *All Chat*. Downstream profit is so high in *All Chat* that they are obtaining a majority of the profit.

6.7. Within-session trends

The analysis so far was conducted on a subsample of the data excluding the first five rounds to focus on subjects' play after they were familiar with the experiment. The sign and the significance of all the results discussed so far are preserved if instead we use all 15 rounds. However, it is worth briefly analyzing the full sample because some interesting trends are revealed.

To uncover these trends, Table 5 repeats the regressions from Table 3 but using the full sample and allowing for different treatment effects in the initial period (rounds 1–5) and the last period (rounds 6–15). We do this by including interactions between the treatment indicators and indicators for the initial period (rounds 1–5) and the last period (rounds 6–15). For example, *No Chat*$_0$, is the interaction between the *No Chat* indicator and an indicator for rounds 1–5, and *No Chat*$_1$ is the interaction between *No Chat* and an indicator for rounds 6–15. The bottom of the table reports the change in the treatment indicator across the two periods along with the appropriate standard error, allowing an assessment of the significance in of the change. For space considerations, we only report one regression for each variable, the one with just treatment indicators and no additional controls.

The results show a fairly distinct pattern. *No Chat* shows few significant changes over time. By contrast, almost all the variables have significant changes, many at the 1% level, for the treatments with communication. What this pattern reveals is that subjects played fairly consistently over the rounds in *No Chat* but took several rounds to settle down to how they eventually played in the communication treatments. Apparently subjects needed more time to understand the value of communication and perhaps build trust in the promises made by their bargaining partner. The differences between *No Chat* and the communication treatments are already apparent in the initial period (we do not report the differences for space considerations), but they are moderate at that point. As play progresses into the later rounds, the communication treatments diverge from *No Chat* and increasingly reveal the distinctive monopolization and bargaining effects we have been highlighting. Subjects in the communication treatments are better able to monopolize the market in the later rounds, shown by a significant decline in X and Q. This monopolization leads to a significant rise in industry profit Π. U is more generous with the D_i over time, leading to significant reductions in T_i, significant reductions in π_U, and significant increases in π_D.

This analysis of within-session trends suggests that our main findings are representative of play by experienced agents and thus should not be expected to disappear over time in real markets. Play in the simple treatment without communication settles down almost immediately to long-run averages. Play in the treatments with communication takes time to settle down, perhaps because the environment is more complex, perhaps because subjects need time to develop trust in trading partners cheap talk.

7. Analysis of Chat Content

Why does cheap talk hardly reduce output in *Bilateral Chat* but lead to complete monopolization in *All Chat*? To answer the question, we analyze the content of the chat, employing formal content analysis (see, for example, Charness and Dufwenberg 2006) as well as reproducing quotes from chats (Kimborough, et al. 2012).

7.1. Bilateral Chat

In our first piece of formal content analysis we will try to determine whether the chat contained meaningful information about the contracts that would be offered that round. The relatively high acceptance rates is consistent with players using chat to reach an agreement which U honored. However, the acceptance rate could have increased through other mechanisms including the possibility that chat, while containing no information about the contract, established a rapport between players that increased acceptance rates.

Following Houser and Xiao (2012), we asked two coders to independently analyze the chat content of the *Bilateral Chat*. Specifically, their task was to read the chat in a given round of play in a given market and guess the vector (x_1, T_1, x_2, T_2) that would most likely result from the chat. If they thought that no plausible guess could be made, they were supposed to enter "n.a." instead of a number. The sequence of markets and rounds were randomized such that the coders could not follow patterns involving certain subjects over time. The coders could not see the offers actually made, only the chat content. They had read the instructions up front and were aware of the communication structure in *Bilateral Chat*. Coders analyzed one complete session of *Bilateral Chat* and five random periods from the remaining sessions of that treatment. The coding was incentivized: five chats were randomly selected and the coders paid for the number of guesses that agreed.

The agreement between the two coders was remarkably high. Across all the chats in *Bilateral Chat*, the fraction in which coders agreed ranged between 85% and 92% across the four elements of the contract vector. Cohen's kappa, which reflects the marginal improvement over chance agreement, regarded as a conservative measure of coder agreement, ranges between $\kappa = 0.83$ and $\kappa = 0.89$. We conclude that coders strongly agree on how to code the data.

Are the coders able to predict offers from the chat content? Yes. Table 6 focuses on the the difference between each coder's guess and the actual offered quantity x_i. Coders' guesses are correct (i.e., 0 difference) in a combined $561/680 = 83\%$ of cases in which a numerical guess was made. Only in a combined 20 chats did coders not find informative statements. Including these "n.a." responses, coders' guesses still are correct in 80% of the cases. A similar picture emerges when we look at the tariffs offered. Hence, the content of chat is meaningful even to outside observers.

The result that the coded chat accurately predicts offers made is remarkable and leads us to strongly reject the null hypothesis that chat is meaningless babble. We will thus inquire further into how subjects use this meaningful chat to facilitate contracting.

The hypothesis in *Bilateral Chat* was that subjects may better coordinate and avoid rejections. We specifically hypothesized that the improved coordination may result from the D_i truthfully communicating their beliefs. Of course we should not expect experimental subjects to use formal theoretical terms ("symmetric," "passive") to express their beliefs. However, indirect evidence on beliefs is sometimes available when the D_i has an opportunity to express a preferences among several potential contracts in the chat.

The following example excerpted from an actual chat in the experiment is consistent with the downstream firm's having passive beliefs. Recall upstream firms are called "manufacturers" and downstream "retailers" in the experiment.

```
Retailer:      2 for 20?
Manufacturer:  can be done
Retailer:      that is good
               barely :)
Manufacturer:  hopefully or wait how about 1 for 20 = 30 for you
Retailer:      no 2 for 20
```

With symmetric beliefs, the $(1,20)$ offer would have been better than $(2,20)$ and would have indeed yielded the retailer an expected profit of 30 as the manufacturer suggests. With passive beliefs, on the other hand, $(1,20)$ yields the retailer an expected profit of only 10 while $(2,20)$ yields 16. The fact that the retailer asked for $(2,20)$ is consistent with his or her having passive beliefs.

The following example is consistent with the retailer's having symmetric beliefs.

```
Manufacturer:  2 for 25?
Retailer:      1 for 20
Manufacturer:  agreed
```

With both passive and symmetric beliefs, the $(2,25)$ offer yields the retailer a profit of 11. With symmetric beliefs, the $(1,20)$ offer is better, yielding an expected profit of 30; whereas with passive beliefs, the expected profit would be 10 and the retailer would have chosen $(2,25)$ instead.

Moving to the manufacturer, subjects in this role sometimes make explicit statements about what rivals are offered in chat with retailers. Sometimes this is an attempt to bolster the retailer's belief in symmetry. For example, retailers occasionally asks about a rival's offer, to which the manufacturer responds (honestly or not) "the same." The manufacturer occasionally mentions symmetry without prompting: "I will offer you and the other retailer 1 unit at a price of 30. That is, you have a guaranteed profit of 20 (or even 30)." In this particular round, the manufacturer was honestly trying to establish the monopoly solution with both retailers.

In some cases, the manufacturer does the opposite, trying to persuade retailers about the asymmetry rather than the symmetry of the offer. For example, in the following chat excerpt, the manufacturer attempts to convince each retailer that he or she will become the monopolist at the expense of the other retailer.

```
Manufacturer:  2 units for 65 ECU, special offer just for you?
Retailer 1:    OK.

Manufacturer:  2 units for 60 ECU, special offer just for you?
Retailer 2:    Yes.  The other retailer gets nothing.
Manufacturer:  Good.
Retailer 2:    Stick to it!
```

Interestingly, while the manufacturer's statements are misleading, they are not lies, strictly speaking. The offers are "special" in the sense of specifying different fixed tariffs. (This may explain the odd choice of offering 65 ECU in one and 60 in the other: perhaps this technicality is enough to dispel the manufacturer's psychic costs of lying.) Of course the statements are misleading since a reasonable interpretation of "special" is that the offer is the only one to include a positive quantity. Retailer 2 attempts to confirm this interpretation, receiving the seemingly reassuring but in fact ambiguous response, "good."[11]

To sum up, the analysis of the chat content in the *Bilateral Chat* treatment suggests three conclusions. First, the content is highly meaningful, explaining the high agreement rates. Second, statements by downstream firms sometimes provide evidence of their out-of-equilibrium beliefs, in some cases evidence of passive beliefs and in other cases of symmetric beliefs. Third, some upstream firms use chat to resolve the commitment problem by persuading downstream firms to have more symmetric beliefs; other upstream firms use the secrecy of contracts to deceive downstream firms.

7.2. All Chat

[CCONSIDER PERFORMING CODING ANALYSIS.] Turning to *All Chat*, note up front that the content analysis conducted for *Bilateral Chat* would not be as compelling for this treatment. In nearly 90% of the *All Chat* observations, U offers $x_i = 1$ to each D_i. It is unlikely that the chat would state otherwise, making the coders' task essentially too easy. We can still look at the content of the chat to understand why *All Chat* results in monopolized markets and why downstream firms gain a larger share of industry profits.

The fact that all three players chat together in *All Chat* can have a range of effects. First, it seems unlikely to result in asymmetric offers. For example, if U suggested contracts involving one unit each but different tariffs, the disadvantaged downstream firm would likely request the same reduced tariff. Second, it seems likely to strengthen U's commitment to the monopoly outcome. Because bilateral communication is impossible, downstream firms can see whether side deals are being cut, which may serve to effectively eliminate such side deals involving increased output. In

[11] Note further in this example that the acceptance of the offers is not supported by symmetric, passive, or wary beliefs. For discussion of possible beliefs other than these, see Eguia et al. (2014).

addition, multilateral communication may introduce other social-psychological effects. For example, Fay et al. (2000) suggest that the nature of communication can fundamentally change when the group increases from two to three: rather than addressing an individual, now communication becomes more of a public address.

[WRITE: SYMMETRY IN CHAT CARRIES OVER TO SYMMETRY IN EXPERIMENTAL RESULTS.]

[WRITE: MONOPOLIZATION IN CHAT, DISCUSSING OFFERS OF TWO TOTAL UNITS, CARRIES OVER TO MONOPOLIZATION IN EXPERIMENTAL RESULTS.]

[CONSIDER BEEFING UP SYMMETRY ANALYSIS IN EXPERIMENTAL DATA, ADDING FORMAL TESTS AND DOING FOR FIXED FEES.] Contract offers in *All Chat* are highly symmetric, considerably more so than the other treatments. In 92% of all quantity offers in *All Chat*, $x_1 = x_2$, whereas this is only true in 70% of quantity offers in *No Chat* and 66% of quantity offers in *Bilateral Chat*. Similar results hold comparing the symmetry of the fixed tariffs across treatments. These comparisons are statistically significant in both probit regressions and Mann-Whitney U tests.

[CONSIDER BEEFING UP MONPOLIZATION ANALYSIS IN EXPERIMENTAL DATA] As a result, we overwhelmingly observe $(1, T)$-$(1, T)$ offers discussed in the chat and, in turn, in the decision stage.

The text of the chats provide further evidence of symmetry and monopolization in *All Chat*. The first round of any session is a particularly useful to analyze because subjects' chat is independent of the web of interactions they will have as the experiment progresses. We find that of the 24 groups that formed markets in the first round of an *All Chat* session, 17 made clear statements about symmetry or joint maximization. In *Bilateral Chat*, such statements were made by only 9 of 24 groups in the first round, significantly fewer than in *All Chat* according to a two-tailed Fisher exact test ($p = 0.041$). Indeed, the very first statement made by any subject in *All Chat* was a retailer asking a manufacturer, "How much will you offer to both of us?" Furthermore, manufacturers often make statements like "one unit for 20 to either retailer" and often use the words "we" and "our" as in, for example, "Our best position is in the market are two units with a turnover of 100."

[THINK MORE ABOUT THIS.] One could counter that all this is still cheap talk and, with contracts being secret, U players have little incentive to stick to the promises made in that chat.

Once the chat is over, the commitment problem and the incentive to deviate from the agreed-upon offers would still be present according to this argument. Note, however, how difficult it would be for U to deviate from the promises made in the chat. Suppose that, following a (1, T) and (1, T) promise in the chat, actual offers are (1, T_1) and (2, T_2). Both a D_2 with initially passive and symmetric believers will conjecture that D_1 also obtained and offer with two units and U cannot talk any more to convince D_2 otherwise. As a result, U would need to request $T_2 < T$, so the deviation (from the chat promise) does not pay. On top of that, D_2 might now be puzzled and distrust the deviation offer, possible resulting in a rejection. Of course, theory still maintains that a D_1 with passive beliefs should reject the (1, T) offer, but D_1 can also follow the reasoning that a deviation offer to D_2 is unlikely to occur and hence accept (1, T). Effectively, the symmetric promises made in the chat serve to establish symmetric beliefs held by the D firms, accordingly, monopolization of the market follows.

8. Conclusion

[TBD] The commitment problem arising when an upstream monopolist supplies several downstream firms (Hart and Tirole, 1900; Rey and Tirole, 2007) is a challenging issue. The firms can jointly earn monopoly rents but, when downstream firms hold the "wrong" beliefs, firms will fail to do so. On top of that, beliefs may be heterogeneous in the field or the laboratory which may cause upstream firms to offer contracts that are rejected with positive probability in equilibrium, implying inefficiencies. Finally, bargaining frictions may occur, suggesting further rejections and smaller industry profits.

We run an experiment on this problem building on previous work by Martin, Normann and Snyder (2001) and introducing cheap talk to this framework. First, we allow for bilateral communication between the upstream firm and the downstream firms, each downstream firm in a separate channel (this form of communication would naturally occur in negotiations the field). Second, all firms in the industry can chat. We compare these treatment involving communication to one without chat.

We find support for the hypothesis that communication results in higher profits and lower rejections rates. How firms communicate is important, though. Bilateral communication improves joint payoffs even though aggregate output is not lower than without communication. When all firms in the industry talk, this effectively leads to a monopolization of the market. This monopolization comes at a cost for the upstream firm, though: downstream firms bargain tougher in this variant, so the upstream firms' profit does not increase.

The conclusion from this is that antitrust policy should be suspicious about communication also in vertically related industries—at least in markets prone to the commitment problem. The problem is not so much that downstream firms engage in explicit horizontal collude here (in fact, they cannot in this model). Neither do they use the upstream firms to establish a hub-and-spoke cartel. Rather, upstream firms manage to overcome the commitment problem with this form of communication. Put differently, the effect of communication is not so much that of a cartel, but more resembling a vertical restraint that causes anticompetitive effects. This negative effect is similar to those resulting from vertical mergers, public contracts or non-discrimination clauses.

Unrestricted communication in particular seems to erode its take-it-or-leave-it bargaining power enough that it would prefer no communication to a meeting among all the firms in the proverbially "smoke-filled room." Thus U would Meetings in among firms in a vertical market may be less cause for antitrust concern This may mitigate any antitrust concern about the anticompetitive Antitrust officials Thus communication in a vertical market raises antitrust concerns, its increased power to commit to the mmcommitment the increased profits from

References

Basayan, Narine, Jacob K. Goeree, Monica Hartmann, John Morgan, Tanya Rosenblat, Maros Servatka, and Dirk Yandell. (2009) "Vertical Integration of Successive Monopolists: A Classroom Experiment," *Perspectives on Economic Education Research* 5(1).

Boone, Jan, Wieland Müller, and Sigrid Suetens. (2014) "Naked Exclusion in the Lab: The Case of Sequential Contracting," *Journal of Industrial Economics* 62: 137–166.

Brosig, Jeanette, Joachim Weimann, and C.-L. Yang. (2001) "Communication, Reputation, and Punishment in Simple Sequential Bargaining Experiments," *Journal of Institutional and Theoretical Economics* 160: 576–606.

Crawford, Vincent. (1998) "A Survey of Experiments on Communication via Cheap Talk," *Journal of Economic Theory* 78: 286–298.

Durham, Yvonne. (2000) "An Experimental Examination of Double Marginalization and Vertical Relationships," *Journal of Economic Behavior and Organization* 42: 207–229.

Eguia, Jon X., Aniol Llorente-Saguer, Rebecca Morton, and Antonio Nicolò. (2014) "Equilibrium Selection in Sequential Games with Imperfect Information," Max Planck Institute for Research on Collective Goods working paper no. 2014–04.

Farrell, Joseph and Matthew Rabin. (1996) "Cheap Talk," *Journal of Economic Perspectives* 10: 103–118.

Fischbacher, Urs. (2007) "Z-Tree: Zurich Toolbox for Readymade Economic Experiments," *Experimental Economics* 10: 171–178.

Greiner, Ben. (2004) "The Online Recruitment System ORSEE 2.0: A Guide for the Organization of Experiments in Economics," University of Cologne Economics Department working paper no. 10.

Fay, Nicolas, Simon Garrod, and Jean Carletta. (2000) "Group Discussion as Interactive Dialogue or as Serial Monologue: The Influence of Group Size," *Psychological Science* 11: 481–486.

Fonseca, Miguel A. and Hans-Theo Normann. (2012) "Explicit vs. Tacit Collusion? The Impact of Communication in Oligopoly Experiments," *European Economic Review* 56: 1759–1772.

Huck, Steffen, Hans-Theo Normann, and Jörg Oechssler. (2004) "Two are Few and Four are Many: On Number Effects in Cournot Oligopoly," *Journal of Economic Behavior and Organization* 53: 435–446.

Hart, Oliver and Jean Tirole. (1990) "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity (Microeconomics)* 205–285.

Isaac, R. Mark, Valerie Ramey, and Arlington W. Williams. (1984) "The Effects of Market Organization on Conspiracies in Restraint of Trade," *Journal of Economic Behavior and Organization* 5: 191–222.

Landeo, Claudia M. and Kathryn E. Spier. (2009) "Naked Exclusion: An Experimental Study of Contracts With Externalities," *American Economic Review* 99: 1850–1877.

Martin, Stephen, Hans-Theo Normann, and Christopher M. Snyder. (2001) "Vertical Foreclosure in Experimental Markets," *Rand Journal of Economics* 32: 466–496.

Mason, Charles F. and Owen R. Phillips. (2000) "Vertical Integration and Collusive Incentives," *International Journal of Industrial Organization* 18: 471–496.

McAfee, R. Preston and Marius Schwartz. (1994) "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity," *American Economic Review* 84: 210–230.

Normann, Hans-Theo. (2011) "Do Vertical Mergers Enable to Raise Rivals' Costs? Experimental Evidence," *Journal of Industrial Economics* 59: 506–527.

O'Brien, Daniel P. and Greg Shaffer. (1992) "Vertical Control with Bilateral Contracts," *Rand Journal of Economics* 23: 299–308.

Radner, Roy and Andrew Schotter. (1989) "The Sealed-Bid Mechanism: An Experimental Study," *Journal of Economic Theory* 48: 179–220.

Rasmusen, Eric B., J. Mark Ramseyer, and John S. Wiley, Jr. (1991) "Naked Exclusion," *American Economic Review* 81: 1137–1145.

Rey, Patrick and Jean Tirole. (2003) "A Primer on Foreclosure," in Mark Armstrong and Robert Porter, eds., *Handbook of Industrial Organization*, vol. 3, North-Holland, 2145–2220.

Rey, Partick and Thibaud Vergé. (2004) "Bilateral Control with Vertical Contracts," *Rand Journal of Economics* 35: 728–746.

Roth, Alvin E. "Bargaining Experiments," in John Kagel and Alvin E. Roth, eds., *Handbook of Experimental Economics*, Princeton University Press, 253–348.

Roth, Alvin E. and J. Keith Murnighan. (1982) "The Role of Information in Bargaining: An Experimental Study," *Econometrica* 50: 1123–1142.

Segal, Ilya and Michael D. Whinston. (2000) "Naked Exclusion: Comment," *American Economic Review* 90: 296–309.

Smith, Angela M. (2011) "An Experimental Study of Exclusive Contracts," *International Journal of Industrial Organization* 29: 4–13.

White, Halbert. (1980) "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica* 48: 817–838.

Zultan, Ro'i. (2012) "Strategic and Social Pre-Play Communication in the Ultimatum Game," *Journal of Economic Psychology* 33: 425–434.

Table 1: Treatment Specifications

	No Chat	*Bilateral Chat*	*All Chat*
Communication	None	U with D_1, U with D_2	All three jointly
Number of subjects	75	69	72
Number of sessions	4	4	4

Table 2: Summary Statistics for Experimental Results

Treatment	X	T_i	a_i	Q	Π	π_U	π_D
No Chat	3.62	32.9	0.72	2.54	70.3	45.3	25.0
$N = 250$	(1.45)	(10.5)	(0.45)	(1.52)	(32.4)	(22.9)	(32.3)
Bilateral Chat	2.80	30.4	0.85	2.37	84.3	49.9	34.5
$N = 230$	(1.12)	(13.2)	(0.36)	(1.10)	(23.1)	(18.9)	(23.2)
All Chat	2.18	22.5	0.92	1.98	92.4	40.8	51.6
$N = 240$	(0.58)	(5.8)	(0.27)	(0.64)	(19.4)	(12.0)	(16.4)
Combined	2.88	28.7	0.83	2.30	82.2	45.3	36.9
$N = 720$	(1.26)	(11.2)	(0.38)	(1.17)	(27.3)	(18.9)	(27.3)

Notes: Entries are means; standard deviations reported in parentheses. Sample excludes the first five rounds of each session. Number of observations for the market-level variables (X, Q, Π, π_U, π_D) listed below row heading. Firm-level variables (T_i and a_i) have twice as many observations, one for each of two firms.

Table 3: Regressions Examining Differences in Means

	Dependent variable								
	X	T_i	T_i	a_i	a_i	Q	Π	π_U	π_D
Regressors									
No Chat	3.62***	32.9***	29.1***	0.72***	0.92***	2.54***	70.3***	45.3***	25.0***
	(0.15)	(0.4)	(1.4)	(0.02)	(0.07)	(0.19)	(1.3)	(1.7)	(2.7)
Bilateral Chat	2.80***	30.4***	28.7***	0.85***	1.00***	2.37***	84.3***	49.9***	34.5***
	(0.17)	(2.3)	(2.1)	(0.04)	(0.07)	(0.12)	(2.4)	(2.7)	(3.1)
All Chat	2.18***	22.5***	22.2***	0.92***	0.95***	1.98***	92.4***	40.8***	51.6***
	(0.03)	(0.8)	(0.7)	(0.00)	(0.07)	(0.03)	(0.5)	(1.0)	(1.4)
Other controls	None	None	$\{1_{x_i}\}$	None	$\{1_{x_i}, T_i\}$	None	None	None	None
Observations	720	1,440	1,440	1,440	1,440	720	720	720	720
R^2	0.22	0.16	0.30	0.05	0.20	0.04	0.12	0.04	0.17
Coefficient differences									
Bilateral Chat–	–0.81***	–2.5	–0.4	0.13***	0.08**	–0.16	14.1***	4.6	9.5**
No Chat	(0.23)	(2.4)	(2.7)	(0.04)	(0.03)	(0.22)	(2.8)	(3.2)	(4.1)
All Chat–	–0.63***	–7.9***	–6.5**	0.07*	–0.05**	–0.39***	8.1***	–9.1***	17.2***
Bilateral Chat	(0.17)	(2.4)	(2.3)	(0.04)	(0.02)	(0.12)	(2.5)	(2.9)	(3.4)
All Chat–	–1.44***	–10.4***	–6.9***	0.20***	0.02	–0.56***	22.2***	–4.5**	26.6***
No Chat	(0.15)	(0.9)	(1.5)	(0.02)	(0.02)	(0.19)	(1.4)	(1.9)	(3.0)

Notes: Each column is an ordinary least squares regression. Exhaustive set of treatment indicators included as regressors and constant omitted. Sample excludes the first five rounds of each session. White (1980) heteroskedasticity-robust standard errors clustered at session level reported in parentheses. The regression for T_i with other controls includes fixed effects for each integer value of x_i. We omit the $x_i = 1$ fixed effect, giving the coefficients on the treatment indicators the interpretation of mean tariffs for contracts offering one unit of quantity. The regression for a_i with other controls includes fixed effects for each integer value of x_i, a linear and squared T_i, and interactions between the fixed effects for x_i and the linear and squared T_i. We omit the $x_i = 1$ fixed effect, giving the coefficients on the treatment indicators the interpretation of mean acceptance rates for contracts offering one unit and zero tariff. Significantly different from 0 in a two-tailed test at the *10% level, **5% level, ***1% level.

Table 4: Regressions Examining Shifts in Quantity Histograms

	Offered quantity		Market quantity	
	$X = 2$	$X \geq 4$	$Q = 2$	$Q \geq 4$
Regressors				
No Chat	0.30***	0.56***	0.32***	0.32***
	(0.08)	(0.08)	(0.05)	(0.08)
Bilateral Chat	0.53***	0.25***	0.50***	0.17***
	(0.09)	(0.07)	(0.06)	(0.05)
All Chat	0.88***	0.06***	0.81***	0.03**
	(0.02)	(0.02)	(0.02)	(0.01)
Observations	1,080	1,080	1,080	1,080
R^2	0.17	0.13	0.13	0.06
Coefficient differences				
Bilateral Chat– *No Chat*	0.23*	–0.31**	0.18**	–0.14
	(0.12)	(0.11)	(0.08)	(0.10)
All Chat– *Bilateral Chat*	0.34***	–0.19**	0.31***	–0.14**
	(0.09)	(0.07)	(0.07)	(0.05)
All Chat– *No Chat*	0.58***	–0.50***	0.49**	–0.28***
	(0.08)	(0.08)	(0.05)	(0.09)

Notes: Each column is an ordinary least squares regression in which the dependent variable is a 0–1 indicator for the event in the column heading. Regression thus interpreted as linear probability model. Omitted treatment is *No Chat*. Sample excludes the first five rounds of each session. White (1980) heteroskedasticity-robust standard errors clustered at session level reported in parentheses. Significantly different from 0 in a two-tailed test at the *10% level, **5% level, ***1% level.

Table 5: Trends in Treatment Effects

	Dependent variable						
	X	T_i	a_i	Q	Π	π_U	π_D
Regressors							
No Chat$_0$	3.80***	38.2***	0.66***	2.33***	64.4***	45.5***	19.0***
	(0.31)	(3.1)	(0.01)	(0.19)	(4.1)	(3.0)	(2.3)
No Chat$_1$	3.62***	32.9***	0.72***	2.54***	70.3***	45.3***	25.0***
	(0.15)	(0.4)	(0.02)	(0.19)	(1.3)	(1.7)	(2.7)
Bilateral Chat$_0$	3.32***	33.5***	0.84***	2.72***	78.9***	53.5***	25.4***
	(0.25)	(2.2)	(0.04)	(0.16)	(3.4)	(2.6)	(2.9)
Bilateral Chat$_1$	2.80***	30.4***	0.85***	2.37***	84.3***	49.9***	34.5***
	(0.17)	(2.3)	(0.04)	(0.12)	(2.4)	(2.7)	(2.1)
All Chat$_0$	2.86***	30.0***	0.83***	2.20***	83.7***	45.7***	37.9***
	(0.25)	(2.3)	(0.22)	(0.10)	(2.2)	(3.0)	(3.5)
All Chat$_1$	2.18***	22.5***	0.92***	1.98***	92.4***	40.8***	51.6***
	(0.03)	(0.8)	(0.00)	(0.03)	(0.5)	(1.0)	(1.4)
Observations	1,080	2,160	2,160	1,080	1,080	1,080	1,080
R^2	0.14	0.13	0.05	0.03	0.11	0.04	0.15
Coefficient differences							
No Chat$_1$−	–0.18	–5.3	0.07***	0.20***	5.8	–0.2	6.0
No Chat$_0$	(0.19)	(3.0)	(0.02)	(0.05)	(4.0)	(4.6)	(3.8)
Bilateral Chat$_1$−	–0.52***	–3.1***	0.01	–0.35***	5.5***	–3.6**	9.1***
Bilateral Chat$_0$	(0.13)	(0.4)	(0.01)	(0.11)	(1.4)	(1.6)	(0.4)
All Chat$_1$−	–0.68***	–7.5***	0.09***	–0.22*	8.8***	–4.9*	13.7***
All Chat$_0$	(0.27)	(0.9)	(0.02)	(0.12)	(1.8)	(2.6)	(2.4)

Notes: Each column is an ordinary least squares regression. Subscripts denote interactions with period indicators, e.g., *No Chat*$_0$ is the interaction between *No Chat* and the initial period consisting of rounds 1–5 and *No Chat*$_1$ is the interaction between *No Chat* and the end period consisting of rounds 6–15. Exhaustive set of treatment indicators included as regressors and constant omitted. Sample includes all 15 rounds. None of the regressions include additional controls for regressors. White (1980) heteroskedasticity-robust standard errors clustered at session level reported in parentheses. Significantly different from 0 in a two-tailed test at the *10% level, **5% level, ***1% level.

Table 6: Differences Between Coders' Guess and Offered Quantity x_i.

	–3 or less	–2	–1	0	1	2	3 or more	n.a.	Total
Coder 1	4	5	29	282	18	4	1	7	350
Coder 2	1	7	28	279	21	0	1	13	350
Total	5	12	57	561	39	4	2	20	700

Appendix: 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 euro for every 40 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.

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.

Note that in every period the manufacturer-retailer groups change. You do not know which retailer or manufacturer you will meet.

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 $\times$ 30 ECU = 60 ECU, while retailer 1 has revenues of 1 $\times$ 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.

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.

Appendix: Supporting Exhibits (Not for Publication)

Table A1: Summary Statistics for Experimental Results, All Rounds

Treatment	X	T_i	a_i	Q	Π	π_U	π_D
No Chat	3.68	34.7	0.70	2.47	68.3	45.3	23.0
$N = 375$	(1.88)	(13.1)	(0.46)	(1.19)	(33.7)	(24.7)	(32.1)
Bilateral Chat	2.98	31.4	0.85	2.49	82.5	51.1	31.4
$N = 345$	(1.29)	(14.3)	(0.36)	(1.19)	(24.6)	(20.2)	(28.0)
All Chat	2.41	25.0	0.89	2.05	89.5	42.5	47.1
$N = 360$	(1.01)	(10.1)	(0.31)	(0.82)	(22.4)	(15.5)	(20.3)
Combined	3.03	30.4	0.81	2.34	79.9	46.2	33.7
$N = 1,080$	(1.54)	(13.2)	(0.39)	(1.28)	(28.9)	(20.8)	(29.1)

Notes: Sample includes all 15 rounds each session. Otherwise analogous to Table 2. See notes to that table.

Table A2: Regressions Examining Differences in Means, All Rounds

	Dependent variable								
	X	T_i	T_i	a_i	a_i	Q	Π	π_U	π_D
Regressors									
No Chat	3.68***	34.7***	30.4***	0.70***	1.23***	2.47***	68.3***	45.3***	23.0***
	(0.19)	(1.2)	(1.5)	(0.02)	(0.09)	(0.19)	(1.8)	(0.5)	(1.8)
Bilateral Chat	2.98***	31.4***	28.9***	0.85***	1.33***	2.49***	82.5***	51.1***	31.4***
	(0.19)	(2.3)	(2.0)	(0.04)	(0.10)	(0.12)	(2.7)	(2.6)	(3.0)
All Chat	2.41***	25.0***	23.9***	0.89***	1.29***	2.05***	89.5***	42.5***	47.1***
	(0.06)	(1.0)	(1.1)	(0.01)	(0.09)	(0.03)	(1.0)	(1.5)	(2.1)
Other controls	None	None	$\{1_{x_i}\}$	None	$\{1_{x_i}, T_i\}$	None	None	None	None
Observations	1,080	2,160	2,160	2,160	2,160	1,080	1,080	1,080	1,080
R^2	0.12	0.09	0.22	0.05	0.20	0.02	0.10	0.03	0.12
Coefficient differences									
Bilateral Chat–	–0.70***	–3.2	–1.6	0.15***	0.11***	0.02	14.2***	5.8**	8.9**
No Chat	(0.27)	(2.6)	(2.6)	(0.04)	(0.03)	(0.23)	(3.2)	(3.2)	(3.5)
All Chat–	–0.57**	–6.4**	–4.9*	0.04	–0.05*	–0.44***	7.0**	–8.6**	15.6***
Bilateral Chat	(0.20)	(2.5)	(2.4)	(0.04)	(0.03)	(0.13)	(2.9)	(3.0)	(3.6)
All Chat–	–1.27***	–9.7***	–6.5***	0.20***	0.06**	–0.42*	21.2***	–2.9*	24.1***
No Chat	(0.20)	(1.5)	(1.9)	(0.02)	(0.02)	(0.19)	(2.0)	(1.6)	(2.7)

Notes: Sample includes all 15 rounds each session. Otherwise analogous to Table 3. See notes to that table.

Table A3: Regressions Examining Shifts in Quantity Histograms, All Rounds

	Offered quantity		Market quantity	
	$X = 2$	$X \geq 4$	$Q = 2$	$Q \geq 4$
Regressors				
No Chat	0.30***	0.54***	0.30***	0.29***
	(0.06)	(0.06)	(0.04)	(0.07)
Bilateral Chat	0.48***	0.30***	0.47***	0.21***
	(0.10)	(0.07)	(0.07)	(0.05)
All Chat	0.79***	0.13***	0.73***	0.07**
	(0.02)	(0.02)	(0.02)	(0.01)
Observations	1,080	1,080	1,080	1,080
R^2	0.17	0.13	0.13	0.06
Coefficient differences				
Bilateral Chat–	0.18	–0.23**	0.17*	–0.08
No Chat	(0.12)	(0.09)	(0.08)	(0.09)
All Chat–	0.32***	–0.17**	0.26***	–0.14**
Bilateral Chat	(0.10)	(0.07)	(0.07)	(0.05)
All Chat–	0.50***	–0.41***	0.43**	–0.22***
No Chat	(0.07)	(0.06)	(0.05)	(0.04)

Notes: Sample includes all 15 rounds each session. Otherwise analogous to Table 4. See notes to that table.