

NBER WORKING PAPER SERIES

COMMUNICATION IN VERTICAL MARKETS:
EXPERIMENTAL EVIDENCE

Claudia Möllers
Hans-Theo Normann
Christopher M. Snyder

Working Paper 22219
<http://www.nber.org/papers/w22219>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2016

The authors are grateful to Maura Doyle, Ali Hortacsu, Francisco Martinez, Patrick Rey and conference participants at CISS Turunc, DFG-ANR Rennes, EARIE Milan, IIOC Boston, LEW London, VfS Berlin (IO Committee) and VfS Muenster for helpful comments. We are grateful to Deutsche Forschungsgemeinschaft (grant DFG NO 429/5-1) for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Claudia Möllers, Hans-Theo Normann, and Christopher M. Snyder. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Communication in Vertical Markets: Experimental Evidence
Claudia Möllers, Hans-Theo Normann, and Christopher M. Snyder
NBER Working Paper No. 22219
May 2016
JEL No. C70,C90,K21,L42

ABSTRACT

When an upstream monopolist supplies several competing downstream firms, it may fail to monopolize the market because it is unable to commit not to behave opportunistically. We build on previous experimental studies of this well-known commitment problem by introducing communication. Allowing the upstream firm to chat privately with each downstream firm reduces total offered quantity from near the Cournot level (observed in the absence of communication) halfway toward the monopoly level. Allowing all three firms to chat together openly results in complete monopolization. Downstream firms obtain such a bargaining advantage from open communication that all of the gains from monopolizing the market accrue to them. A simple structural model of Nash bargaining fits the pattern of shifting surpluses well. We conclude with a discussion of the antitrust implications of open communication in vertical markets.

Claudia Möllers
University of Duesseldorf
Universitätsstr. 1
40225 Düsseldorf Germany
moellers@dice.hhu.de

Hans-Theo Normann
University of Duesseldorf
Universitätsstr. 1
40225 Düsseldorf Germany
normann@dice.hhu.de

Christopher M. Snyder
Department of Economics
Dartmouth College
301 Rockefeller Hall
Hanover, NH 03755
and NBER
chris.snyder@dartmouth.edu

1. Introduction

Whether vertical mergers can have anticompetitive effects remains a central question in the largest antitrust cases. For example, in January 2011, the U.S. Department of Justice applied the “most intense scrutiny ever for a planned media merger” before approving the takeover of NBC Universal (an upstream content provider) by Comcast (a downstream cable distributor) subject to a list of conditions (Arango and Stelter 2011). In April 2015, the European Competition Commission charged Google with the violation of favoring its affiliates over competitors in search displays (Kanter and Scott 2015).

An influential strand of the theoretical literature (summarized in Rey and Tirole 2007) connects the anticompetitive effects of vertical restraints to their ability to solve a commitment problem. An upstream monopolist serving downstream competitors might wish to offer contracts restricting output to the joint-profit maximum. It may fail to do so, however, because it has an incentive to behave opportunistically, offering one of the downstream firms a contract increasing their bilateral profits at the expense of all other downstream firms (the same logic extending to the bilateral contract with each downstream firm). In Hart and Tirole (1990), a vertical merger helps to solve this commitment problem by removing its incentive to behave opportunistically in a way that would harm the downstream unit with which it shares profits. While the upstream firm benefits from solving the commitment problem, overall the vertical merger has an anticompetitive effect on the market because prices rise and output falls. Similar anticompetitive effects can arise with vertical restraints aside from mergers including resale price maintenance (O’Brien and Shaffer 1992, Rey and Vergé 2004) and non-discrimination clauses (McAfee and Schwartz 1994).

The commitment problem is a somewhat delicate theoretical proposition. Depending on downstream firms’ beliefs after receiving a deviating secret contract offer—not pinned down in a perfect Bayesian equilibrium—there can be multiple equilibria, with the commitment effect arising in some and not in others (McAfee and Schwartz 1994, and Rey and Vergé 2004). With *symmetric beliefs*, downstream firms reject deviating contracts generating negative profits for rivals because they infer that rivals received the same deviating contract. In this way, symmetric beliefs afford the upstream firm the ability to commit to monopolizing the market. With *passive beliefs*, on the other hand, deviation does not change downstream firms beliefs, increasing their willingness to accept deviating contracts, impairing the upstream firm’s commitment power.

In the absence of a widely accepted refinement of perfect Bayesian equilibrium providing a firm theoretical foundation for selecting one or another equilibrium in this context, Martin, Normann, and Snyder (2001) turned to experiments to gauge the significance of the commitment problem. In their baseline treatment in which an upstream monopolist makes secret offers of nonlinear tariffs to two downstream firms, labeled *SECRAN*, they found that markets were rarely monopolized; industry profits averaged only two thirds of the joint maximum. By contrast, markets were regularly monopolized when either the upstream monopoly was vertically integrated with a downstream firm or when contracts were public. The experiments thus support the view that the commitment problem is genuine.

In this paper, we return to an experimental study of vertical markets with a new focus—on whether allowing firms to communicate can help them solve the commitment problem without resorting to vertical restraints. For the sake of comparison, we start with the same *SECRAN* treatment as Martin, Normann, and Snyder (2001). In addition to this baseline treatment without communication, we run a series of three treatments in which players can communicate whatever messages they want via a messenger-like tool. The communication treatments involve different levels of openness. One allows the upstream firm to engage in private two-way chat with each downstream firm. Another allows all three firms to engage in completely open (three-way) chat. A third is a hybrid of the other two, allowing players the option of using either or both of two- or three-way communication.

Communication is cheap talk in our experiments, so standard results (Crawford and Sobel 1982) leave open the possibility that adding this form of communication may have no effect on equilibrium. Yet we have a number of good reasons to believe communication might have real effects in our experiments. First, the vertical contracting game involves considerable strategic uncertainty. A downstream firm has to form out-of-equilibrium beliefs and other firms have to conjecture what this belief is (or what its distribution is if beliefs are heterogeneous). Communication could resolve some of this strategic uncertainty, perhaps by allowing downstream firms to communicate their beliefs to the upstream firm so that the offer can be tailored to those beliefs, perhaps by allowing the upstream firm to persuade downstream firms to hold more propitious beliefs (in particular the symmetric beliefs mentioned previously as leading to joint monopolization of the market). Second, communication could help solve the commitment problem by allowing

the upstream firm to make promises. Promises about rival contracts are not legally enforceable in our experiments but could still afford some commitment power if making a bald-faced lie involves a substantial psychological cost. Third, communication has been shown in previous experiments to reduce bargaining frictions (Roth 1995). On the other hand, communication could conceivably work in the opposite direction, impairing commitment. A conspiracy between the upstream and a downstream firm to deviate to a contract increasing their bilateral profits at the expense of the downstream rival would be easier to hatch if they could communicate privately. Of course, open communication precludes conspiracy, so open communication should either aid commitment or at worst have no effect. When firms are given the option of using either private or open communication, whether or not they are tempted to conspire, undermining commitment, is an interesting empirical question, which can be addressed by the hybrid treatment.

Along with the theoretical motives we just described for studying the effects of communication, we also have practical policy motives. Communication between vertically related firms is presumably the rule rather than the exception in the field, so introducing it in the lab adds an important practical element to existing experiments. While a conversation between an upstream and a downstream firm would not violate antitrust law, communication in an open forum involving horizontally along with vertically related firms might raise antitrust concerns. Whether such communication has the potential to restrain competition has so far not been studied.

Our experimental results reveal a remarkably consistent pattern: increasing the openness of communication has a monotonic effect across virtually every market outcome and treatment we study. In the treatment without communication, the same severe commitment problem observed in Martin, Normann, and Snyder (2001) occurs: aggregate offered quantity is again much closer to the Cournot than the monopoly level. Two-way communication mitigates but does not solve the commitment problem, cutting the distance between aggregate offered quantity and the monopoly quantity about in half. Three-way communication cuts the remaining distance again in half, resulting in nearly complete monopolization of the market, particularly in the late rounds of play. Results for the hybrid treatment are between the other two, somewhat closer to the treatment with open communication. Further, we find that more open communication leads to more fluid bargaining, captured by an increasing rate of contract acceptance. The increase in acceptance rate, due in part to increasing confidence in the upstream firm's commitment to monopolize the market, is

also due in part to a reduction in the upstream firm’s tariff demands. Overall, the increase in acceptance rates offsets the reduction in tariffs, so consequently upstream profits remain essentially unchanged across treatments. The gains from using more open communication to better monopolize the market thus accrue almost entirely to downstream firms. Compared to the treatment without communication, completely open communication more than doubles downstream profits.

From a policy perspective, our results imply that some forms of communication can effectively function as an anticompetitive vertical restraint. In particular, allowing an upstream firm to discuss contracts with several downstream firms in a “smoke-filled room” (or simply to exchange public pronouncements) has the potential to substantially restrict output. On the other hand, if firms already have such forums for open communication, vertical mergers and restraints themselves may not raise further antitrust concerns.

Regarding its relationship to the literature, our paper is the first experimental study of communication in a vertically related market. Our paper is closest to the one on which we build, Martin, Normann, and Snyder (2001), which provides an experimental test of the theories of anticompetitive vertical restraints (vertical mergers, public contracts) put forth by the papers mentioned earlier (Hart and Tirole 1990, O’Brien and Shaffer 1992, McAfee and Schwartz 1994, Rey and Vergé 2004, Rey and Tirole 2007; see Avenel 2012 and Rey and Caprice 2015 for more recent developments). Other experiments in vertically related markets include Mason and Phillips’ (2000) study of equilibrium when the upstream input is demanded by a Cournot duopoly in one market and perfectly competitive firms in another. Durham (2000) and Badasyan et al. (2009) analyze whether vertical merger mitigates the double-marginalization problem. Normann (2011) investigates whether vertical merger has an anticompetitive “raising rivals’ cost” effect in a bilateral duopoly. None of these papers studies communication, the focus of the present paper.

Also related is the experimental literature on exclusive dealing (Landeo and Spier 2009, Smith 2011, Boone, Müller, and Suetens 2014). As in our setting, in this literature the vertical contract exerts an externality on other downstream firms. The nature of the externality is different; rather than secretly oversupplying a rival—not an issue because downstream firms are assumed to operate in different final-good markets—an initial exclusive contract diverts demand that would otherwise prompt a more efficient upstream firm to enter, which then would supply other downstream firms at lower prices. Landeo and Spier (2009) and Smith (2011) show that communication between down-

stream firms reduces entry-detering exclusion by helping them coordinate on a strategy of waiting for the efficient entrant. Communication may also serve a coordination role in our setting but the effect on social welfare is the opposite, anticompetitive in our setting rather than procompetitive.

Our paper contributes to a large literature on cheap talk in experimental games. Theory suggests that potential gains from cheap talk are greatest in games of common rather than conflicting interests (Farrell and Rabin 1996). Consistent with theory, experiments find large gains from cheap talk in coordination games (see Crawford 1998 for a survey). However, cheap talk also increases the rate of cooperation in dilemma games (Dawes, McTavish and Shaklee 1977, Isaac, Ramey and Williams 1984, Balliet 2010) in which neoclassical theory would suggest agreements to cooperate should be worthless. Our result that communication aids monopolization has a similar flavor, although decision making is more complex in our setting: final output is the result of a negotiation between upstream and downstream firms rather than being one firm's unilateral choice.¹

Cheap talk has been found to achieve superior outcomes in trust games (Charness and Dufwenberg 2006). Our vertically related markets also have an element of trust: accepting a contract offer may only be profitable if the downstream firm trusts the upstream firm's promise (implicit or explicit) to restrict output traded to the rival firm. Our setting is different from the standard trust game, however. Whereas in a typical two-player trust game the proposer needs to trust the responder (for example to return money), in our setting the responder (downstream firm) needs to trust the proposer (upstream firm). Promises relate not to the upstream firm's future action but to the offer made simultaneously to the rival downstream firm (who is a trustor at the same time).

Within the literature on cheap talk in experimental games, ours is closest to studies of the effect of cheap talk on bargaining. Adding a round of face-to-face communication before offers are made results in near perfect rates of agreement (Roth 1995). Typed messages—the sort of communication also used in our experiments—does not improve efficiency as much but still improves upon no communication (Brosig, Ockenfels and Weimann 2003, Andersson et al. 2010, Zultan 2012).

¹Several experimental industrial organization papers have the flavor of communication in a dilemma game. Anderson and Wengström (2007) analyze costly communication in Bertrand duopoly, finding that prices are higher and collusion more stable when communication is costly. Hinlopen and Soetevent (2008) and Bigoni et al. (2012) evaluate lenience programs in laboratory experiments with communication. 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 is credible. They analyze what type of communication is most effective in achieving cooperation in this setup.

Ours is the first to study how cheap talk between vertically related players affects bargaining with externalities. In this setting, the openness of communication becomes an important treatment variable. We find that private communication improves efficiency somewhat and open communication still more, reaching 92% agreement rates.

2. Theoretical Framework

2.1. Model

Consider a simplified version of the model due to Rey and Tirole (2007).² 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.³

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 ($a_i = 0$) 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 $Q^m = \operatorname{argmax}_Q P(Q)Q$ 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

$$BR(q) = \operatorname{argmax}_{\tilde{q}} P(\tilde{q} + q)\tilde{q},$$

²Rey 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 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.

³Assume $P(Q)$ 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 .

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.

2.2. Commitment Problem with Secret Contracts

To understand the nature of the commitment problem with secret contracts, suppose first that contracts are public, meaning that each D_i can see the contract offered to its rival. If so, U can extract the monopoly profit in equilibrium. For example, by offering the contract $(Q^m/2, \Pi^m/2)$ to each D_i . The D_i earn zero profit whether or not they accept so they accept in equilibrium.

Secret contracts transform the model into a dynamic game of imperfect information. The relevant solution concept is perfect Bayesian equilibrium, requiring strategies to be best responses 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, and 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, that is, 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 $P(2x^d)x^d$, which is obviously no greater than the fixed fee $\Pi^m/2$ that U charged in the equilibrium contract.⁴

Another assumption, called *passive beliefs*, is that after receiving a deviating offer, D_i continues to believe its rival receives the equilibrium contract. These beliefs make deviation particularly attractive, rendering the monopoly outcome unstable. Formally, there will always exist a strictly profitable deviation unless equilibrium firm quantity q^* is best response to itself, that is, $q^* = BR(q^*)$. But as we saw above, the Cournot output q^c is the unique quantity satisfying this equation. Hence the equilibrium contract offer is (q^c, π^c) , which both D_i accept. Here we see the commitment problem: if the D_i have passive beliefs, U cannot restrict output to the monopoly level despite being an upstream monopolist.⁵

⁴ D_i would reject a tariff greater than $P(2x^d)x^d$ if it believes D_i 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.

⁵While symmetric and passive beliefs are the main cases typically studied, other beliefs are possible. McAfee and Schwartz (1994) proposed *wary beliefs*, 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. In

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. Each D_i holds symmetric beliefs with probability $s \in [0, 1]$ and passive beliefs with $1 - s$. 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 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.

2.3. Communication and the Commitment Problem

We modify the benchmark model by adding a communication stage prior to contract offers. Since this is just cheap talk, it is always possible that communication—whether between two or among all three parties—changes nothing. The outcome of the communication stage can always be a babbling equilibrium with completely uninformative communication.

On the other hand, it is conceivable that communication could enhance U 's commitment power. In two-way communication, D_i could extract a promise from U not to oversupply its rival. While this would be an empty promise coming from a neoclassical agent, a behavioral agent may face psychic costs from renegeing on an explicit promise.⁶ Simply discussing beliefs may resolve a lot of strategic uncertainty and perhaps persuade D_i to hold favorable (symmetric) beliefs.

It is also conceivable that two-way communication could exacerbate the commitment problem. A deviating contract specifying a higher output and tariff than expected may be unappealing. U might be able to increase the appeal by adding an explanation that the deviation is the best response to the equilibrium offer, a special deal just for D_i . Two-way communication may destabilize the monopoly outcome.

While the effect of two-way communication on U 's commitment power is ambiguous, open

most of the rest of the paper, for brevity, statements that apply equally to wary and passive beliefs will just mention passive beliefs. Rey and Vergé (2004) show that wary and passive beliefs lead to different equilibrium outcomes if downstream firms engage in Bertrand competition.

⁶See Gneezy, Rockenbach and Serra-Garcia (2013), Serra-Garcia, van Damme and Potters (2013) and the references cited therein for recent studies on lying aversion.

communication among all three market participants seems likely to only enhance U 's commitment power. U can describe exactly the symmetric offers it will make and can urge the D_i to reject any other offers. The downstream firms observe everything U says, so they can verify that U has no opportunity to cut side deals with rivals or convince rivals to accept deviating offers.

2.4. Communication and Bargaining

Communication may also affect the way parties bargain. To focus exclusively on these effects, assume away the commitment problem for now by positing that contracts offer $x_i = Q^m/2$ units to each D_i , so that firms end up monopolizing the market. The only issue is the tariff offered (T_i) and whether the contract ends up being accepted (a_i).

In the benchmark model where no communication was possible, U issues take-it-or-leave-it offers to the D_i , allowing it to extract the full monopoly surplus $\Pi^m/2$ from each D_i . This was the equilibrium outcome we found with symmetric beliefs. Of course, this extreme theoretical predication may not materialize in practical markets or experiments because of fairness and other considerations. In practice, the familiar results from the ultimatum game may be observed with positive surplus afforded the responder (here represented by tariffs lower than $\Pi^m/2$) and contract rejections for less than equal divisions for the responder. In competitive settings such as ours, we may expect outcomes closer to the extreme theoretical prediction (Roth, Prasnikar, Okuno-Fujiwara and Zamir 1991).

Allowing two-way communication may have no effect on bargaining if parties simply babble during the communication stage, leaving U still with 100% of the bargaining power. On the other hand, two-way communication could function as a bargaining protocol with a more even distribution of power. Suppose, for example, that two-way communication is characterized by Nash bargaining between U and D_i over the contract $(Q^m/2, T_i)$, taking as given U offers $Q^m/2$ units to D_i and this is accepted. If U and D_i have equal bargaining power, they split the gains from agreement between them, $\Pi^m/2$ in half, effected by setting tariff $T_i = \Pi^m/4$. More generally, letting $\alpha \in [0, 1]$ be U 's bargaining power vis-a-vis D_i and $d \geq 1$ be the number of downstream firms, Nash bargaining leads to a tariff maximizing the Nash product

$$T_i^\alpha \left(\frac{\Pi^m}{d} - T_i \right)^{1-\alpha}, \quad (1)$$

as U gains revenue T_i and D_i gains its share of the monopoly profit Π^m/d less the payment T_i . Taking the first-order condition with respect to T_i and rearranging yields $T_i^* = \alpha\Pi^m/d$.

Moving to open communication among all three parties, again the outcome may be the same as with no communication because parties may just babble. If they engage in meaningful communication, however, the resulting three-way bargain can lead to a different division of surplus than two-way bargaining. To see this, again consider Nash bargaining. While there are several ways to extend the Nash product when multiple parties bargain, a natural one in our setting maintains a constant ratio between the bargaining weight for U and for an individual downstream firm for any number d of downstream firms, leading to the following Nash product:⁷

$$(dT_i)^{\frac{\alpha}{\alpha+d(1-\alpha)}} \prod_{i=1}^d \left(\frac{\Pi^m}{d} - T_i \right)^{\frac{1-\alpha}{\alpha+d(1-\alpha)}} = \left[(dT_i)^\alpha \left(\frac{\Pi^m}{d} - T_i \right)^{d(1-\alpha)} \right]^{\frac{1}{\alpha+d(1-\alpha)}}. \quad (2)$$

Maximizing this expression yields the equilibrium tariff under open communication:

$$T_i^{**} = \frac{\alpha\Pi^m}{d\alpha + d^2(1-\alpha)}. \quad (3)$$

Comparing equations (1) and (2), we have $T_i^* = [\alpha + (1-\alpha)d]T_i^{**} \geq T_i^{**}$, with strict inequality as long as $d \geq 2$ and $\alpha < 1$. In this case, our bargaining model predicts that the tariff should fall in moving from bilateral to open bargaining. While theories can easily be offered predicting the opposite (Shapley value, for instance), the variant of Nash bargaining offered here is at least one plausible model offering a comparative-static result for which we will later find support in the experiments.⁸

⁷Laurrelle and Valenciano (2008) provide a noncooperative foundation for the generalized Nash bargaining formulae in (1) and (2). In the limit as the probability of bargaining breakdown vanishes, the payoffs in a stationary subgame perfect equilibrium converge to those emerging from maximization of the Nash product, where the weights are given by the probability that the party is selected to be the proposer in a round. Translated into their terms, our specification would be equivalent to assuming that the ratio between the probability of selecting U for the proposer and of selecting a given D_i does not vary with d .

⁸The natural alternative to our specification of Nash bargaining maintains a constant ratio between U 's bargaining weight and the sum of all downstream firms' bargaining weights rather than a given downstream firm's. One can show that the equilibrium tariff is constant in d under this variant, which does not find support in the data.

3. Experimental Design

We build on the experimental design of Martin, Normann, and Snyder (2001). We will maintain their baseline treatment—which they 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 producer 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 $\pi_{D_i} = P(Q)q_i - a_i T_i$ for D_i . Let $\pi_D = \pi_{D_1} + \pi_{D_2}$ 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 $Q^m = 2$. The Cournot outcome involves market output $Q^c = 4$, firm output $q^c = 2$, and industry profit $\Pi^c = 72$.

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. 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 learned his profit; U was told his own and each of the two downstream firm’s profits that round. All these design features were explained to subjects in the instructions.

We conducted four different treatments. 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 remaining treatments introduced the possibility of communication using an instant-messaging technology via a chat window. In *Two Chat*, U could engage in private, two-way com-

munication 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 *Three Chat*, U , D_1 , and D_2 could freely communicate with each other. Whatever a player typed in his chat window was displayed to all three players in the market. It was not possible to exclude one of the players and engage in two-way chat. *Choose Chat* allowed each player to send each message via whichever communication channel—private communication between vertical levels as in *Two Chat* or the open communication as in *Three Chat*—he wanted. All channels were open in separate windows allowing receivers to know whether the message was sent privately or publicly.

Every round of *Two Chat*, *Three Chat*, and *Choose Chat* began with a communication stage prior to U 's making contract offers. Except for threats to be carried out outside the lab or information that could be used identify subjects, the content of the chat was unrestricted. In *Two Chat*, subjects had 90 seconds to chat during the first five rounds, reduced to one minute for the last ten rounds. The communication stage lasted 60 seconds in all 15 rounds in *Three Chat* and 90 seconds in all 15 rounds in *Choose Chat*. Subjects could not leave the chat stage before the time expired. The design of the communication treatments was otherwise identical to *No Chat*.

Subjects were invited using the ORSEE system (Greiner 2015). Upon arrival in the lab, each was assigned to a cubicle and provided with instructions, reproduced in Appendix B, available online. 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. The experiments were programmed in Z-tree (Fischbacher 2007).

It is possible for downstream firms to earn negative payoffs. To offset this possibility as well as 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 about 14 euros on average.

We conducted a total of 16 sessions, four sessions for each of the four treatments. All sessions were run at DICElab of the University of Duesseldorf from November 2013 to February 2015.

Each session lasted for about one hour. In total, 285 subjects participated.

4. Results

To streamline the discussion of our results, we will confine the initial discussion to the distinct treatments *No Chat*, *Two Chat*, and *Three Chat*. Once the relationship between *Two Chat* and *Three Chat* is understood, we can study which one the hybrid treatment *Choose Chat* is closer to.

The top part of Table 1 can be interpreted as a summary statistics for the main experimental variables. It regresses these variables (X , T_i , a_i , ...) on an exhaustive set of treatment indicators, suppressing the constant. This specification allows us to recover the treatment means of the variables as the coefficients on the indicators. The advantage of the regressions is that the supplied standard errors allow statistical tests of the differences between the means (bottom part). We compute White (1980) heteroskedasticity-robust standard errors clustered by session, allowing for dependence among observations arising from the same set of interacting subjects.

Comparing treatments *No Chat* and *SECRAN* from Martin, Normann, and Snyder (2001) provides a consistency check. Total offered quantity, $X = x_1 + x_2$ averaged 3.64 in *SECRAN*,⁹ nearly identical to the 3.68 in *No Chat* (see the first column of Table 1). The averages for total accepted quantity $Q = q_1 + q_2$ are also almost identical—2.41 in *SECRAN* versus 2.47 in *No Chat*—as are the averages for industry profit Π —68.2 in *SECRAN* versus 68.3 in *No Chat*. Upstream firms earned somewhat higher profit π_U in *SECRAN* (mean 51.2) compared to *No Chat* (mean 45.3). The remarkable consistency between *SECRAN* and *No Chat* suggests that *No Chat* is a good baseline for comparing treatments with communication.

4.1. Offered Quantity

We begin by analyzing total offered quantity, X . This single variable captures whether U is able to solve the commitment problem. Table 1 shows that the mean of X is highest in *No Chat*, 3.68, falling to 2.98 in *Two Chat*, falling further to 2.41 in *Three Chat*, close to the monopoly output

⁹The means for *SECRAN* reported here differ from those reported in Table 2 of Martin, Normann, and Snyder (2001). To reduce noise from inexperienced play, they dropped the first five rounds of each session. We are primarily interested in communication, which may have the largest effects in early rounds of play, so have chosen to focus on results for all rounds. To facilitate comparison, the last row of Table 1 here reports means from the Martin, Normann, and Snyder (2001) data retaining the first five rounds. These data are reported in histogram form in Figures 3–6.

of 2. These results are consistent with more open 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 circles shifts down from *No Chat* to *Two Chat* to *Three Chat*, providing confidence that decline toward the monopoly level is not due to an outlying session.

The bottom part of Table 1 provides formal statistical tests of the differences between treatment means. It reports differences between all combinations of treatment-indicator pairs, providing the appropriate standard errors for these differences. The fall in the mean of X from *No Chat* to *Two Chat* of 0.70, statistically significant at the 1% level, represents 40% of the gap between *No Chat* and the monopoly output. The fall from *Two Chat* to *Three Chat* of 0.57, statistically significant again at the 1% level, brings offered quantity 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 light grey bars, representing *Two 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 *Three Chat* piles almost all the mass in the monopoly ($X = 2$) bin.

Table 2 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 a 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 reported 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. *Three Chat* is 32 percentage points more likely to generate monopoly offers than *Two Chat*, a difference significant at the 1% level. *Two Chat* is 18 percentage points more likely to generate monopoly offers than *No Chat*, although this difference does not achieve significance at the 10% level.¹⁰ The next column regresses an indicator for the event $X \geq 4$, that is, that the offers total to at least the Cournot output. *Three Chat* is 17 percentage points less likely than *Two Chat* to have offers this high, and *Two*

¹⁰As we will see in Table 4, the difference is significant at the 10% level after dropping the first five rounds, reflecting noisier play, from the sample.

Chat is 23 percentage points less likely than *No Chat* to have offers this high, both differences significant at the 5% level. We conclude that increasing the openness of communication from *No Chat* to *Two Chat* to *Three Chat* results in a substantial and generally statistically significant shift in the mass from the Cournot to the monopoly bin.

Column 5 of Table 2 measures symmetry implicit in offered quantities. It estimates a linear model of the probability that the two contract offers involve symmetric quantities, $x_1 = x_2$. The estimate on the *No Chat* indicator implies that 68% of the offers in that treatment involve symmetric quantities. A similar percentage (in fact slightly lower) involve symmetric quantities in *Two Chat*. Evidently the private communication channel helps with monopolization but not symmetry. The open communication associated with *Three Chat* promotes symmetry: 88% of the offers involve symmetric quantities, over 20 percentage points more than *No Chat* or *Two Chat*, differences statistically significant at the 1% level. As the next column shows, the results for symmetry are similar if we take a stricter definition of symmetry, requiring all contractual terms (x_i and T_i) to be the same.

4.2. 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 can help measure how communication affects the division of surplus in the experiment. The mean reported in Table 1 falls from 34.7 ECU in *No Chat* to 31.4 in *Two Chat* to 26.9 in *Three Chat*. The means in *No Chat* and *Two Chat* are not significantly different from each other, but the mean in *Three Chat* is significantly lower than the others at the 5% level. The scatter plots of session-level means in Figure 3B tell a similar story: the group of filled dots shows a general downward shift from *No Chat* to *Two Chat* to *Three Chat*, although there is considerable spread in the dots for *Two Chat*.

Definitive inferences are difficult to draw from the raw means of T_i , however, because x_i varies systematically across treatments as well. To understand why this fact can pollute inferences, 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.¹¹ To purge these quantity effects, the third column of Table 1 restricts

¹¹With passive beliefs, the computation is less clear because the generosity of a contract depends on whether it is

the sample to contracts with $x_i = 1$. Now we see a decrease in the mean of T_i of 6.4 from *No Chat* to *Two Chat*, significant at the 1% level, and a further decrease of 3.2 from *Two Chat* to *Three Chat*, significant at the 5% level. These results suggest that starting from a situation in which U makes contract offers to the D_i , layering increasingly open communication allows the D_i to extract a greater share.

The fall in T_i from *No Chat* to *Three Chat* holding x_i constant is an intriguing result. The drop in tariff from *No Chat* to *Two Chat* is consistent with previous experimental work: introducing pre-play communication in the ultimatum game leads to more generous splits for the responder (Zultan 2012, using video chat). The further fall in T_i from *Two Chat* to *Three Chat* is to our knowledge a new experimental result. It is consistent with the Nash-bargaining model offered in Section 2.4. With two downstream firms ($d = 2$) and an upstream firm with less than complete bargaining power ($\alpha < 1$), the tariff under bilateral negotiations is strictly higher than that under open negotiations, that is, $T_i^* > T_i^{**}$. In effect, the additional downstream firm participating in open negotiations forces U to split surplus among three rather than just two parties.

Table 3 provides more evidence on the fit of various bargaining models to tariffs. For comparison, the last row shows the mean tariffs in the actual data in the *Two Chat* and *Three Chat* treatments, restricting the sample to offers with $x_i = 1$. The first row shows predicted tariff values, \hat{T}_i , from a Nash bargaining model positing a bargaining-power term for U of $\alpha = 0.5$, consistent with equal surplus division. Predicted tariffs match the comparative-static property of actual tariffs, falling from *Two Chat* to *Three Chat*, although predicted tariffs considerably underestimate actual ones in the *Three Chat* treatment. The next row continues with the Nash-bargaining model but now allows α to be a free parameter. We estimate α using non-linear least squares, in effect searching for the value providing the best fit between predicted and actual tariffs. The estimate is $\hat{\alpha} = 0.60$ with a standard error (clustered across sessions) of 0.02. Using the estimated $\hat{\alpha}$ in place of the posited $\alpha = 0.5$ results in a slightly worse fit between predicted and actual tariffs for *Two Chat* but a much improved fit for *Three Chat*.

To provide a counterpoint to Nash bargaining, the next two rows show fitted values for an alternative bargaining model, Shapley value. The row with $\alpha = 0.5$ is the standard version of Shapley in which all permutations of players used to compute marginal contributions are equally an equilibrium or out-of-equilibrium offer.

likely. The model gets the wrong comparative-static result, predicting a rise in tariffs with more open communication. The next row analyzes a generalized version of Shapley value, introducing a bargaining-power-like parameter that can be estimated to give it a better chance to fit the tariff data. Intuitively, in this row α is the probability that U comes after a given D_i in a permutation, giving U a better chance to make a positive marginal contribution. More detail on this generalization of Shapley value, as well as a discussion of its foundation in a model of asymmetric distributions of bargaining arrival times due to Kalai and Samet (1987), is provided in Appendix A. Non-linear least squares produces an estimate of $\hat{\alpha}$ of 0.39. In effect, the estimated version of the Shapley-value model tries to moderate the grossly overestimated tariffs in *Three Chat* by reducing U 's bargaining power. While this helps the fit in *Three Chat*, predicted tariffs now substantially undershoot actual in *Two Chat*. Thus the incorrect comparative-static result that tariffs rise with more open communication persists.

Overall, Table 3 shows that the version of Nash bargaining with the estimated $\hat{\alpha}$, besides getting the qualitative result right that tariffs fall from *Two Chat* to *Three Chat*, provides a reasonably good quantitative fit for tariffs in each treatment. The Bayesian Information Criterion (BIC) in the last column provides one gauge of fit across these non-nested models. An increase in BIC of 10 is typically taken as “very strong” evidence against the model with the higher BIC (Kass and Raftery 1995). Here we see that BIC is hundreds or thousands of points higher in alternatives to Nash bargaining with the estimated $\hat{\alpha}$.

4.3. Acceptance Behavior

Having analyzed upstream behavior, we next turn to downstream behavior, embodied in the acceptance decision a_i in Table 1. The acceptance rate rises from 70% in *No Chat* to 85% in *Two Chat* to 89% in *Three Chat*. Table 1 shows that the 15% increase from *No Chat* to *Two Chat* is significant at the 1% level but the further increase from *Two Chat* to *Three Chat* is insignificant. Figure 3C shows this same pattern holds in the session-level means.

The raw means of a_i provide a reduced-form measure of how acceptance rates vary with communication when the contract offers underlying the acceptance decision are also allowed to vary. The fifth column of Table 1 sheds light on how acceptance rates vary with communication holding contract offers constant. This column regresses a_i on the treatment indicators controlling for the

contract's terms in a semi-parametric way by restricting the sample to observations with $x_i = 1$ and including a second-order polynomial in standardized values \tilde{T}_i of the tariff (standardized by subtracting the sample mean and dividing by the variance). With this sample restrictions and added controls, the coefficients on the treatment indicators can be interpreted as the acceptance rates of a contract offering one unit at the sample mean tariff.

Controlling for contract offer reduces the gap between the *No Chat* and *Two Chat* acceptance rates as well as the *Two Chat* and *Three Chat* acceptance rates. We conclude, therefore, that the main reason acceptance rates rise from *No Chat* to *Two Chat* to *Three 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.

4.4. Market Output

The rest of the variables for which we provide summary statistics in Table 1 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 mean for market output Q in *No Chat*, 2.47, is about the same as in *Two Chat*, 2.49. The constancy of the mean between these treatments masks a significant change to the distribution of Q , shown in Figure 4B. Moving from *No Chat* to *Two Chat* concentrates the distribution from above and below on the mode at the monopoly outcome. The concentration from above is inherited from the effect that communication helps monopolize the market resulting in more offers of $X = 2$. 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 2, the monopoly outcome ($Q = 2$) is 17 percentage points more likely in *Two Chat* than *No Chat*, significant at the 10% level.

It should be emphasized that firms and consumers are not indifferent between treatments with the same mean for Q . Due to the concavity of industry profit in Q , a treatment which averages together values of Q well above the monopoly level with zero values from contract rejections will be much less profitable than a treatment in which Q varies less around its mean of 2.49. The opposite is true for consumer surplus, which is convex in Q . These facts will come into play in the

analyses of profits and consumer surplus in following subsections.

Moving from *Two Chat* to *Three Chat* reduces the mean of Q by 0.44 according to Table 1, significant at the 1% level. The mean of Q is 2.05 in *Three Chat*, very close to the monopoly output. Examining the full distribution of Q , it turns out the monopoly outcome ($Q = 2$) is 26 percentage points more likely in *Three Chat* than *Two Chat* according to Table 2, and Cournot or higher outputs ($Q \geq 4$) 14 percentage points less likely, both differences significant at the 5% level or better.

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

4.5. Profits

An analysis of profits will let us put a monetary value on the differences across treatments uncovered so far. First consider industry profit, Π . Table 1 shows that the mean rises from 68.3 to 82.5 to 89.5 ECU. The table shows that the 14.2 increase in the mean of Π from *No Chat* to *Two Chat* and 7.0 increase from *Two Chat* to *Three Chat* are statistically significant at the 5% level or better. 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 *Three Chat*, 89.5, is close to the monopoly profit of 100.

Moving to the allocation of profit across industry levels, U 's profits change non-monotonically across the treatments, increasing from 45.3 in *No Chat* to 51.1 in *Two Chat* and then falling to 42.5 in *Three Chat*. The substantial increase in the acceptance rate offsets a small decrease in tariff to cause the 5.8 increase in π_U from *No Chat* to *Two Chat*, significant at the 5% level. The adverse bargaining effects for U in moving from *Two Chat* to *Three Chat* ends up reducing π_U by 8.6, significant at the 5% level. The first rise and then fall leads to a fairly similar value of π_U between *No Chat* and *Three Chat*.

Although U 's profit level changes non-monotonically, its profit share, $s_U = \pi_U/\Pi$, shows a monotonic pattern in Table 1, falling from 0.63 in *No Chat* to 0.59 in *Two Chat* to 0.47 in *Three Chat*. More and more open communication leads U to obtain a smaller share of a growing pie. The

biggest drop in s_U (and only significant one), however, occurs in the move from *Two Chat* to *Three Chat*. As discussed in Section 2.4, the move from *Two Chat* to *Three Chat* could represent a fundamental change in the bargaining game, which under Nash bargaining could erode U 's bargaining surplus. The fall in s_U in *Three Chat* supports this theory.

Downstream firms gain both in absolute and relative terms from more and more open communication. Table 1 shows that the sum of downstream profits, π_D , rises from 23.0 ECU in *No Chat* to 31.4 in *Two Chat* to 47.1 in *Three Chat*, both increases significant at the 5% level or better, as shown in Table 1. Downstream profit is so high in *Three Chat* that they obtain a majority of the profit (53% compared to U 's 47%).

4.6. Consumer Surplus

The last column of Table 1 presents results for consumer surplus, CS . The mean of CS falls from 39.7 ECU in *No Chat* to 34.1 in *Two Chat* to 17.1 in *Three Chat*. Figure 3F disaggregates the means by session. The 5.6 fall from *No Chat* to *Two Chat* is not statistically significant, but the 17.1 fall from *Two Chat* to *Three Chat* is, at the 1% level. This large decline in CS between these treatments is due in part to the large reduction in the mean of Q , from 2.49 to 2.05, as consumers prefer higher quantities. Another factor relates to the convexity of CS in Q , which implies that consumers prefer more rather than less variance in Q . The reduction in the spread of Q from *Two Chat* to *Three Chat* shown in Figure 4 leads to a further reduction in CS between those treatments. This factor leads to the fall in CS moving from *No Chat* to *Two Chat* despite the increase in mean Q between the treatments. Hence we see that more and more open communication can lead to substantial consumer harm.

The monotonic increase in profit and decrease in consumer surplus offset each other, leading to fairly small changes in mean welfare across treatments. While U 's ability to monopolize the market is improved, reducing welfare, the decline in rejections (and decline in variance of Q , which is socially beneficial because, like profit, welfare is concave in Q) keeps welfare from falling very far in *Three Chat*. Whether these fairly benign welfare results carry over to markets outside the lab depends on how relatively important in real markets are the offsetting factors found in the lab. The possibility that enhanced monopolization may be the dominant factor in real markets, coupled with the unambiguous and large harm to consumers found in our experiments, leave ample cause

for policy concern.

4.7. Choose Chat Treatment

We now pick up the analysis of the *Choose Chat* treatment. The results show a clear pattern. For every variable in Table 1, the *Choose Chat* mean is between the means of the treatments of which *Choose Chat* is a hybrid, that is, the *Two Chat* and *Three Chat* treatments. For example, the 2.55 mean of X in *Choose Chat* is between the 2.98 for *Two Chat* and 2.41 for *Three Chat*. Comparing the *Choose Chat* – *Two Chat* difference to the *Three Chat* – *Choose Chat* difference at the bottom of Table 1, in every column the magnitude of the *Choose Chat* – *Two Chat* difference is weakly larger, meaning that the results for *Choose Chat* are closer to *Three Chat* than *Two Chat*.

Evidently, allowing players the option to communicate both privately and openly affords almost as much commitment power as restricting them to communicate openly. The results suggest that open communication can lead to monopolization even if, as is realistic, the upstream and downstream firms are also free to communicate privately. While antitrust outlaws private communication between downstream firms, it would be difficult to outlaw communication between levels as might be required to work out supply arrangements. Here we see that this permissible form of communication can help firms monopolize the market even in the absence of direct communication between downstream firms, which would be illegal, and that monopolization need not be subverted by legal, private communication between upstream and downstream firms.

4.8. Trends Within Session

The analysis so far has considered average effects over all rounds of play. In this subsection we explore whether the results show convergence or divergence trends as players gain experience in the market from early to late rounds. To uncover these trends, Table 4 repeats the regressions from Table 1 interacting the treatment indicators with indicators for the initial and end periods. 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 of the change.

The results show a fairly consistent trend. *No Chat* shows few significant changes over time. By

contrast, almost all the variables for the treatments with communication have significant changes, many at the 1% level. 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 functionality of communication. 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. This monopolization leads to a significant rise in industry profit Π , and a significant fall in CS . U is more generous with the D_i over time, leading to significant reductions in T_i , significant reductions in π_U , significant increases in π_D , and significant reductions in s_U .

The main change in *No Chat* is a 7 percentage point increase in the acceptance rate, leading to a 0.20 increase in Q , both trends statistically significant at the 1 % level. Thus, as players gain experience in *No Chat*, output diverges further from the monopoly output. The opposite happens in the communication treatments, as lower offered quantity translates into lower output. The mean of Q falls from early to late period across all of them, by as much as 0.35 units (in *Two Chat*, significant at the 1% level in that case). The combined effect of the increase in Q in *No Chat* and its decrease in the communication treatments results in the mean of Q being significantly higher in *No Chat* than in any of the communication treatments—even *Two Chat*—in the late period. This result leads us to conclude with even more confidence that communication leads to monopolization, whether measured by offered or realized quantity.

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. 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.

5. Analysis of Chat Content

In this section we draw further insights about the effect of different forms of communication by analyzing the content of the chat itself. We take a series of approaches: inspecting some simple descriptive statistics, gauging how much information an outside party can glean from the chat

(see, for example, Houser and Xiao 2011), studying representative quotes (Kimbrough, Smith and Wilson 2008), and finally running regressions of experimental outcomes on chat covariates.

5.1. Descriptive Statistics

Table 5 provides descriptive statistics characterizing chat across communication treatments. Each treatment may span several columns, both separating and combining channels operating simultaneously in a market to provide a full picture of the nature of chat.

The first set of variables are indicators for a message being sent in a round of chat: $Any Mes_U$ is an indicator for a message being sent by U , $Any Mes_D$ by one or both D_i , and $Any Mes$ by any player. Virtually all chat rounds (98% or higher) had at least some chat across all treatments. A conspicuous finding in *Choose Chat*, looking at the $Any Mes$ variable, is that subjects relied on the open more often than the private channel.

The next set of variables, $Num Mes$, record the number of messages sent by one level or the other or in total. In *Two Chat*, U sent 2.5 messages and each D_i sent 3.0 messages on average each round. The averages are almost identical in *Three Chat* (the downstream firms together sent 6.0 messages, implying 3.0 per individual D_i). In *Choose Chat*, players sent about this same number of messages via the open channel, but because they could also use the private channel, players ended up sending more messages in this than the other communication treatments.

The $Init$ variables indicate which level (U or D) initiated the chat, if any. In *Two Chat*, each bilateral chat was about equally likely to have been initiated by either side. In *Three Chat*, the probability of initiating chat, 29% for the upstream and 71% for the downstream firms, is close to what one would expect if each of the three players had an equal chance of being the first mover. The same is true for *Choose Chat* regarding the open channel, although the private channel was more likely to be initiated by a downstream firm.

Finally, the last set of rows presents correlations between the existence or extent of chat from the two sides. A positive correlation would be consistent with more chat from one side stimulating chat from the other, a negative correlation with chat from one side crowding out the other. Across all treatments the correlation is positive, suggesting that messages typically induce replies.

5.2. Coder Agreement

We next study whether the content of communication in *Two Chat* conveyed meaningful information about the terms of the contracts that would be offered that round. Following Houser and Xiao (2011), we asked two coders to independently analyze the chat content of *Two Chat*, *Choose Chat* and *Three 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, x_2, T_1, T_2) that would most likely result from the chat. If they thought that no plausible guess could be made, they were asked to enter “n.a.” instead of a number. They had read the instructions for the experiment up front and were aware of the communication structure in the treatments. At no point in time could the coders see the offers actually made. The coding was incentivized: five chats were randomly selected and the coders paid for the number of guesses that agreed with each other. For all treatments with communication, the same coders analyzed one complete session and five random rounds from the remaining three sessions. The sequence of markets and rounds were randomized such that the coders could not follow patterns involving certain subjects over time.

Figure 5 presents the results. Panels A and B show that communication was remarkably informative in *Two Chat*. Over 80% of the coders’ guesses for x_i matched the actual offer; over 95% of these also agreed with the other coder’s guess. Nearly two thirds of coder’s guesses for T_i matched the actual offer, and nearly 95% of these again agreed with the other coder’s guess. What makes the accuracy of T_i coding particularly noteworthy is that this variable could take on any of the large number of integers between 0 and 120. In the minority of the cases in which a coder’s did not match actual, their guesses still agreed with each other more often than not, suggesting that the chat was informative but misleading. This sort of misleading chat was fairly rare, for example accounting for fewer than 12% of coder’s guesses for x_i . Panels C–F show similar results for *Choose Chat* and *Three Chat*.

The accuracy of the chat coding leads us to strongly reject the null hypothesis that chat is meaningless babble in either the private or the open channel. We will thus inquire further into how subjects use this meaningful chat to facilitate contracting in the various communication treatments.

5.3. Representative Quotes

We hypothesized that better communication may allow subjects to better coordinate and avoid rejections. Specifically, heterogeneous beliefs about rival contract offers could generate considerable strategic uncertainty, some of which could be resolved by having parties communicate about 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 D_i has an opportunity to express a preference among several potential contracts.

The cleanest belief expressions are provided by the private communication channel in *Two Chat*. The following excerpt from a chat round is consistent with downstream firms’ having passive beliefs. (Recall upstream firms are called “producers” and downstream “retailers” in the experiment.)

```
Retailer: 2 for 20?  
Producer: can be done  
Retailer: that is good  
          barely :)  
Producer: 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 producer 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 having passive beliefs.

The following excerpt is consistent with the retailer’s having symmetric beliefs.

```
Producer: 2 for 25?  
Retailer: 1 for 20  
Producer: 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 producer, 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 producer responds (honestly or not) "the same." The producer 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 producer was honestly trying to establish the monopoly solution with both retailers.

In some cases, the producer 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 producer attempts to convince each retailer that he or she will become the monopolist at the expense of the other retailer (recall the producer in *Two Chat* had a separate window for each retailer).¹²

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

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

Taken together, these excerpts show that one can find examples of every possible use of communication in *Two Chat*, in some cases to enhance commitment leading to monopolization, in some cases undermining it with deviating "special offers". However, the results for the mean of market outcomes show that the net effect is that private communication moderately enhances monopolization. Examples like the second evidently outweigh those like the third.

The quantitative results from Section 4 suggest that open communication in *Three Chat* somehow bolsters a commitment to symmetric offers that monopolize the market. One mechanism for enhanced commitment depends more on what is left unsaid than what is said: the producer's opportunity to negotiate side deals is cut off in *Three Chat*. As the following excerpt illustrates, the active exchange of public messages seems to have a role in helping subjects arrive at monopolizing offers, perhaps an independent role, perhaps complementing the curtailment of private messages:

¹²Though the producer'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 producer'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." Note further that acceptance of the offers in the example of this passage is not supported by symmetric, passive, or wary beliefs. For discussion of possible beliefs other than these, see Eguia, Llorente-Saguer, Morton and Nicolò (2014).

Producer: Hi. How many do you both want?
 Retailer 1: Yo, producer, why don't you offer 1 each for 25, then you
 make a nice profit and so do we ;-)
 Retailer 2: An idea: to get a profit as large as possible, I suggest we
 sell only 2 units.
 Retailer 1: This is what i thought
 Producer: I agree
 Retailer 1: perfect

Note the use of plurals “both,” “we,” and “our,” suggesting symmetry of offers, as well as references to joint maximization.¹³

The hybrid *Choose Chat* treatment uses the mechanisms available in *Two Chat* and *Three Chat*. The analysis of these private and open channels gives a pure insight about how parties used the channels. Excerpts from *Choose Chat* yield additional insights where subjects switch between these channels.

Public chat

Retailer 2: hey, one for 23 for both
 Retailer 1: 1 for 25 for both.
 Producer: each 1 for 35.
 Retailer 2: 35 is too much ;)
 Producer: each 1 for 30?
 Retailer 2: way too much

Private chat

Retailer 1: I'd also buy 2 for 60 if the other one does not go all the way to 30
 Producer: OK
 Retailer 1: But then you have to promise to me
 Producer: 2 for 60 then?
 Retailer 1: okay

Here, a retailer tries to convince the producer to sidestep the public negotiations and do business exclusively with him. There is also evidence in the chats that some participants want to prevent precisely this sort of sidestepping with calls like “producer, type something here” or “always cooperate” into the public chat window.

5.4. Regressions on Chat Covariates

In the final piece of chat-content analysis, we regress various experimental outcomes on variables characterizing the chat from Table 5 among others. The variables are endogenous so their coeffi-

¹³Another mechanism for public communication to bolster commitment could involve various social-psychological effects. For example, Fay, Garrod and Carletta (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.

cients will not have causal interpretations, but will still reveal interesting correlations.

A conspicuous and statistically significant finding is that $Num\ Mes_D$ is associated with lower offered quantities, X , in all treatments and also with lower x_i in *Two Chat*. Evidently, more downstream chat helps arrive at quantities closer to the monopoly or at least is correlated with those outcomes. Whether the upstream firm initiates chat and how many messages it sends are not measurably associated with offered quantities. Another significant association that is somewhat robust is that $Num\ Mes_U$ is positively associated with s_U . More chat seems to help U extract a greater profit share.

Perhaps the most interesting findings are in the columns for *Choose Chat* including the *Any Private* variable, an indicator for whether any player used the private channel in the chat round. Resorting to the private channel is associated with a huge increase in X by 0.59 units, significant at the 1% level. Resorting to the private channel is also associated with a huge increase in s_U , by 16 percentage points, also significant at the 1% level. These results are suggestive of the possibility that U resorts to the private channel to cut side deals that allows it to extract a greater share at the same time it expands output by undermining commitment. As Table 5 shows, U resorts to the open channel twice as often as to the private channel, so the option to use the private channel in *Choose Chat* does not destroy commitment power completely.

6. Conclusion

In this paper, we introduce communication to a strategically complex vertical market. One upstream and two downstream firms can jointly earn monopoly rents but they may well fail to do so due to a commitment problem (Hart and Tirole 1990, Rey and Tirole 2007). The relevance of this commitment problem in turn depends on technical modeling assumptions: the (possibly heterogeneous) beliefs players maintain may suggest different equilibria in which the market may or may not be monopolized. In addition to players holding different expectations, bargaining frictions may add to the intricacy of the setup. Communication has the potential to overcome these problems. Our experimental treatments vary the openness or transparency of communication among the three players. The first treatment allows the upstream firm to engage in private two-way chat with each downstream firm. A second one lets all three firms engage in completely open (three-way) chat.

The third is a hybrid of the other two, allowing players the option of using either or both of two- or three-way communication.

Our first result is that increasing the openness of communication has a monotonic effect on market performance. Industry profits realize a minimum in the treatment without communication, increase for private two-way chat and the hybrid treatment, and attain a maximum for the open (three-way) chat. We thus find support for the hypothesis that communication can solve the commitment problem and results in higher profits. How firms communicate is important, though, and only when all three players can talk openly we observe full monopolization of the markets.

A second finding is a bargaining effect. More open communication leads to an increasing rate of contract acceptance. The increase in acceptance rate is partly due to a reduction in the upstream firm's tariff demands. Overall, increasing the openness of communication monotonically reduces the share of industry profits the upstream firm accrues. The additional profits from being able to better monopolize the market almost entirely go to downstream firms. A simple structural model of Nash bargaining fits the pattern of shifting surpluses well.

What are the positive and normative implications of our experimental results for real-world markets? It is reasonable to assume that open communication is not a practical option because firms cannot commit not to engage in private communication on the side. This leaves no communication, two-way chat and the hybrid form as practical communication structures. Both upstream and downstream profits are higher with two-way chat and the hybrid variant, thus firms prefer some form of communication to the treatment without, suggesting that some form of communication would endogenously emerge in the market. Given that upstream and downstream firms differ in their preferences over two-way chat versus the hybrid form of communication, it may be difficult to predict which would emerge without making additional assumptions. For instance, if private and public communication channels exist in the market, it may be difficult for parties to commit not to use them, in which case the hybrid variant would be a natural communication structure. Given that there are plausible conditions under which this form of communication may endogenously emerge, the monopolizing effects of communication and the steep decline in consumer surplus in this variant may be cause for antitrust concern.

Appendix A: Generalizing Shapley Value

This appendix presents a generalization of Shapley value allowing for asymmetric weights. We follow Kalai and Samet's (1987) foundation of this version of the Shapley value in a model of asymmetric arrival times.

To this end, assume that coalitions are formed from permutations arising from players randomly arriving at a location. Let A_U be U 's arrival time, exponentially distributed with rate parameter λ_U , and let A_i be the arrival time for a given D_i , exponentially distributed with rate parameter λ_D , symmetric across downstream firms. Assume arrival times are independent. Define $\alpha = \Pr(A_U > A_i)$. Using standard results for exponential distributions, one can show

$$\alpha = \frac{\lambda_D}{\lambda_D + \lambda_U}. \quad (\text{A1})$$

U 's marginal contribution to its coalition is 0 if it comes first in the permutation and Π^m otherwise. Thus U 's generalized Shapley value from a bargain in which U and d downstream firms participate is

$$\Pi^m \Pr\left(A_U > \min_{i \in \{1, \dots, d\}} \{A_i\}\right) = \Pi^m \left[1 - \Pr\left(A_U < \min_{i \in \{1, \dots, d\}} \{A_i\}\right)\right] \quad (\text{A2})$$

$$= \Pi^m \left(\frac{d\lambda_D}{\lambda_U + d\lambda_D}\right) \quad (\text{A3})$$

$$= \Pi^m \left(\frac{\alpha d}{1 - \alpha + \alpha d}\right), \quad (\text{A4})$$

where (A3) follows from standard results for exponential distributions and (A4) from (A1).

The tariff implementing the equilibrium surplus share in (A4) is

$$T_i^* = \frac{\Pi^m}{2} \left(\frac{\alpha d}{1 - \alpha + \alpha d}\right). \quad (\text{A5})$$

This equation provides the fitted tariff values for the rows in Table 3 for the Shapley value.

These formulas nest the standard Shapley value with symmetric weights, which can be recovered by substituting $\alpha = 1/2$. Take the case of $d = 1$, corresponding to the bilateral bargaining of *Two Chat*. U 's share of the monopoly profit Π^m then is $1/2$ and the equilibrium tariff is $\Pi^m/4$. Take the case of $d = 2$, corresponding to the open communication of *Three Chat*. U 's share of the monopoly profit rises to $2/3$ and the equilibrium tariff to $\Pi^m/3$.

The fact that U 's share and tariffs rise with d generalizes beyond the symmetric case of $\alpha = 1/2$. For any $\alpha \in (0, 1)$, one can show that equations (A4) and (A5) are increasing in d . This provides a contrasting comparative-static result to that derived in the text for the generalized Nash product.

References

- Andersson, Ola and Erik Wengström. (2007) “Do Antitrust Laws Facilitate Collusion? Experimental Evidence on Costly Communication in Duopolies,” *Scandinavian Journal of Economics*, 109(2): 321–339.
- Andersson, Ola, Matteo M. Galizzi, Tim Hoppe, Sebastian Kranz, Karen van der Wiel, and Erik Wengström. (2010) “Persuasion in experimental ultimatum games,” *Economics Letters*, 108(1): 16–18.
- Arango, Tim and Brian Stelter. (2011) “Comcast Receives Approval for NBC Universal Merger,” *New York Times* January 19, p. B9.
- Avenel, Eric. (2012) “Upstream Capacity Constraint and the Preservation of Monopoly Power in Private Bilateral Contracting,” *Journal of Industrial Economics*, 60: 578–598.
- Badasyan, 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).
- Balliet, Daniel. (2010) “Communication and Cooperation in Social Dilemmas: A Meta-Analytic Review,” *Journal of Conflict Resolution*, 54(1): 39–57.
- 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, Jeannette, Axel Ockenfels and Joachim Weimann. (2003) “The Effect of Communication Media on Cooperation,” *German Economic Review*, 4: 217–241.
- Charness, Gary and Martin Dufwenberg. (2006) “Promises and Partnership,” *Econometrica*, 74: 1579–1601.
- Cooper, David J. and Kai-Uwe Kühn. (2009) “Communication, Renegotiation, and the Scope for Collusion,” *American Economic Journal: Microeconomics*, 6(2): 247–78.
- Crawford, Vincent and Joel Sobel. (1982) “Strategic Information Transmission,” *Econometrica*, 50: 1431–51.
- Crawford, Vincent. (1998) “A Survey of Experiments on Communication via Cheap Talk,” *Journal of Economic Theory*, 78: 286–298.
- Dawes, Robyn M., Jean McTavish, and Harriet Shaklee. (1977) “Behavior, Communication, and Assumptions About other People’s Behavior in a Commons Dilemma Situation,” *Journal of Personality and Social Psychology*, 35(1): 1–11.
- 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.
- 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.
- Fischbacher, Urs. (2007) “Z-Tree: Zurich Toolbox for Readymade Economic Experiments,” *Experimental Economics*, 10: 171–178.
- 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.
- Gneezy, Uri, Bettina Rockenbach, and Marta Serra-Garcia. (2013) “Measuring Lying Aversion,” *Journal of Economic Behavior and Organization*, 93 293–300.
- Greiner, Ben. (2015) “Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE,” *Journal of the Economic Science Association*, 1: 114–125.
- Hart, Oliver and Jean Tirole. (1990) “Vertical Integration and Market Foreclosure,” *Brookings Papers on Economic Activity (Microeconomics)*, 205–285.
- Houser, Daniel and Erte Xiao. (2011) “Classification of Natural Language Messages Using A Coordination Game,” *Experimental Economics*, 14: 1–14.
- 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.
- Kalai, Ehud and Dov Samet. (1987) “On Weighted Shapley Values,” *International Journal of Game Theory* 16: 205–222.
- Kanter, James and Mark Scott. (2015) “Challenge to Google,” *New York Times*, April 16, page B1.
- Kass, Robert E. and Adrian E. Raftery. (1995) “Bayes Factors,” *Journal of the American Statistical Association*, 90: 773–795.
- Kimbrough, Erik, O., Vernon Smith and Bart Wilson. (2008) “Historical Property Rights, Sociality, and the Emergence of Impersonal Exchange in Long-Distance Trade,” *American Economic Review*, 98(3): 1009–1039.
- Landeo, Claudia M. and Kathryn E. Spier. (2009) “Naked Exclusion: An Experimental Study of Contracts With Externalities,” *American Economic Review*, 99: 1850–1877.
- Laruelle, Annick and Federico Valenciano. (2008) “Noncooperative Foundations of Bargaining Power in Committees and the Shapley-Shubik Index,” *Games and Economic Behavior* 63: 341–353.

- 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.
- Rey, Patrick and Stéphane Caprice. (2015) “Buyer Power from Joint Listing Decisions,” forthcoming in: *Economic Journal*.
- Rey, Patrick and Jean Tirole. (2007) “A Primer on Foreclosure,” in Mark Armstrong and Robert Porter, eds., *Handbook of Industrial Organization*, vol. 3, North-Holland, 2145–2220.
- Rey, Patrick and Thibaud Vergé. (2004) “Bilateral Control with Vertical Contracts,” *Rand Journal of Economics*, 35: 728–746.
- Roth, Alvin E. (1995) “Bargaining Experiments,” in John Kagel and Alvin E. Roth, eds., *Handbook of Experimental Economics*, Princeton University Press, 253–348.
- Roth, Alvin E., Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir. (1991) “Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study,” *American Economic Review*, 81: 1068–1095.
- Serra-Garcia, Marta, Eric van Damme, and Jan Potters. (2013) “Lying About What You Know or About What You Do?,” *Journal of the European Economic Association*, 11(5): 1204–1229.
- 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: Regressions Examining Differences in Means

	Dependent variable											
	X	T_i	T_i	a_i	a_i	a_i	ϱ	Π	π_U	π_D	s_U	CS
<i>No Chat</i>	3.68*** (0.19)	34.7*** (1.2)	33.3*** (0.9)	0.70*** (0.02)	0.77*** (0.03)	2.47*** (0.19)	68.3*** (1.8)	45.3*** (0.5)	23.0*** (1.8)	0.63*** (0.01)	39.7*** (5.7)	
<i>Two Chat</i>	2.98*** (0.19)	31.4*** (2.3)	26.9*** (1.2)	0.85*** (0.04)	0.85*** (0.03)	2.49*** (0.12)	82.5*** (2.7)	51.1*** (2.6)	31.4*** (3.0)	0.59*** (0.03)	34.1*** (4.4)	
<i>Choose Chat</i>	2.55*** (0.15)	26.9*** (1.2)	24.8*** (1.1)	0.89*** (0.01)	0.89*** (0.02)	2.20*** (0.07)	87.4*** (2.3)	46.2*** (2.4)	41.2*** (3.1)	0.52*** (0.03)	22.6*** (3.2)	
<i>Three Chat</i>	2.41*** (0.06)	25.0*** (1.0)	23.7*** (0.9)	0.89*** (0.01)	0.86*** (0.02)	2.05*** (0.03)	89.5*** (1.0)	42.5*** (1.5)	47.1*** (2.1)	0.47*** (0.02)	17.1*** (0.6)	
Other controls	None	None	None	None	$\tilde{T}_i, \tilde{T}_i^2$	None	None	None	None	None	None	
Sample	All	All	$x_i = 1$	All	$x_i = 1$	All	All	All	All	$\Pi > 0$	All	
Observations	1,425	2,850	1,797	2,850	1,797	1,425	1,425	1,425	1,425	1,324	1,425	
R^2	0.12	0.09	0.14	0.05	0.17	0.02	0.09	0.02	0.11	0.09	0.06	
Coefficient differences												
<i>Two Chat – No Chat</i>	-0.70*** (0.27)	-3.2 (2.6)	-6.4*** (1.5)	0.15*** (0.04)	0.09* (0.05)	0.02 (0.23)	14.2*** (3.2)	5.8** (3.2)	8.9** (3.5)	-0.03 (0.03)	-5.6 (7.2)	
<i>Choose Chat – No Chat</i>	-1.13*** (0.25)	-7.7*** (1.7)	-8.5*** (1.4)	0.19*** (0.02)	0.12*** (0.04)	-0.27 (0.20)	19.1*** (2.9)	0.9 (2.5)	18.2*** (3.6)	-0.11*** (0.03)	-17.1** (6.5)	
<i>Three Chat – No Chat</i>	-1.27*** (0.20)	-9.7*** (1.5)	-9.6*** (1.2)	0.20*** (0.02)	0.09*** (0.04)	-0.42* (0.19)	21.2*** (2.0)	-2.9* (1.6)	24.1*** (2.7)	-0.15*** (0.02)	-22.6*** (5.7)	
<i>Choose Chat – Two Chat</i>	-0.43* (0.24)	-4.5* (2.5)	-2.1 (1.6)	0.04 (0.04)	0.03 (0.03)	-0.29* (0.14)	4.9 (3.5)	-4.9 (3.5)	9.7** (4.3)	-0.07* (0.04)	-11.5* (5.4)	
<i>Three Chat – Two Chat</i>	-0.57** (0.20)	-6.4** (2.5)	-3.2* (1.5)	0.04 (0.04)	0.00 (0.03)	-0.44*** (0.13)	7.0** (2.9)	-8.6** (3.0)	15.6*** (3.6)	-0.12** (0.03)	-17.0*** (4.4)	
<i>Three Chat – Choose Chat</i>	-0.14 (0.17)	-1.9 (1.5)	-1.1 (1.4)	-0.00 (0.02)	0.03 (0.02)	-0.15* (0.07)	2.1 (2.5)	-3.8 (2.8)	5.9 (3.7)	-0.05 (0.03)	-5.5 (3.3)	

Notes: Each column is an ordinary least squares regression. Exhaustive set of treatment indicators included as regressors and constant omitted. Sample includes all 15 rounds in each session. Sample for s_U column excludes observations with $\Pi = 0$ for which s_U undefined. A small subset (6%) of these involve $\pi_D < 0$; we set $s_U = 1$ for these. White (1980) heteroskedasticity-robust standard errors clustered at session level reported in parentheses. Regressions for T_i and a_i run for all contract offers and for with $x_i = 1$. The regression for a_i with other controls includes standardized tariff \tilde{T}_i and its square, giving coefficients on the treatment indicators the interpretation of mean acceptance rates for contracts offering mean tariff. Significantly different from 0 in a two-tailed test at the * 10% level, ** 5% level, *** 1% level.

Table 2: Linear Probability Models of Outcome Variables

	Measuring monopolization				Measuring symmetry	
	Offered quantity, X		Market output, Q		$x_1 = x_2$	$x_1 = x_2,$ $T_1 = T_2$
	$X = 2$	$X \geq 4$	$Q = 2$	$Q \geq 4$		
<i>No Chat</i>	0.30*** (0.06)	0.54*** (0.06)	0.30*** (0.04)	0.29*** (0.07)	0.68*** (0.01)	0.55*** (0.01)
<i>Two Chat</i>	0.48*** (0.10)	0.30*** (0.07)	0.47*** (0.07)	0.21*** (0.05)	0.63*** (0.07)	0.39*** (0.05)
<i>Choose Chat</i>	0.71*** (0.06)	0.17*** (0.04)	0.67*** (0.06)	0.10*** (0.03)	0.76*** (0.04)	0.65*** (0.05)
<i>Three Chat</i>	0.79*** (0.02)	0.13*** (0.02)	0.73*** (0.02)	0.07** (0.01)	0.88*** (0.02)	0.76*** (0.03)
Observations	1,425	1,425	1,425	1,425	1,425	1,425
R^2	0.16	0.13	0.12	0.06	0.05	0.08
Coefficient differences						
<i>Two Chat – No Chat</i>	0.18 (0.12)	-0.23** (0.09)	0.17* (0.08)	-0.08 (0.09)	-0.05 (0.07)	-0.16*** (0.05)
<i>Choose Chat – No Chat</i>	0.41*** (0.09)	-0.37*** (0.08)	0.37*** (0.07)	-0.19** (0.07)	0.08* (0.04)	0.10* (0.05)
<i>Three Chat – No Chat</i>	0.50*** (0.07)	-0.41*** (0.06)	0.43** (0.05)	-0.22*** (0.04)	0.21*** (0.02)	0.21*** (0.03)
<i>Choose Chat – Two Chat</i>	0.23* (0.12)	-0.14 (0.08)	0.21** (0.09)	-0.11* (0.06)	0.13 (0.08)	0.26*** (0.07)
<i>Three Chat – Two Chat</i>	0.32*** (0.10)	-0.17** (0.07)	0.26*** (0.07)	-0.14** (0.05)	0.25*** (0.07)	0.37*** (0.05)
<i>Three Chat – Choose Chat</i>	0.08 (0.07)	-0.03 (0.05)	0.05 (0.07)	-0.03 (0.03)	0.12** (0.04)	0.11* (0.05)

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. Exhaustive set of treatment indicators included as regressors and constant omitted. Sample includes all 15 rounds in 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 3: Tariffs Predicted by Various Bargaining Models

Model	Mean \hat{T}_i in subsample		BIC
	<i>Two Chat</i>	<i>Three Chat</i>	
Nash bargaining			
Posit $\alpha = 0.50$	25.0	16.7	70,351
NLLS estimate $\hat{\alpha} = 0.60$	30.2	21.6	6,998
Shapley value			
Posit $\alpha = 0.50$	25.0	33.3	108,258
NLLS estimate $\hat{\alpha} = 0.39$	19.7	28.2	7,298
Actual data	26.9	23.7	

Notes: Sample restricted to offers involving $x_i = 1$ in *Two Chat* and *Three Chat* treatments only. Each row is a different model, for which we display fitted tariff values \hat{T}_i for the two included treatments as well as the Bayesian Information Criterion (BIC) to compare model fits. For rows involving an estimate $\hat{\alpha}$, estimation performed using non-linear least squares, equivalent to maximum likelihood assuming $\varepsilon_i = T_i - \hat{T}_i$ has standard normal distribution.

Table 4: Trends in Treatment Effects

	Dependent variable											
	X	\tilde{T}_i	\tilde{T}_i	a_i	a_i	a_i	\tilde{Q}	Π	π_U	π_D	s_U	CS
<i>No Chat0</i>	3.80*** (0.31)	38.2*** (3.1)	35.0*** (2.2)	0.66*** (0.01)	0.75*** (0.04)	2.33*** (0.19)	64.4*** (4.1)	45.5*** (3.0)	19.0*** (2.3)	0.66*** (0.02)	36.5*** (4.9)	
<i>No Chat1</i>	3.62*** (0.15)	32.9*** (0.4)	32.4*** (1.2)	0.72*** (0.02)	0.77*** (0.04)	2.54*** (0.19)	70.3*** (1.3)	45.3*** (1.7)	25.0*** (2.7)	0.61*** (0.03)	41.3*** (6.4)	
<i>Two Chat0</i>	3.32*** (0.25)	33.5*** (2.2)	26.8*** (1.0)	0.84*** (0.04)	0.86*** (0.03)	2.72*** (0.16)	78.9*** (3.4)	53.5*** (2.6)	25.4*** (2.9)	0.61*** (0.03)	42.5*** (5.4)	
<i>Two Chat1</i>	2.80*** (0.17)	30.4*** (2.3)	26.9*** (1.3)	0.85*** (0.04)	0.85*** (0.04)	2.37*** (0.12)	84.3*** (2.4)	49.9*** (2.7)	34.5*** (2.1)	0.59*** (0.03)	29.9*** (4.2)	
<i>Choose Chat0</i>	2.90*** (0.23)	29.6*** (1.4)	27.7*** (1.2)	0.85*** (0.03)	0.90*** (0.01)	2.32*** (0.09)	82.1*** (3.6)	48.2*** (1.8)	33.9*** (4.0)	0.57*** (0.03)	28.9*** (4.0)	
<i>Choose Chat1</i>	2.37*** (0.13)	25.6*** (1.3)	23.7*** (1.2)	0.92*** (0.01)	0.88*** (0.02)	2.14*** (0.11)	90.1*** (1.9)	45.2*** (2.9)	44.8*** (3.3)	0.50*** (0.03)	19.5*** (3.9)	
<i>Three Chat0</i>	2.86*** (0.25)	30.0*** (1.5)	28.4*** (1.7)	0.83*** (0.02)	0.87*** (0.03)	2.20*** (0.10)	83.7*** (2.2)	45.7*** (3.0)	37.9*** (3.5)	0.54*** (0.03)	24.7*** (3.2)	
<i>Three Chat1</i>	2.18*** (0.03)	22.5*** (0.7)	21.9*** (0.4)	0.92*** (0.00)	0.85*** (0.02)	1.98*** (0.03)	92.4*** (0.5)	40.8*** (1.0)	51.6*** (1.4)	0.44*** (0.01)	13.4*** (1.0)	
Other controls	None	None	None	None	$\tilde{T}_i, \tilde{T}_i^2$	None	None	None	None	None	None	
Sample	All	All	$x_i = 1$	All	$x_i = 1$	All	All	All	All	$\Pi > 0$	All	
Observations	1,425	2,850	1,797	2,850	1,797	1,425	1,425	1,425	1,425	1,324	1,425	
R^2	0.15	0.12	0.20	0.05	0.17	0.03	0.10	0.03	0.14	0.12	0.08	
Coefficient differences												
<i>No Chat1 - No Chat0</i>	-0.18 (0.19)	-5.3* (2.9)	-2.6 (2.9)	0.07*** (0.02)	0.02 (0.06)	0.20*** (0.05)	5.8 (4.0)	-0.2 (4.6)	6.0 (3.8)	-0.05 (0.05)	4.8 (3.4)	
<i>Two Chat1 - Two Chat0</i>	-0.52*** (0.13)	-3.1*** (0.4)	0.2 (0.3)	0.01 (0.01)	-0.01 (0.01)	-0.35*** (0.11)	5.5*** (1.4)	-3.6** (1.6)	9.1*** (0.4)	-0.03 (0.02)	-12.6*** (3.3)	
<i>Choose Chat1 - Choose Chat0</i>	-0.53*** (0.15)	-3.9** (1.5)	-4.0*** (1.1)	0.07* (0.03)	-0.03 (0.02)	-0.18 (0.15)	8.0*** (2.6)	-2.9 (1.9)	10.9*** (3.6)	-0.07*** (0.02)	-9.4* (4.9)	
<i>Three Chat1 - Three Chat0</i>	-0.68*** (0.27)	-7.5*** (0.9)	-4.7** (2.1)	0.09*** (0.02)	-0.01 (0.02)	-0.22* (0.12)	8.8*** (1.8)	-4.9* (2.6)	13.7*** (2.4)	-0.10*** (0.02)	-11.3** (4.1)	

Notes: Each column is an ordinary least squares regression including interactions between a set of treatment indicators and a set of (initial, end) period indicators. Interactions denoted with subscripts; for example, *No Chat0* is the interaction between *No Chat* and the initial period consisting of rounds 1–5, and *No Chat1* is the interaction between *No Chat* and the end period consisting of rounds 6–15. Exhaustive set of interactions included as regressors and constant omitted. Sample includes all 15 rounds. Sample for s_U column excludes observations with $\Pi = 0$ for which s_U undefined. A small subset (6%) of these involve $\pi_D < 0$; we set $s_U = 1$ for these. Two regressions run for \tilde{T}_i and a_i , one for all contract offers and one for contract offers with $x_i = 1$. The regression for a_i with other controls includes standardized tariff \tilde{T}_i and its square, giving coefficients on the treatment indicators the interpretation of mean acceptance rates for contracts offering mean tariff. 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: Chat Content Descriptive Statistics

	Two Chat		Choose Chat			Three Chat
	Channels separated	Channels combined	Private channels separated	Private channels combined	Open channel	
Means of indicators for some message sent						
<i>Any Mes_U</i>	0.97	0.99	0.32	0.43	0.86	0.97
<i>Any Mes_D</i>	0.99	1.00	0.39	0.55	0.93	0.99
<i>Any Mes</i>	0.99	1.00	0.42	0.58	0.95	0.99
Means of number of messages sent						
<i>Num Mes_U</i>	2.5	5.0	0.8	1.5	3.0	4.6
<i>Num Mes_D</i>	3.0	6.0	1.0	2.0	6.6	8.6
<i>Num Mes</i>	5.5	11.0	1.8	3.5	9.6	13.1
Means of indicators for chat initiation						
<i>Init_U</i>	0.52	0.48	0.15	0.16	0.28	0.30
<i>Init_D</i>	0.50	0.57	0.27	0.42	0.70	0.72
Correlations across market levels						
<i>Any Mes_U with Any Mes_D</i>	0.29	<i>a</i>	0.72	0.66	0.46	0.47
<i>Num Mes_U with Num Mes_D</i>	0.43	0.43	0.72	0.76	0.61	0.38
Observations	690	345	690	345	345	360

Notes: Sum of *Init_U* and *Init_D* down column can exceed 1 because time was measured in discrete units (seconds), resulting in some ties for initiator. Sum of *Init_U* and *Init_D* down column can be less than 1 when that channel was not used, so there was no chat initiator, for some observations. "Correlation undefined because variance of *Any Mes_D* equals 0.

Table 6: Regressions Including Chat Covariates

	Two Chat		Choose Chat		Three Chat				
	x_i	s_U	X	s_U	X	s_U			
Constant	1.50*** (0.09)	3.04*** (0.23)	0.65*** (0.10)	2.81*** (0.20)	2.49*** (0.18)	0.50*** (0.04)	0.41*** (0.03)	2.92*** (0.18)	0.54*** (0.04)
$Num\ Mes_U$	0.03 (0.04)	0.04 (0.06)	0.02* (0.01)	0.04 (0.04)	0.02 (0.03)	0.02* (0.01)	0.01 (0.01)	-0.08** (0.02)	0.00 (0.01)
$Num\ Mes_D$	-0.03*** (0.01)	-0.04** (0.01)	-0.02 (0.01)	-0.05** (0.01)	-0.04** (0.01)	-0.00 (0.00)	-0.00 (0.00)	-0.06* (0.02)	-0.01 (0.01)
$Init_U$	-0.00 (0.07)								
<i>Any Private</i>					0.59*** (0.08)		0.16*** (0.03)		
Observations	690	345	330	345	345	327	327	360	346
R^2	0.00	0.01	0.02	0.03	0.10	0.02	0.06	0.08	0.05

Notes: Each column is an ordinary least squares regression. Sample includes all 15 rounds in 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.

Figure 1: Vertical Structure




Figure 2: Experimental Market Demand and Profit








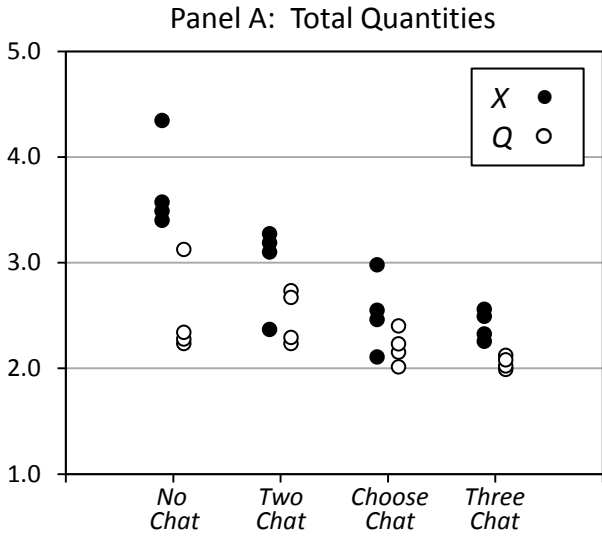


Figure 3: Session-Level Means



Notes: Small vertical jitter added to overlapping means where necessary to distinguish the four sessions for each treatment. In Panels B and C, black dots represent all observations and white dots just for contract offers involving one unit.

Figure 4: Quantity Histograms






Figure 5: Accuracy and Mutual Agreement of Coded Chat



Notes: *Two Chat* sample consists of 350 contract offers, *Choose Chat* of 370, and *Three Chat* of 360, each of which is assessed by two coders. Grey bar is proportion of sample for which coder's guess of contractual variable agrees with other coder and black for which his guess disagrees with the other coder. N.a. indicates an affirmative statement that coder could not guess variable based on chat content.