Competition in concentrated markets: Three Essays on Experimental Industrial Organization

DISSERTATION

zur Erlangung des akademischen Grades
doctor rerum politicarum (dr. rer. pol.)
im Fach Volkswirtschaftslehre

eingereicht an der Wirtschaftswissenschaftlichen Fakultät
Heinrich-Heine-Universität Düsseldorf

von: Claudia Möllers, M.Sc.
geboren am 19.02.1985 in Lingen (Ems)

Erstgutachter: Prof. Dr. Hans-Theo Normann
Zweitgutachter: Prof. Dr. Alexander Rasch

Abgabedatum: 31.10.2016

– to my parents –

Acknowledgment

This thesis was carried out during my time as a doctoral researcher at *Düsseldorf Institute for Competition Economics* (DICE) at the Heinrich-Heine University of Düsseldorf. I want to thank my colleagues at DICE for insightful discussions, professional and personal support. Additionally, I want to thank the Deutsche Forschungsgemeinschaft for the financial support.

I want to thank my co-authors Volker Benndorf, Christopher M. Snyder, Torben Stühmeier and Tobias Wenzel and my second supervisor Alexander Rasch. I was fortunate to work with great and experienced researchers from whom I learned a lot. Further, I want to thank Sander Onderstal for the opportunity to stay in the University of Amsterdam and for being a fantastic office mate. My work benefited also from discussions and established contacts from international conferences and workshops as well as visitors at DICE.

In particular, I am deeply grateful to Hans-Theo Normann. From the first game theory lecture I attended, he supported and encouraged me throughout the last six years. Integrating exiting research questions in his lecture, raised my interest in experimental economics and industrial organization in the first place. He gave me the opportunity to visit two great universities in Melbourne and Amsterdam and meet well-known, international researchers. Accompanying my time as a PhD student, he did not only help and support me as my co-author, but was always available as my supervisor. He gave me professional advice, discussed research questions with me and helped me to proceed. Thank you for your support.

Also, I want to thank my friends and family. I consider myself lucky to have so many inspiring, caring and reliable friends in my life, thank you. My family supported me throughout my life, gave me the opportunity to do what I like and were there for me anytime. Thank you, Franziska, Günter, Kerstin and Matthias. Finally, I want to thank Pedro for his patience, encouragement and love.

Contents

List of Tables

List of Figures

Chapter 1

General Introduction

The study of markets and functioning competition has become a central topic in microeconomics over the past 50 years. On a game-theoretic foundation, theorists started to analyze the relevance and impact of different parameters on competition: the market structure, e.g. vertical integration, number of sellers, product differentiation, the cost function and so on. Traditionally, theories rely on the assumption that agents are rational which might be appropriate in some cases but in others the assumption is less plausible. Chamberlin (1948) conducted the first market experiment to demonstrate that perfect competition is not obtained in real markets. In contrast, Smith (1962) implemented a double auction procedure in an experiment and obtained the competitive outcome.

There is an ongoing debate (see for example Smith, 1980, Falk and Fehr, 2003, Levitt and List, 2007a, 2007b, 2008, Falk and Heckman, 2009, Camerer, 2011) whether laboratory experiments are an appropriate method to get further insides into the functioning of "real" markets. While research in laboratories is indispensable in natural sciences, the method established in social sciences much later. The main critique is the relevance of obtained results or the realism of experiments in the laboratory.

Vernon Smith replies: "First, if the purpose of the experiment is to test a theory, are the elements of alleged unrealism in the experiment parameters of the theory? If not, then the criticism must be directed to the theory as much as to the experiment. Laboratory experiments are normally as "rich" as the theories they test. Second, are there field data to support the criticism, i.e., data suggesting that there may be differences between laboratory and field behavior. If not, then the criticism is pure speculation; if so, then it is important to parametrize the theory to include the behavior in question." (Smith, 1980, p. 350)

Plott states: "While laboratory processes are simple in comparison to naturally occurring processes, they are real processes in the sense that real people participate for real and substantial profits and follow real rules in doing so. It is precisely because they are real that they are interesting" (Plott, 1982, p. 1486).

In fact, a growing literature compares findings from the lab and the field supporting the relevance of obtained results from the laboratory (see Camerer, 2011 for an overview). Actually, the tight control of the parameters in the lab is a huge advantage in order to derive conclusions about causality. "The issue of realism, however, is not a distinctive feature of lab versus field data. The real issue is determining the best way to isolate the causal effect of interest." (Falk and Heckman, 2009, p. 536)

Consider the market for fuel as an illustrating example. The German cartel office concluded from a sector inquiry in 2011 that the market for fuel is dominated by five firms owning approximately 65% of gas stations in Germany. While there was evidence of coordinated behavior, the investigation did not find proof for explicit collusion.[1] Therefore, it is still an open question whether gas stations communicate their prices and if so, how the coordination works. The inquiry resulted in further initiatives on the part of the cartel office.

First, in order to increase price transparency for consumers, the cartel office installed in 2013 the "Markttransparenzstelle für Kraftstoffe" which provides real-time price information via application, navigation system or internet. One year after the implementation, average prices were unaffected by the intervention, also two years later there was no significant impact of the "Markttransparenzstelle für Kraftstoffe" on the price level. However, prices became more dispersed, that is, the variation of prices got larger. Consequently, the aim to decrease prices for consumers was not met.[2]

Second, a concern about anticompetitive behavior in the market for fuel remained and resulted in a second sector inquiry. Since 2012 the wholesale market for fuel is under investigation.[3] In particular, the pass-on of crude oil price movements to final consumers is a question of interest. However, even with perfect collusion in the wholesale market, it is a rather open question how it translates into monopolization downstream when facing non-integrated, competing gas stations which account for 30% of the German market.

Related to the last point, another reason to investigate the wholesale market for fuel is complaints by non-integrated gas stations. They claim that vertically integrated firms hinder competition downstream by setting unfair prices on the wholesale level. The refineries were accused to sell gas to independent gas stations at a higher price than the own gas stations request from the final consumer.[4]

In my thesis, I studied concentrated markets and the anticompetitive effects of changes in parameters and conditions using laboratory experiments. As the example

[1]See *Fuel Sector Inquiry*, May 1, 2011, available at `http://www.bundeskartellamt.de/EN/Economicsectors/MineralOil/mineraloil_node.html`

[2]For more information see details available at `http://www.bundeskartellamt.de/EN/Economicsectors/MineralOil/MTU-Fuels/mtufuels_node.html`, the report after one year is provided here: `http://www.bundeskartellamt.de/SharedDocs/Publikation/DE/Berichte/Ein_Jahr_MTS-K_Marginalsp.pdf?__blob=publicationFile&v=10`, the report after two years at: `http://www.bundeskartellamt.de/SharedDocs/Publikation/DE/Berichte/Zweiter_Jahresbericht_MTS-K.pdf?__blob=publicationFile&v=3`.

[3]See *Launch of Sector Inquiry into Refineries and Oil Wholesale Sector*, September 27, 2012, available at `http://www.bundeskartellamt.de/SharedDocs/Meldung/EN/Pressemitteilungen/2012/27_09_2012_SU-Raffinerien.html?nn=3589696`

[4]See *Bundeskartellamt examines cases of unfair hindrance of independent petrol stations*, April 4, 2012, available at `http://www.bundeskartellamt.de/SharedDocs/Meldung/EN/Pressemitteilungen/2012/04_04_2012_Freie-Tankstellen.html?nn=3589696`

shows, in “real” markets, many influences are present at the same time. It seems not easy to uncover the mechanisms behind observations and derive causal inferences. Why did the “Markttransparenzstelle für Kraftstoffe” not result in lower average prices? Do consumers search non-optimally? Is it evidence in favor of collusion among gas stations? Might collusion on the wholesale level be the reason for stable average gas prices? Or do integrated firms hinder competition downstream? In order to implement policy interventions targeting an increase in consumer surplus it is insightful to analyze the effects isolated.

The second chapter of my thesis “Communication in Vertical Markets: Experimental Evidence” is co-authored with Hans-Theo Normann and Christopher M. Snyder. We consider an upstream monopolist which supplies two non-integrated downstream firms. The monopolist 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-in-Nash bargaining fits the pattern of shifting surpluses well. Using third-party coders, unsupervised text mining, among other approaches, we uncover features of the rich chat data that are correlated with market outcomes. We conclude with a discussion of the antitrust implications of open communication in vertical markets.

The third chapter “Search Costs in Concentrated Markets – An Experimental Analysis” is co-authored with Torben Stühmeier and Tobias Wenzel. We experimentally study the role of search cost in duopoly markets where sellers may be able to coordinate pricing decisions. We vary the level of search cost and whether sellers can communicate. While we find that consumers are more likely to invest in search when the cost is reduced, we find that a reduction of search cost does not influence prices. This effect does not change with the availability of seller communication. Our results suggest that policy interventions that aim to increase the competitiveness of markets via reducing search cost may not be effective in concentrated markets.

The fourth chapter is called “Reputation and Foreclosure with Vertical Integration – Experimental Evidence”. I study the role of reputation building on input foreclosure in vertically related markets. In one-shot interactions, upstream firms can choose to build a reputation by revealing their price history to the current upstream competitor. In particular, integrated firms can establish a reputation to foreclose the input market—an outcome that would otherwise not be tenable due to a commitment problem. I get three main results: First, withdrawal from the input market is three times more common with reputation building of the integrated firm. Second, the anticompetitive effects of vertical integration are much stronger when the integrated firm builds a reputation. Third, integrated firms choose to build a reputation significantly more often than non-integrated firms. Withdrawal of the integrated firm which results in monopolization upstream occurs in markets with reputation building ten times more often.

Finally, in the fifth chapter, I conclude my thesis.

References

Camerer, Colin. (2011) "The promise and success of lab-field generalizability in experimental economics: A critical reply to Levitt and List," Working Paper.

Chamberlin, Edward H. (1948) "An experimental imperfect market," *The Journal of Political Economy*, 1: 95–108.

Falk, Armin and Ernst Fehr. (2003) "Why labour market experiments?" *Labour Economics* 10(4): 399–406.

Falk, Armin and James J. Heckman. (2009) "Lab experiments are a major source of knowledge in the social sciences," *science*, 326(5952): 535–538.

Levitt, Steven D. and John A. List. (2007a) "Viewpoint: On the Generalizability of Lab Behaviour to the Field," *Canadian Journal of Economics*, 40(2): 347–370.

Levitt, Steven D. and John A. List. (2007b) "What do Laboratory Experiments Measuring Social Preferences Reveal about the Real World," *Journal of Economic Perspectives*, 21(2): 153–174.

Levitt, Steven D. and John A. List. (2008) "Homo Economicus Evolves," *Science*, 319(5865): 909–910.

Plott, Charles R. (1982) "Industrial organization theory and experimental economics," *Journal of Economic Literature* 20(4): 1485–1527.

Smith, Vernon L. (1962) "An experimental study of competitive market behavior," *The Journal of Political Economy*, 1: 111–137.

Smith, Vernon L. (1980) "Relevance of laboratory experiments to testing resource allocation theory," *Evaluation of Econometric Models*, Academic Press, 345–377.

Chapter 2

Communication in Vertical Markets: Experimental Evidence

Co-authored with Hans-Theo Normann and Christopher M. Snyder

2.1 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 an out-of-equilibrium belief and other firms have to conjecture what this belief is (or what the distribution of beliefs are in the case of heterogeneous beliefs). Communication could resolve some of this strategic uncertainty. 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,[1] 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 leaves upstream profits essentially unchanged; the increase in industry profit accrues almost entirely to downstream firms.

That different communication treatments led to dramatically different divisions of surplus between upstream and downstream firms initially surprised us as we had not designed the treatments to look for such effects. In Section 2.5, we propose a simple bargaining model providing a straightforward explanation. In the absence of communication, the upstream firm makes take-it-or-leave-it offers; opening a com-

[1]Lee and Whang's (2000) seminal article categorizes the kinds of information shared across vertical levels (inventories, sales, sales forecasts, order tracking, production plans, quality metrics), providing anecdotes for each involving well-known firms. Moving from anecdotal to survey evidence, 62% of the sample in Vanpoucke, Boyer, and Vereecke (2009) reported communicating with firms along the supply chain.

munication channel affords participating subjects an opportunity to bargain. We assume bargaining outcomes are given by the widely used "Nash-in-Nash" solution concept proposed by Horn and Wolinksy (1988), recently given non-cooperative foundations by Collard-Wexler, Gowrisankaran, and Lee (2016). According to this solution concept, each bargain maximizes the Nash product assuming that other bargains occurring simultaneously are efficiently consummated. Our bargaining model delivers the same pattern of surplus division observed in the experiments: opening a two-way communication channel in the model causes the upstream firm to lose bargaining power and moving to three-way communication reduces upstream surplus yet further.

Section 2.6 delves into the content of communication to uncover correlations between content features and market outcomes. To deal with the difficulty in quantifying the rich content data, we take several analytical approaches: counting messages, employing third-party coders, and using text-mining methods to extract keywords. The communication stage appears to function like a bargaining process, with discussions successfully converging to a contract that is the one that ends up being offered. When the upstream firm is successful at committing to the monopoly outcome, his or her messages tend to mention deals given to all both downstream firms and market prices. Commitment sometimes breaks down when a subject tries to strike an exclusive deal to sell the entire industry quantity, inevitably leading to oversupply as exclusion proves unenforceable.

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, the vertical contract exerts an externality on other downstream firms. The nature of the externality is different: rather than secretly oversupplying a rival, 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 downstream firms reduces entry-deterring exclusion.

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).[2] 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:

[2]The closest in this literature is contemporaneous research by Grandjean et al. (2014). They report on three-player experiments involving a different base game from ours but similar communication treatments. Their base game is a coordination game with multiple Pareto-ranked equilibria, in which the Pareto-optimal one susceptible to coalitional deviations. They find that play of the Pareto-optimal equilibrium is promoted by open communication similar to our *Three Chat*. For discussion of the possible beliefs in dynamic games with imperfect information, see Eguia et al. (2014).

final output is the result of a negotiation between upstream and downstream firms rather than being one firm's unilateral choice.[3,4]

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). 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.

2.2 Theoretical Framework

2.2.1 Market Model

Consider a simplified version of the model due to Rey and Tirole (2007).[5] The market has a vertical structure shown in Figure 2.1, with a monopoly upstream

[3]Several 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. Hinloopen and Soetevent (2008) and Bigoni et al. (2012) evaluate leniency 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.

[4] Cheap talk has been found to achieve superior outcomes in trust games (Charness and Dufwenberg, 2006). Although our vertically related markets are different from the standard trust game, they 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.

[5]Rey and Tirole (2007) is itself a simplified version of a number of earlier papers including Hart and Tirole (1990) and McAfee and Schwartz (1994). We modify Rey and Tirole (2007) in three ways. First, contracts here specify a single bundle at a fixed tariff rather than a tariff function. Second, downstream firms 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.

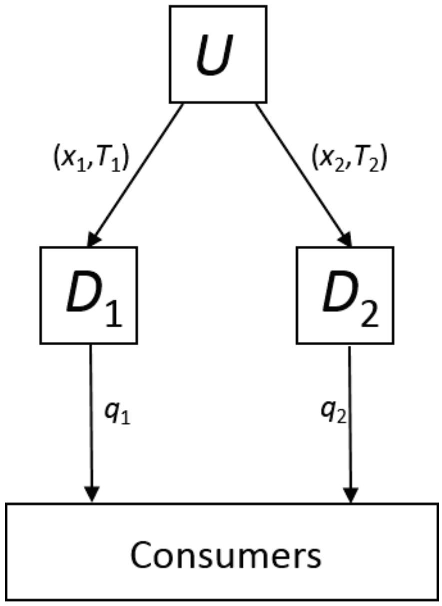

Figure 2.1: Vertical Structure

firm, U, and two downstream firms, D_i, $i = 1, 2$. The upstream firm produces an intermediate product at zero cost. The downstream firms transform this product on a one-for-one basis, also at zero cost, into a final good sold to consumers. Consumers have inverse demand $P(Q)$ for this homogeneous final good.[6]

The timing is as follows. First, U offers contracts (x_i, T_i) to each D_i specifying a quantity x_i and fixed tariff T_i. Second, the D_i simultaneously decide whether to accept $(a_i = 1)$ or reject $(a_i = 0)$ 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 $a_1T_1 + a_2T_2$ for U and $P(Q)q_i - a_iT_i$ for D_i.

To set some benchmarks, let $Q^m = \arg\max_Q P(Q)Q$ 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 2.1 were compressed into a single level. That is, defining the best-response function

$$BR(q) = \arg\max_{\tilde{q}} P(\tilde{q} + q)\tilde{q},$$

[6]Assume $P(Q)$ has properties ensuring that the Cournot game formed by compressing the vertical structure in Figure 2.1 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.

2.2.2 Commitment Problem with Secret Contracts

To understand the nature of the commitment problem with secret contacts, 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.[7]

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.[8]

[7] 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.

[8] While 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 most of the

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.

2.2.3 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 reneging on an explicit promise.[9] 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 communication among all three market participants seems likely to only enhance U's commitment power. U can describe exactly the symmetric offers

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.

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

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.

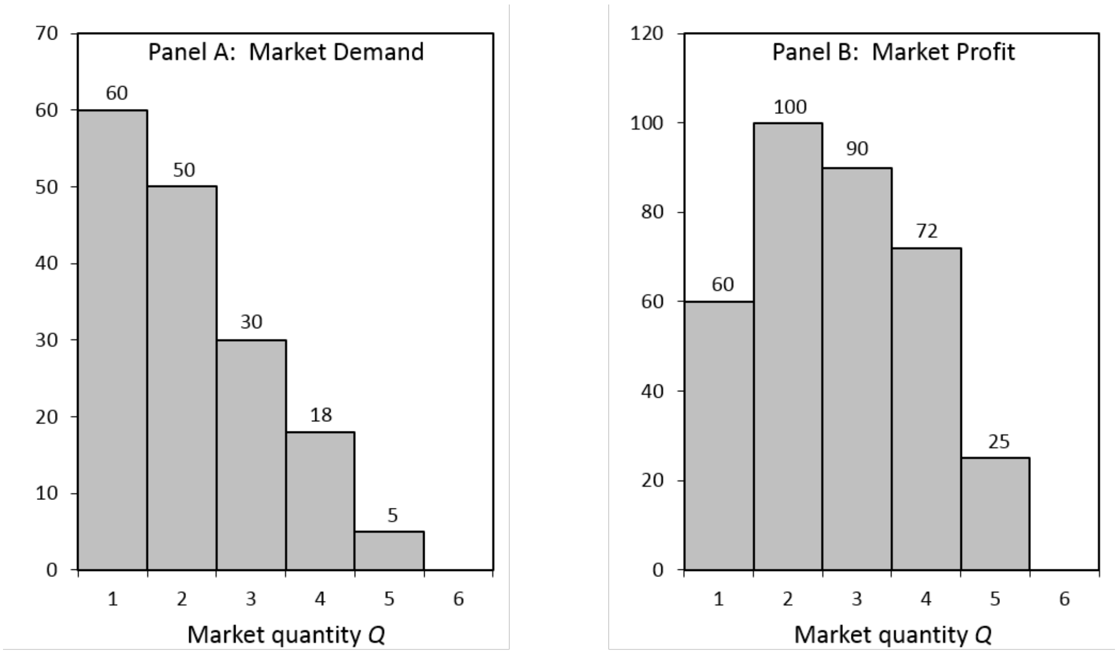

Figure 2.2: Experimental Market Demand and Profit

2.3 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 2.1, 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 2.2A. All firms produce at zero cost. Thus profits are $\pi_U = a_1 T_1 + a_2 T_2$ for U and $\pi_{Di} = P(Q)q_i - a_i T_i$ for D_i. Let $\pi_D = \pi_{D1} + \pi_{D2}$ denote total downstream profit and $\Pi = \pi_U + \pi_D$ denote market profit. Figure 2.1B 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 communication with each D_i. D_1 and D_2 could not communicate with each other, and D_i could not observe U's communications with his competitor. U had separate chat windows for each D_i on its screen; each D_i had only one chat window on its screen through which it communicated to U. In *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 chat stage prior to U's making contract offers. Except for threats to be carried out outside the lab or information that could be used to identify subjects, the content of the chat was unrestricted. The duration of the chat stage turned out to be a delicate experimental-design choice. We wanted to allow subjects enough time for full communication yet not so much that they became bored, possibly leading to distorted behavior. Equalizing chat time across treatments leads to different chat time per channel in treatments with different channels; we specified chat times balancing these considerations. In *Two Chat*, subjects had 90 seconds to chat during the first five rounds, reduced to 60 seconds 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*.[10] Subjects could not leave the chat stage before the time expired. Apart from the added chat stage, 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. Including the show-up fee, participants earned about 14 euros on average (19 euros on average for subjects playing the role of U, 12 for subjects playing the role of D).

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.

[10]The shorter chat time in certain rounds appears not to have unduly constrained subjects. For example, in rounds 6–15 of *Two Chat* having the shorter chat stage, a subject were typing at the last second in only 15% of the cases, similar to the 12% in rounds 1–5 with the longer chat stage. A majority of these last-second messages involved a confirmation or pleasantry. That boredom is a possible concern is embodied in one subject's chat message, "Booooooooooooooooooooooooooooooooooooooooooring."

2.4 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 2.1 can be interpreted as 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, provided in the bottom part of the table. We compute White (1980) heteroskedasticity-robust standard errors clustered by session, allowing for dependence among observations arising from the same set of interacting subjects, throughout the analysis.

Comparing the results for the *No Chat* treatment to those for *SECRAN* from Martin, Normann, and Snyder (2001) provides a consistency check. Total offered quantity $X = x_1 + x_2$ averaged 3.64 in *SECRAN*,[11] nearly identical to the 3.68 in *No Chat* (see the first column of Table 2.1). 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.

[11] The 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. Martin, Normann, and Snyder (2001) report results for all rounds, not in Table 2, but in Figures 3–6, in the form of histograms.

Table 2.1: Regressions Examining Differences in Means

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

Notes: Each column is an ordinary least squares regression. Specification includes an exhaustive set of treatment indicators (*No Chat*, *Two Chat*, *Choose Chat*, *Three Chat*) and omits the constant, allowing one to read coefficients as sample means. 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.

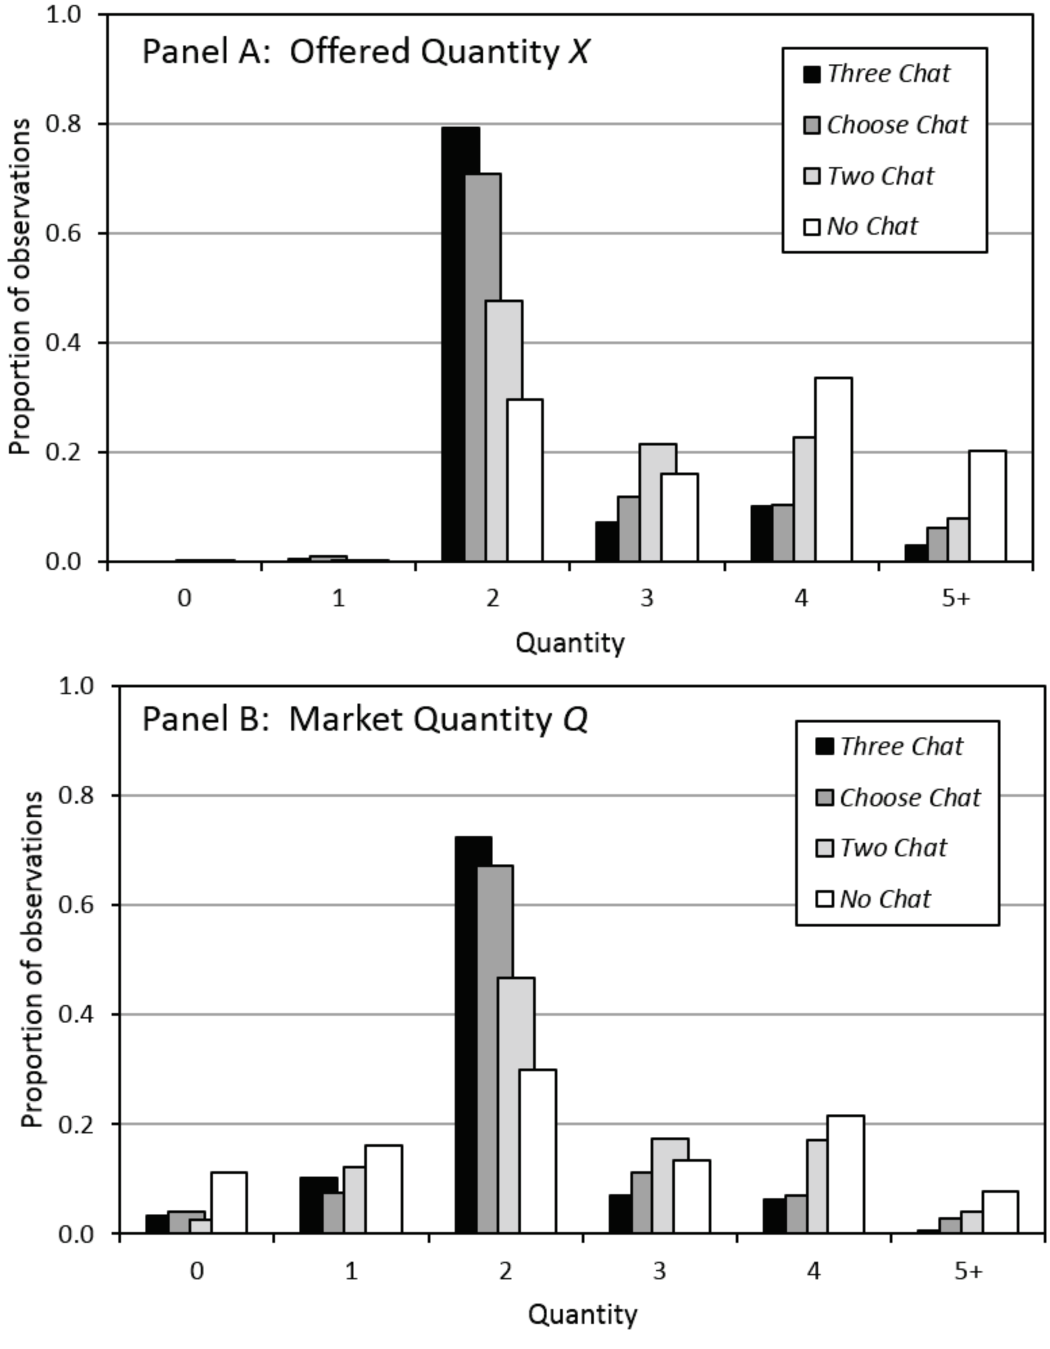

Figure 2.3: Quantity Histograms

2.4.1 Offered Quantity

We begin by analyzing total offered quantity, X. This single variable captures whether U is able to solve the commitment problem. Table 2.1 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 of 2. These results are consistent with more open communication facilitating commitment and monopolization.

The bottom part of the table 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 $X = 2$ (although a formal statistical test rejects equality at the 1% level).[12]

Figure 2.3 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 gray 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 2.2 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 2.3A in the $X = 2$ bin) directly from the coefficients on the treatment indicators. The reported standard errors 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.[13] 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 Chat* is 23 percentage

[12] As we will see, the mean of market quantity Q, 2.05 in *Three Chat*, is yet closer to the monopoly level of 2, the difference now only significant at the 6% level.

[13] As we will see in Table 2.3, the difference is significant at the 10% level after dropping the first five rounds, reflecting noisier play, from the sample.

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.

The fifth column of results in Table 2.2 measures symmetry implicit in offered quantities. It presents estimates from a linear model of the probability that the two contract offers involve symmetric quantities, $x_1 = x_2$. As noted in Section 2.2, theory predicts that the *No Chat* treatment should yield a symmetric equilibrium whether players hold symmetric or passive beliefs—the beliefs were shown to affect equilibrium quantities, not the symmetry between them. The estimate on the *No Chat* indicator implies that 68% of the offers in that treatment involve symmetric quantities. A large majority of observations thus comport with the symmetry prediction. Yet from a more pessimistic view, almost a third of the observations are asymmetric (perhaps not an overwhelming rejection of the theory given the noisy nature of experimental play).

That off-equilibrium-path outcomes are actually observed in the experiment provides an opportunity to learn about out-of-equilibrium beliefs. In Section 2.2, we hypothesized that, in an environment of heterogeneous and fluid beliefs, communication could serve to coordinate players on symmetric beliefs, which are beneficial for monopolization. The quantity offers observed by downstream firms in the *Two Chat* treatment would not justify their shifting toward more symmetric beliefs. The estimate on the *Two Chat* indicator in Table 2.2 shows that the percentage of symmetric offers did not increase but in fact slightly declined relative to *No Chat*. Evidently the private communication channel helps with monopolization but not by promoting symmetric beliefs, if anything impairing symmetry. By contrast, 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 last column of results 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.

Table 2.2: Linear Probability Models of Outcome Variables

	Measuring monopolization				Measuring symmetry	
	Offered quantity, X		Market output, Q			
	$X = 2$	$X \geq 4$	$Q = 2$	$Q \geq 4$	$x_1 = x_2$	$x_1 = x_2$, $T_1 = T_2$
No Chat	0.30***	0.54***	0.30***	0.29***	0.68***	0.55***
	(0.06)	(0.06)	(0.04)	(0.07)	(0.01)	(0.01)
Two Chat	0.48***	0.30***	0.47***	0.21***	0.63***	0.39***
	(0.10)	(0.07)	(0.07)	(0.05)	(0.07)	(0.05)
Choose Chat	0.71***	0.17***	0.67***	0.10***	0.76***	0.65***
	(0.06)	(0.04)	(0.06)	(0.03)	(0.04)	(0.05)
Three Chat	0.79***	0.13***	0.73***	0.07**	0.88***	0.76***
	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(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						
Two Chat – *No Chat*	0.18	−0.23**	0.17*	−0.08	−0.05	−0.16***
	(0.12)	(0.09)	(0.08)	(0.09)	(0.07)	(0.05)
Choose Chat – *No Chat*	0.41***	−0.37***	0.37***	−0.19**	0.08*	0.10*
	(0.09)	(0.08)	(0.07)	(0.07)	(0.04)	(0.05)
Three Chat – *No Chat*	0.50***	−0.41***	0.43**	−0.22***	0.21***	0.21***
	(0.07)	(0.06)	(0.05)	(0.04)	(0.02)	(0.03)
Choose Chat – *Two Chat*	0.23*	−0.14	0.21**	−0.11*	0.13	0.26***
	(0.12)	(0.08)	(0.09)	(0.06)	(0.08)	(0.07)
Three Chat – *Two Chat*	0.32***	−0.17**	0.26***	−0.14**	0.25***	0.37***
	(0.10)	(0.07)	(0.07)	(0.05)	(0.07)	(0.05)
Three Chat – *Choose Chat*	0.08	−0.03	0.05	−0.03	0.12**	0.11*
	(0.07)	(0.05)	(0.07)	(0.03)	(0.04)	(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. Specification includes an exhaustive set of treatment indicators (*No Chat*, *Two Chat*, *Choose Chat*, *Three Chat*) and omits the constant, allowing one to read coefficients as sample frequencies. 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.

2.4.2 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 2.1 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.

Definitive inferences are difficult to draw from the raw means of T_i, however, because x_i varies systematically across treatments as well.[14] To purge these quantity effects, the third column of Table 2.1 restricts 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. We will return to this result in Section 2.5, showing how it can be rationalized in a standard bargaining model.

2.4.3 Acceptance Behavior

Having analyzed upstream behavior, we next turn to downstream behavior, embodied in the acceptance decision a_i in Table 2.1. The acceptance rate rises from 70% in *No Chat* to 85% in *Two Chat* to 89% in *Three Chat*. Table 2.1 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.

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 2.1 sheds light on how

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

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.

2.4.4 Market Output

The rest of the variables for which we provide summary statistics in Table 2.1 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 2.3B. 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 2.2, 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 2.1, 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 2.2, 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.

2.4.5 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 2.1 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 (although a formal statistical test rejects equality at the 1% level).

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 2.1, 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 further in Section 2.5, the move from *Two Chat* to *Three Chat* could represent a change in the structure of bargaining, which, if bargaining is characterized by the Nash-in-Nash solution, ends up eroding U's bargaining surplus. This bargaining theory explains the fall in s_U in *Three Chat*.

So far we have examined how s_U changes across treatments. We have not remarked yet on the fact that the mean of s_U in *No Chat* is 63%, considerably less than the 100% theory would predict for that treatment in which U makes take-it-or-leave-it offers. It is standard in ultimatum games to find a more equitable split of surplus than the subgame-perfect equilibrium predicts (see Roth's 1995 review). Martin, Normann, and Snyder (2001) found similar results in their analogous *SECRAN* treatment, devoting all of Section 6 to evaluating alternative explanations.

Downstream firms gain both in absolute and relative terms from more and more open communication. Table 2.1 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 2.1. Downstream profit is so high in *Three Chat* that they obtain a majority of the profit (53% compared to U's 47%).

2.4.6 Consumer Surplus

The last column of Table 2.1 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*. 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 2.3B 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.

2.4.7 *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 2.1, 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 2.1, 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.

2.4.8 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 2.3 repeats the regressions from Table 2.1 interacting the treatment indicators with indicators for the initial and end periods. For example, $\textit{No Chat}_0$, is the interaction between the *No Chat* indicator and an indicator for rounds 1–5, and $\textit{No Chat}_1$ 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.

Table 2.3: Trends in Treatment Effects

Dependent variable:	X	T_i		a_i		Q	Π	π_U	π_D	s_U	CS
Sample:	Full	Full	$x_i = 1$	Full	$x_i = 1$	Full	Full	Full	Full	$\Pi > 0$	Full
No Chat$_0$	3.80***	38.2***	35.0***	0.66***	0.75***	2.33***	64.4***	45.5***	19.0***	0.66***	36.5***
	(0.31)	(3.1)	(2.2)	(0.01)	(0.04)	(0.19)	(4.1)	(3.0)	(2.3)	(0.02)	(4.9)
No Chat$_1$	3.62***	32.9***	32.4***	0.72***	0.77***	2.54***	70.3***	45.3***	25.0***	0.61***	41.3***
	(0.15)	(0.4)	(1.2)	(0.02)	(0.04)	(0.19)	(1.3)	(1.7)	(2.7)	(0.03)	(6.4)
Two Chat$_0$	3.32***	33.5***	26.8***	0.84***	0.86***	2.72***	78.9***	53.5***	25.4***	0.61***	42.5***
	(0.25)	(2.2)	(1.0)	(0.04)	(0.03)	(0.16)	(3.4)	(2.6)	(2.9)	(0.03)	(5.4)
Two Chat$_1$	2.80***	30.4***	26.9***	0.85***	0.85***	2.37***	84.3***	49.9***	34.5***	0.59***	29.9***
	(0.17)	(2.3)	(1.3)	(0.04)	(0.04)	(0.12)	(2.4)	(2.7)	(2.1)	(0.03)	(4.2)
Choose Chat$_0$	2.90***	29.6***	27.7***	0.85***	0.90***	2.32***	82.1***	48.2***	33.9***	0.57***	28.9***
	(0.23)	(1.4)	(1.2)	(0.03)	(0.01)	(0.09)	(3.6)	(1.8)	(4.0)	(0.03)	(4.0)
Choose Chat$_1$	2.37***	25.6***	23.7***	0.92***	0.88***	2.14***	90.1***	45.2***	44.8***	0.50***	19.5***
	(0.13)	(1.3)	(1.2)	(0.01)	(0.02)	(0.11)	(1.9)	(2.9)	(3.3)	(0.03)	(3.9)
Three Chat$_0$	2.86***	30.0***	28.4***	0.83***	0.87***	2.20***	83.7***	45.7***	37.9***	0.54***	24.7***
	(0.25)	(1.5)	(1.7)	(0.02)	(0.03)	(0.10)	(2.2)	(3.0)	(3.5)	(0.03)	(3.2)
Three Chat$_1$	2.18***	22.5***	21.9***	0.92***	0.85***	1.98***	92.4***	40.8***	51.6***	0.44***	13.4***
	(0.03)	(0.7)	(0.4)	(0.00)	(0.02)	(0.03)	(0.5)	(1.0)	(1.4)	(0.01)	(1.0)
Other controls	None	None	None	None	$\tilde{T}_i, \tilde{T}_i^2$	None	None	None	None	None	None
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											
No Chat$_1$ − *No Chat*$_0$	−0.18	−5.3*	−2.6	0.07***	0.02	0.20***	5.8	−0.2	6.0	−0.05	4.8
	(0.19)	(2.9)	(2.9)	(0.02)	(0.06)	(0.05)	(4.0)	(4.6)	(3.8)	(0.05)	(3.4)
Two Chat$_1$ − *Two Chat*$_0$	−0.52***	−3.1***	0.2	0.01	−0.01	−0.35***	5.5***	−3.6**	9.1***	−0.03	−12.6***
	(0.13)	(0.4)	(0.3)	(0.01)	(0.01)	(0.11)	(1.4)	(1.6)	(0.4)	(0.02)	(3.3)
Choose Chat$_1$ − *Choose Chat*$_0$	−0.53***	−3.9**	−4.0***	0.07*	−0.03	−0.18	8.0***	−2.9	10.9***	−0.07***	−9.4*
	(0.15)	(1.5)	(1.1)	(0.03)	(0.02)	(0.15)	(2.6)	(1.9)	(3.6)	(0.02)	(4.9)
Three Chat$_1$ − *Three Chat*$_0$	−0.68***	−7.5***	−4.7**	0.09***	−0.01	−0.22*	8.8***	−4.9*	13.7***	−0.10***	−11.3**
	(0.27)	(0.9)	(2.1)	(0.02)	(0.02)	(0.12)	(1.8)	(2.6)	(2.4)	(0.02)	(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 Chat*$_0$ is the interaction between *No Chat* and the initial period consisting of rounds 1–5, and *No Chat*$_1$ is the interaction between *No Chat* and the end period consisting of rounds 6–15. Specification includes an exhaustive set of treatment indicators (*No Chat*, *Two Chat*, *Choose Chat*, *Three Chat*) and omits the constant, allowing one to read coefficients as sample means. 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 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.

2.5 Rationalizing Effects on Surplus Division

We designed our experiments to test the hypothesis that communication can help vertically related firms monopolize a market by solving a commitment problem. The results, as seen, bear this hypothesis out. We found another set of results for which we did not have *a priori* hypotheses—results related to the division of surplus between upstream and downstream firms varied across communication treatments—which were strong, systematic and beg explanation. It is worth recapitulating these results. They were clearest in the case of tariff levels (T_i) accompanying monopoly quantity offers ($x_i = 1$). Holding offered quantities constant fixes total surplus, so that changes in T_i represent a free transfer of surplus between upstream to downstream. We found that the move from *No Chat* to *Two Chat* reduced T_i by 6.4 and from *Two Chat* to *Three Chat* reduced T_i by a further 3.2. The results for the tariffs are mirrored in profit shares: moving from *No Chat* to *Two Chat* reduced s_U by 3 percentage points and from *Two Chat* to *Three Chat* by a further 11 percentage points. In sum, the move from *No Chat* to *Two Chat* to *Three Chat* shifted surplus from the upstream to downstream firms.

In this section we show that these experimental results can be rationalized as bargaining effects in a standard bargaining model. We assume that opening up a communication channel sets up a bargaining process among the subjects involved. We further assume that bargaining is characterized by the Nash-in-Nash solution proposed by Horn and Wolinksy (1988), now a widely used bargaining concept in applied industrial organization as evidenced by the scores of references in Collard-Wexler, Gowrisankaran, and Lee's (2016) review paper. The Nash-in-Nash solution (short for "Nash bargains nested within a Nash equilibrium") turns out to predict the precise pattern of the variation in the division of surplus across communication treatments we observe in the experiment.

2.5.1 Nash-in-Nash Bargaining

This subsection provides some theoretical background on Nash-in-Nash bargaining and its application to our experimental setting. To focus exclusively on implications for surplus division, we assume away the commitment problem for now by positing that contracts offer $x_i = Q^m/2 = 1$ unit 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). We assume that opening a communication channel sets up a bargaining process among the subjects involved. Whatever contracts parties agree to in the communication stage are the contracts U then offers. This is a key assumption that requires some discussion. In theory, U could regard the chat as cheap talk, and instead make the take-it-or-leave-it offer it would have in the absence of communication. We can provide several explanations for why chat settles on the actual contract offer. The results from ultimatum games suggest that responders react negatively to offers violating their internal expectations of fairness; it is likely the D_i would react even more negatively to contracts violating their explicitly stated expectations. If U faces a sufficiently high cost of reneging on a promise (see the references to recent empirical work measuring lying aversion in footnote 9), he or she will be inclined to offer the agreed-to contract. Whatever the reason, chat did settle on the offered contract in a large majority of cases, as we will document in Section 2.6.2.

In the absence of communication in *No Chat*, there is no bargaining. In this case, U issues take-it-or-leave-it offers to the D_i, allowing it to extract all of the gross profit ($\Pi^m/2 = 50$) that D_i earns in the monopoly outcome with an equilibrium tariff of $T_i^* = 50$. 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 50) 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 et al., 1991).

The bilateral communication channels in *Two Chat* set up two separate bargaining processes, one between U and each D_i. The Nash-in-Nash solution posits that U and D_1 reach an agreement maximizing their joint payoffs, arriving at a tariff maximizing their Nash product assuming that the other bargain between U and D_2, occurring simultaneously, also reaches an agreement maximizing those parties' joint payoffs. That U and D_1 split surplus between them according to the Nash product is the Nash bargaining alluded to by "Nash-in-Nash"; that U and D_1 assume the other bargain is consummated in the rational (jointly efficient) way is the Nash equilibrium alluded to by "Nash-in-Nash." In our setting, it can easily be seen that acceptance ($a_i = 1$) and trade of one unit ($q_i = 1$) is jointly efficient for U and each D_i because it is true regardless of what happens in the other bargain: if U and D_1 trade, their joint surplus equals 60 if U and D_2 happen not to come to an agreement and 50 if they do. Given that U and D_2 end up trading one unit, the joint payoff to be split between U and D_1 equals 50. Their Nash product is $T_i^{1/2}(50 - T_i)^{1/2}$ if they have equal bargaining power. The equilibrium tariff maximizing this Nash product is $T_1^* = 25$ (and $T_2^* = 25$ by symmetry). More generally, letting $\alpha \in [0, 1]$ be U's bargaining power vis-á-vis D_i, the relevant Nash product is $T_i^\alpha(50 - T_i)^{1-\alpha}$, maximized by equilibrium tariff $T_i^* = 50\alpha$.

That the outcome of the bilateral bargains is characterized by the Nash-in-Nash solution is another key assumption behind our explanation of surplus division. This assumption is less strong than it may appear. Collard-Wexler, Gowrisankaran, and Lee (2016) provide sufficient conditions for the Nash-in-Nash solution to be the limit of the unique equilibrium of a generalized Rubinstein (1982) process in which upstream and downstream firms alternate offers.[15] In Appendix A1, we show that these sufficient conditions for uniqueness are satisfied in our setting.

Moving to *Three Chat*, the open communication channel in this treatment sets up a single three-way bargain. Since there is no other simultaneously occurring bargain in this case, the Nash-in-Nash solution concept reduces simply to Nash bargaining, maximizing the Nash product of three players' payoffs. If they have equal bargaining power, this Nash product simplifies to $(2T_i)^{1/3}(50-T_i)^{2/3}$, maximized by equilibrium tariff $T_i^* = 50/3 = 16.\overline{6}$. More generally, they may have asymmetric bargaining powers. There are several ways to generalize bargaining weights for Nash product involving more than two players. A natural generalization 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:[16]

$$(dT_i)^{\frac{\alpha}{\alpha+d(1-\alpha)}} \prod_{i=1}^{d} \left(\frac{\Pi^m}{d} - T_i \right)^{\frac{1-\alpha}{\alpha+d(1-\alpha)}} = \left[(dT_i)^{\alpha} \left(\frac{\Pi^m}{d} - T_i \right)^{d(1-\alpha)} \right]^{\frac{1}{\alpha+d(1-\alpha)}}. \quad (2.1)$$

Maximizing this expression and substituting the experimental parameters $\Pi^m = 100$ and $d = 2$ yields equilibrium tariff $T_i^* = 50\alpha/(2-\alpha)$ in the *Three Chat* treatment.

[15] The limit is the usual one in analyzing Rubinstein (1982) processes, taking the time between offers to zero. The authors restrict attention to perfect Bayesian equilibria with an additional refinement.

[16] Laurelle and Valenciano (2008) provide a noncooperative foundation for the generalized Nash bargaining formulae in (2.1). 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.

Perhaps the most natural alternative to our specification of bargaining weights maintains a constant ratio between U's bargaining weight and the sum of all downstream firms' bargaining weights rather than an individual downstream firm's. One can show that the model would predict equal tariffs in the *Two Chat* and *Three Chat* treatments under this variant, which is rejected by the experimental results, whereas our specification rationalizes them. Indeed any nontrivial linear combination of our and the alternative bargaining weights would generate the comparative statics for tariffs observed across experimental treatments.

Comparing equilibrium tariffs across treatments, the model predicts T_i^* falls from 50 in *No Chat* to 25 in *Two Chat* to $16.\overline{6}$ in *Three Chat* if subjects have equal bargaining power. In the general case of asymmetric bargaining weights, T_i falls from 50 in *No Chat* to 50α in *Two Chat* to $50\alpha/(2-\alpha)$ in *Three Chat*. As long as $\alpha < 1$ so that U does not have all the bargaining power, the tariff is predicted to strictly fall from *No Chat* to *Two Chat* to *Three Chat* in the general case. It is clear in theory why moving from *No Chat* to *Two Chat* should reduce the tariff: moving from U's making take-it-or-leave-it offers to affording some bargaining power to D_i should be expected to reduce upstream and increase downstream surplus generally, regardless of the assumed bargaining concept, whether Nash-in-Nash or some other. The tariff reduction moving from *Two Chat* to *Three Chat* rests more heavily on the Nash-in-Nash assumption; as we will see, other solution concepts need not deliver this prediction. Behind the Nash-in-Nash solution is the intuition that including more downstream firms in the more comprehensive bargains just means that U has to split the surplus among more parties.

The intuition can be different with a different solution concept. Consider an alternative we propose here, which we will label Nash-in-Shapley bargaining. As Nash-in-Nash, Nash-in-Shapley posits that individual bargains are consummated efficiently assuming others are as well; the only difference is that the incremental surplus generated by each bargain is divided using the Shapley value rather than Nash product. In our setting, the two concepts turn out to yield identical outcomes in *Two Chat*. They diverge with *Three Chat*. With Nash-in-Nash, U is harmed by combining the separate bilateral bargains in one grand bargain because the surplus is fairly divided among whichever players happen to be "in the room." With Nash-in-Shapley, U benefits from combining bargains. The formula builds in the idea that if one of the downstream firm rejects its contract, the others observe this in real time and move from bargaining over the division of a surplus of 60 rather than 50. While Nash-in-Nash and Nash-in-Shapley both have reasonable economic intuition behind them, in the end it is an empirical question which fits the data better.

Table 2.4: Tariffs Predicted by Various Bargaining Models

	Mean $\hat{T}_i$ in subsample		
Bargaining model	*Two Chat*	*Three Chat*	BIC
Nash-in-Nash			
Posit $\alpha = 0.50$	25.0	16.7	70,351
NLLS estimate $\hat{\alpha} = 0.60$	30.2	21.6	6,998
Nash-in-Shapley			
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 $\epsilon_i = T_i - \hat{T}_i$ has standard normal distribution.

2.5.2 Structural Estimates

Table 2.4 provides structural evidence on how well these bargaining models fit the data on 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 Nash-in-Nash bargaining 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 Nash-in-Nash bargaining 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-in-Nash bargaining, the next two rows show fitted values for the Nash-in-Shapley alternative. The row with $\alpha = 0.5$ is the standard version of the Shapley value in which all permutations of players used to compute marginal contributions are equally 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. Appendix-A2 provides the details of this generalization, based on Kalai and Samet (1987). Non-linear least squares produces an estimate of $\hat{\alpha}$ of 0.39. In effect, the estimated version of Nash-in-Shapley bargaining 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 2.4 shows that the model of Nash-in-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 any of the alternatives to Nash-in-Nash bargaining with the estimated $\hat{\alpha}$ involve hundreds or thousands of points higher values of BIC.

Table 2.5: Descriptive Statistics on Message Counts

	Two Chat		*Choose Chat*				
	Channels separated	Channels combined	Private channels separated	Private channels combined	Open channel	All channels combined	*Three Chat*
Means of indicators for some message sent							
$Any\ Mes_U$	0.97	0.99	0.32	0.43	0.86	0.97	0.85
$Any\ Mes_D$	0.99	1.00	0.39	0.55	0.93	0.99	0.98
Any Mes	0.99	1.00	0.42	0.58	0.95	0.99	0.98
Means of number of messages sent							
$Num\ Mes_U$	2.5	5.0	0.8	1.5	3.0	4.6	2.4
$Num\ Mes_D$	3.0	6.0	1.0	2.0	6.6	8.6	6.0
Num Mes	5.5	11.0	1.8	3.5	9.6	13.1	8.3
Means of indicators for chat initiation							
$Init_U$	0.52	0.48	0.15	0.16	0.28	0.30	0.29
$Init_D$	0.50	0.57	0.27	0.42	0.70	0.72	0.71
Correlations across market levels							
$Any\ Mes_U$ with $Any\ Mes_D$	0.29	*a*	0.72	0.66	0.46	0.47	0.36
$Num\ Mes_U$ with $Num\ Mes_D$	0.43	0.43	0.72	0.76	0.61	0.38	0.50
Observations	690	345	690	345	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. [a]Correlation undefined because variance of $Any\ Mes_D$ equals 0.

2.6 Analysis of Chat Content

In this section we draw further insights about the effect of communication by analyzing the content of the chat itself. The rich content data does not lend itself to easy quantification (Kimbrough et al., 2008), so in this section we take a series of approaches to do so: counting messages, employing third-party coders, and mining the text for keywords. Unlike the results reported to this point, the results in this section should be interpreted as associations, not causal relationships.

2.6.1 Message Counts

Table 2.5 provides descriptive statistics on aspects of the unstructured chat text that are amenable to simple counting. Apart from *Three Chat*, the other chat treatments involve multiple communication channels operating simultaneously. To provide a full picture of the nature of chat in these treatments, we provide analyses separating and combining the channels in a series of columns.

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.

Table 2.6 regresses various experimental outcomes on variables characterizing the chat from Table 2.5 among others. The regressors are endogenous so their coefficients will not have causal interpretations, but will still reveal interesting correlations and provide some measure of the strength of these 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 level 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. It appears that vertical pairs sometimes resort to the private channel to cut side deals that secretly expand traded output, thereby expropriating surplus from the other downstream firm. If so, the descriptive statistics from Table 2.5 tell us that the D_i initiate many more of these side deals than U. That the D_i initiate private communication that ends up reducing aggregate downstream surplus is reminiscent of the equilibrium outcome in the related Prisoners Dilemma, in which players destroy joint surplus in equilibrium by finking on each other. What may come as more of a surprise is how infrequent the side deals are: as Table 2.5 shows, the D_i resort to the open channel twice as often as the private channel, so the option to use the private channel in *Choose Chat* does not destroy commitment power completely.

Table 2.6: Regressions on Message-Count Covariates

	Two Chat			*Choose Chat*				*Three Chat*	
	x_i	X	s_U	X	X	s_U	s_U	X	s_U
Constant	1.50***	3.04***	0.65***	2.81***	2.49***	0.50***	0.41***	2.92***	0.54***
	(0.09)	(0.23)	(0.10)	(0.20)	(0.18)	(0.04)	(0.03)	(0.18)	(0.04)
$Num\ Mes_U$	0.03	0.04	0.02*	0.04	0.02	0.02*	0.01	−0.08**	0.00
	(0.04)	(0.06)	(0.01)	(0.04)	(0.03)	(0.01)	(0.01)	(0.02)	(0.01)
$Num\ Mes_D$	−0.03***	−0.04**	−0.02	−0.05**	−0.04**	−0.00	−0.00	−0.06*	−0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.02)	(0.01)
$Init_U$	−0.00								
	(0.07)								
Any Private					0.59***		0.16***		
					(0.08)		(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.

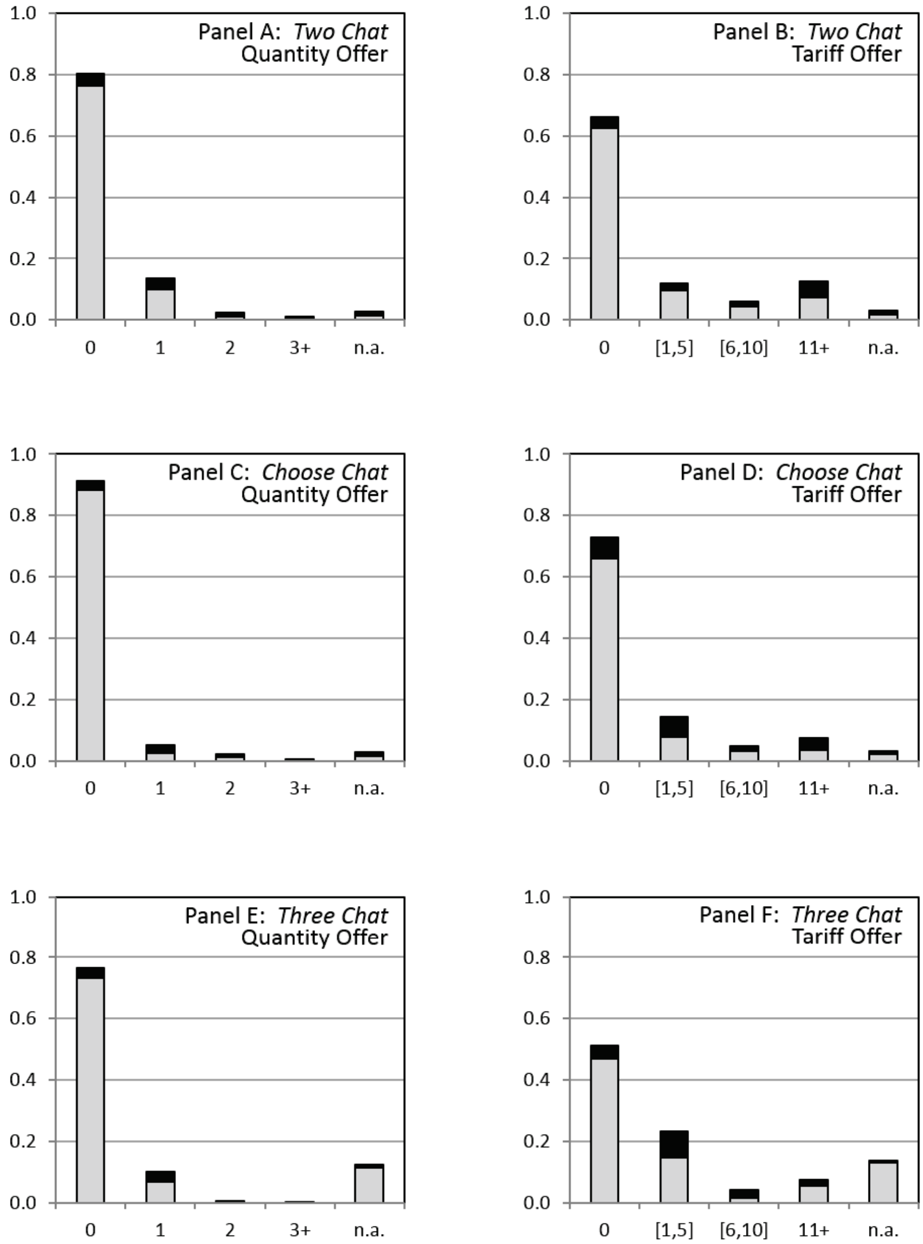

Notes: *Two Chat* sample consists of 350 contract offers, *Choose Chat* of 370, and *Three Chat* of 360, each of which is assessed by two coders. Grey bar is proportion of sample for which coder's guess of contractual variable agrees with other coder and black for which his guess disagrees with the other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

Figure 2.4: Accuracy and Mutual Agreement of Coded Chat

2.6.2 Coder Exercises

To probe more deeply into the chat content, in this section we report on several exercises using input from external coders. 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.

Our first use of the coder data is to determine whether chat content conveyed meaningful information about the terms of the contracts that would be offered that round. Figure 2.4 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. More typically, it appears that subjects used the chat stage to come to an agreement about contractual terms that are then reflected in U's offers.

Table 2.7 compares downstream acceptance behavior depending on whether U's contract fulfilled downstream expectations from the chat stage. The table restricts attention to just those observations that the coders' provided an integer guess for x_i and T_i and the coders' guesses matched for both. Presumably the coders' guess provides a good proxy of what downstream expectations were for U's contract offer.

The regression in each column implements a linear probability model in which a_i is the dependent variable, specified so that coefficients on the treatment indicators can be interpreted as average acceptance rates in each treatment. Column (1) shows that when downstream expectations are met, the contract is almost certain to be accepted across all treatments, more certain in more open communication treatments. In *Three Chat*, 100% of such contracts were accepted. Offers that do not fulfill expectations in columns (2) and (3) are accepted less often. We see that only 42% of offers that differ in both terms from expectations are accepted in *Three Chat*. As the the last column of the table shows, the reduction in acceptance rate from column (1) to column (3) is large and statistically significant for all three treatments. Following this last column down to the bottom part of the table, we see that this decline in acceptance is significantly greater for the most open form of communication (*Three Chat*) compared to the other two (*Two Chat* and *Choose Chat*). Downstream rejection seems to be a mechanism for enforcing agreements made in the chat stage, a mechanism that is strongest in the *Three Chat* treatment, suggestion why commitment is strongest in that treatment.

Table 2.7: Variation of Acceptance Rate with Fulfillment of Chat Expectations

	Offer terms matching coders' guess			
	Both x_i, T_i	One of x_i, T_i	Neither x_i, T_i	Difference
	(1)	(2)	(3)	(1) − (3)
Two Chat	0.92***	0.73***	0.58***	0.34***
	(0.01)	(0.05)	(0.01)	(0.02)
Choose Chat	0.96***	0.85***	0.67***	0.30**
	(0.01)	(0.04)	(0.11)	(0.11)
Three Chat	1.00***	0.85***	0.42***	0.58***
	(0.00)	(0.05)	(0.09)	(0.09)
Observations	603	156	50	
R^2	0.02	0.02	0.03	
Coefficient differences				
Choose Chat − *Two Chat*	0.04***	0.12	0.09	−0.05
	(0.01)	(0.07)	(0.11)	(0.11)
Three Chat − *Two Chat*	0.08***	0.12	−0.16	0.24**
	(0.01)	(0.07)	(0.09)	(0.09)
Three Chat − *Choose Chat*	0.04***	0.00	−0.25	0.29*
	(0.01)	(0.07)	(0.14)	(0.14)

Notes: Each column is an ordinary least squares regression using acceptance a_i as the dependent variable. Regression thus interpreted as linear probability model. Sample begins with the subset of observations from communication treatments that were subjected to coding (one complete session and five randomly selected periods from the three other sessions) and drops all but ones in which coders' integer guesses match each other for both x_i and T_i. First three columns consider different subsamples of this restricted sample depending on how many of the coders' guesses for terms x_i and T_i match U's offers. Specification includes an exhaustive set of indicators for the communication treatments (*Two Chat*, *Choose Chat*, *Three Chat*) and omits the constant, allowing one to read coefficients as sample frequencies. 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.

2.6.3 Mining Text for Keywords

Perhaps the deepest insight into the association between chat content and outcomes comes from the final analytical approach reported in this subsection. We are interested in determining whether there was something unique about the chat leading to monopoly offers ($X = 2$) compared to chat that did not. We proceed by using text-mining methods for extracting keywords from a body of text, referred to as a corpus.

To describe the methods, it is easiest to work through a concrete example. To make the comparisons as clean as possible, focus just the messages sent by U in the *Two Chat* treatment leading up to symmetric offers. This yields two corpora to compare, chat associated with low-quantity offers $x_1 = x_2 = 1$ (label this corpus L) and chat associated with high-quantity offers $x_1 = x_2 = 2$ (label this corpus H). We measure the "keyness" of word w in corpus L relative to H using Huerta's (2008) relative rank difference, computed as follows. Generate ranks $r_L(w)$ for all words w in corpus h according to frequency, ranging from 1 for the most common to $\underline{r}_L$ for the least. Similarly, generate ranks $r_H(w)$ for all words w in corpus h. The difference in the rank of w in corpus L relative to corpus H is defined as

$$\frac{r_H(w) - r_L(w)}{r_L(w)}. \tag{2.2}$$

When w does not appear in H, $\underline{r}_H$ will substitute for $r_H(w)$. Huerta's (2008) measure captures two essential properties for word w to be key: first, that w appears more frequently in L than H (captured by the numerator of expression (2.2)); and, second, that w is commonly used in L (captured by the denominator of (2.2)).

The keywords extracted from corpus L relative to H using this method are provided in the first box in Table 2.8. The box shows keywords that are among the top 50 most common in L for which expression (2.2) exceeds 3.5, omitting conjunctions, prepositions, and articles. The first box should be compared to the third box in the first row, providing keywords from the same exercise swapping the corpora (i.e., keywords from H relative to L). As one would expect—but reassuring that the extraction method is giving sensible results—words related to the number of units in the offer ("unit," "one," and "1" in the first box; "2" in the third box) emerge as key. The rest of the words tell us something deeper about chat content. By far the most key word in the first box, with a relative rank difference of 43.0, is *ihr*, the plural form of you in German, translated in the table as "you both." Although U is privately communicating with a single retailer in this *Two Chat* treatment, this word apparently bolsters commitment by indicating that whatever is being written applies to the other retailer (and presumably vice versa). Keywords "also" and "both" might have this same effect. Use of verbs conjugated in the third person ("gets," "gives") presumably reference the other retailer. Together, these keywords suggest that commitment to the monopoly outcome may be bolstered by referring to the retailer left out of the private communication channel in *Two Chat*. Other words may contribute to commitment as well: "price" get retailers thinking about the high market price that can result from monopolization; "guaranteed" is a direct reference to commitment.

Quite a different picture emerges in the third box. These keywords suggest a conspiracy hatched in the private communication channel for them to trade "2" units at a tariff of "50," splitting the profits between them "50:50," leaving the other retailer with "nothing." The first-person pronoun "we" and verb form *haben*, translated "(we) have," seem to contribute to this conspiracy. The loss of commitment in these cases (recall this is only a minority of observations) appears to due to an overreach, U's attempt to duplicate the monopoly industry outcome with each retailer.

The second and fourth boxes in the first row of Table 2.8 repeats the exercise for *Two Chat* messages sent by downstream firms. The keyword lists are shorter and the relative rank differences smaller than just seen for messages sent by U, suggesting that U was the driving force behind the direction chat took. The number of units in the offer shows up in the keyword lists but few other words besides.

The last row of the table looks at chat content in the *Three Chat* treatment. The keywords associated with low-output offers are similar to those in *Two Chat*. We see references to the single unit involved in the offer ("1," "unit," "one") and "you both" again shows up as key. Turning to the keywords associated with the high-output offers, besides the reference to the number of units ("units," '2"), there is less of a clear pattern. Gone are the conspiratorial keywords seen in *Two Chat*; of course the absence of a private communication channel in *Three Chat* would preclude such a conspiracy. Instead, we see words that could reflect a frustration at being unable to achieve a satisfactory outcome, such as "loss," "only," and "hampers," likely the effect rather than the cause of an inability to commit.

Table 2.9 turns from symmetric quantity offers to report keywords extracted from chat leading up to asymmetric quantity offers. Whereas before we treated all chat exchanged in a market in a round was treated together, here we separate the chat in the two private channels, putting the chat in the channel with the low offer ($x_i = 1$) in corpus L and the chat in the channel with the high offer in corpus H, so the interpretation of L and H is slightly different in this than the previous table. Because so few offers in *Three Chat* were asymmetric, we restrict attention to the *Two Chat* treatment. The length of the keyword list is the reverse of before, now much longer for messages sent by downstream firms than upstream, suggesting that what downstream firms write generates offer asymmetry. As expected, the number of units in the offer constitute some of the keywords. Chat leading to the low offer features references to the two retailers ("each," "us") as well as a consideration of counterfactuals ("otherwise," "would"). Chat leading to the high offer suggests selfish considerations, referring to "me" rather than "us," perhaps indicating that the other retailer receive "nothing."

Overall, the text-mining exercise shows that when U was able to successfully commit to the monopoly outcome, the messages it sent featured references to the other retailer, to market outcomes, and to guarantees. The trigger breaking down commitment in *Two Chat* in some instances appears to have been one of the bilateral pair suggesting an exclusive deal cutting out the other retailer. U sometimes tried to initiate purportedly exclusive deals with both retailers simultaneously, leading to $X = 4$ units. When a downstream firm was the initiator, it appears that U was sometimes happy to play along but then not follow through on the exclusion, leading to asymmetric offers and a total offer of $X = 3$ units. The finding that exclusive deals exacerbate the commitment problem is not anticipated by theory. In the model of Section 2.2, commitment is eroded not by active attempts to cut exclusive deals but by passive beliefs that the bilateral pair can do nothing to reduce the amount sold by the rival, so they may as well best respond. The chat in *Three Chat* conforms more closely with the theory. Cutting special deals is difficult in that treatment because all chat is public. The chat associated with unsuccessful attempts at monopolization in *Three Chat* appears to reflect frustration at an inability to escape equilibrium forces leading to a Pareto inferior outcome for the industry.

Table 2.8: Keywords Mined from Chat Leading to Symmetric Offers

	Low-quantity offers with $x_1 = x_2 = 1$ (corpus L)				High-quantity offers with $x_1 = x_2 = 2$ (corpus H)			
	Messages sent by U		Messages sent by D_i		Messages sent by U		Messages sent by D_i	
Treatment	Word	Rank diff. relative to H	Word	Rank diff. relative to H	Word	Rank diff. relative to L	Word	Rank diff. relative to L
Two Chat	you both	43.0	times	7.1	50:50	19.4	risk	9.5
	unit	23.3	other	6.3	50	18.7	2	5.0
	one	7.9	unit	4.3	2	14.0		
	also	7.7			nothing	9.0		
	1	6.3			profit share	7.2		
	me	5.9			we	6.3		
	ever	4.9			(we) have	5.2		
	(you) have	4.5						
	gets	4.5						
	gives	4.1						
	price	3.7						
	guaranteed	3.7						
	both	3.7						
Three Chat	1	29.0	40	21.0	other	51.0	pay	9.9
	20	10.6	:)	5.8	units	31.4	2	4.9
	unit	8.3	each	5.4	total price	15.5	only	4.0
	you both	5.5	retailer	5.3	loss	15.5	hampers	3.6
	ever	3.8	one	3.9	2	15.5		
					amount	7.9		
					going	7.9		
					better	3.9		

Notes: Words ranked by frequency within chat corpus c. Lower numbers for rank r_c indicates a more common word. Displayed are words whose absolute rank satisfies $r_c \leq 50$ and whose rank differential relative to comparison corpus c' satisfies $(r_c - r_{c'})/r_c \geq 3.5$. Conjunctions, prepositions, and articles omitted. Comparisons hold constant treatment (*Two Chat* or *Three Chat*), source of message (upstream or downstream), and symmetry of offered quantities, only varying total quantity involved in offer ($X = 2$ versus $X = 4$). Words are translations from the original German.

Table 2.9: Keywords Mined from Chat Leading to Asymmetric Offers

	Low-quantity offer with $x_i = 1$ (corpus L)				High-quantity offer with $x_i = 2$ (corpus H)			
	Upstream messages		Downstream messages		Upstream messages		Downstream messages	
Treatment	Word	Rank diff. relative to H	Word	Rank diff. relative to H	Word	Rank diff. relative to L	Word	Rank diff. relative to L
Two Chat	how	48.0	how	15.8			45	9.7
			only	9.3			makes	9.7
			going	9.3			units	8.6
			are	7.1			nothing	5.8
			each	7.1			60	5.8
			us	7.1			2	5.5
			10	6.1			35	4.2
			1	5.3			me	4.2
			otherwise	5.0				
			would	5.0				

Notes: Words ranked by frequency within chat corpus c. Lower numbers for rank r_c indicates a more common word. Displayed are words whose absolute rank satisfies $r_c \leq 50$ and whose rank differential relative to comparison corpus c' satisfies $(r_c - r_{c'})/r_c \geq 3.5$. Conjunctions, prepositions, and articles omitted. Comparisons hold constant treatment (*Two Chat*), source of message (upstream or downstream), and asymmetry of offered quantities, only varying quantity involved in individual offer ($x_i = 1$ versus $x_i = 2$). Words are translations from the original German.

2.7 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.

The last section delved into content analysis using a variety of analytical approaches. Our analysis of message counts found that more messages correlated with successful monopolization. There was also a positive correlation between messages sent by a subject and that individual's bargaining share. The exercise employing third-party coders confirmed that chat functioned like a bargaining process, with discussions successfully converging to a contract that is the one that ends up being offered. Departures from these expectations were significantly less likely to be accepted. The keyword-mining exercise found that when the upstream firm was successful at committing to the monopoly outcome, his or her messages tend to mention deals given to all both downstream firms and market prices. Commitment sometimes breaks down when a subject tries to strike an exclusive deal to sell the entire industry quantity, inevitably leading to oversupply as the exclusion proves to be unenforceable.

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.

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.

Bigoni, Maria, Sven-Olof Fridolfsson, Chloé Le Coq, and Giancarlo Spagnolo. (2012) "Fines, leniency, and rewards in antitrust," *Rand Journal of Economics*, 43(2): 368–390.

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.

Collard-Wexler, Allan, Gautam Gowrisankaran, and Robin S. Lee. (2016) " 'Nash-in-Nash' Bargaining: A Microfoundation for Applied Work," Harvard University working paper.

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.

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.

Grandjean, Gilles, Marco Mantovani, Ana Mauleon, and Vicent Vannetelbosch. (2014) "Whom Are You Talking With? An Experiment on Credibility and Communication Structure," CORE discussion paper no. 2014/42.

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.

Hinloopen, Jeroen, and Adriaan R. Soetevent. (2008) "Laboratory evidence on the effectiveness of corporate leniency programs," *Rand Journal of Economics*, 39(2): 607–616.

Horn, Henrick and Asher Wolinsky. (1988) "Bilateral Monopolies and Incentives for Merger," *Rand Journal of Economics* 19: 408–419.

Houser, Daniel and Erte Xiao. (2011) "Classification of Natural Language Messages Using A Coordination Game," *Experimental Economics*, 14: 1–14.

Huerta, Juan M. "Relative Rank Statistics for Dialog Analysis," *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, 965–972.

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.

Lee, Hau L. and Seungjin Whang. (2000) "Information Sharing in a Supply Chain," *International Journal of Technology Management* 20: 373–387.

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, Partick 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.

Vanpoucke, Evelyne, Kenneth K. Boyer, and Ann Vereecke. (2009) “Supply Chain Information Flow Strategies: An Empirical Taxonomy,” *International Journal of Operations and Production Management* 29: 1213–1241.

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.

Appendix A: Additional Bargaining Results

This appendix presents several theoretical results on bargaining referenced in the text.

A1. Verifying Uniqueness Conditions

Theorem 4.3 of Collard-Wexler, Gowrisankaran, and Lee (2016) provides sufficient conditions for the Nash-in-Nash solution to characterize any equilibrium of a generalized Rubinstein (1982) process in which upstream and downstream firms alternate offers. The following assumptions together make up the sufficient conditions: Gains from Trade (A.GFT), Strong Conditional Decreasing Marginal Contribution (A.SCDMC), and Limited Negative Externalities (A.LNEXT).

Before proceeding to verify that these assumptions hold in our setting, we need to introduce some of the authors' notation, adapted to our setting. Let j indicate the consummation of a successful bargain between U and D_j, $j \in \{1, 2\}$, resulting in the trade of one unit. If all efficient trades are made, the outcome is denoted $\mathcal{G} = \{1, 2\}$. An arbitrary outcome is denoted by the set $\mathcal{A} \subseteq \mathcal{G}$. U's gross surplus (payoff not including transfers T_j) in outcome $\mathcal{A}$ is denoted $\pi_U(\mathcal{A})$ and D_j's is $\pi_{D_j}(\mathcal{A})$. Let $\Delta\pi_U(\mathcal{A}, \mathcal{B}) = \pi_U(\mathcal{A}) - \pi_U(\mathcal{A} \setminus \mathcal{B})$, $\mathcal{B} \subseteq \mathcal{A} \subseteq \mathcal{G}$, denote the marginal contribution of agreements $\mathcal{B}$ to the gross surplus U earns from $\mathcal{A}$. Define $\Delta\pi_{D_j}(\mathcal{A}, \mathcal{B})$ analogously.

A preliminary result will be helpful in verifying the assumptions. In our setting, U's only source of surplus is T_i. It otherwise earns no gross surplus, and its production is costless. Thus, for all $\mathcal{B} \subseteq \mathcal{A} \subseteq \mathcal{G}$,

$$\pi_U(\mathcal{A}) = \Delta\pi_U(\mathcal{A}, \mathcal{B}) = 0. \tag{A1}$$

We can now proceed to verify the three assumptions. Translated into our setting, the first assumption, A.GFT, holds for a representative downstream firm, say D_1 if

$$\Delta\pi_U(\mathcal{G}, \{1\}) + \Delta\pi_{D_1}(\mathcal{G}, \{1\}) > 0. \tag{A2}$$

Equation (A1) implies $\Delta\pi_U(\mathcal{G}, \{1\}) = 0$. Further, $\Delta\pi_{D_1}(\mathcal{G}, \{1\}) = \pi_{D_1}(\mathcal{G}) - \pi_{D_1}(\mathcal{G} \setminus \{1\}) = \pi_{D_1}(\mathcal{G})$. But $\pi_{D_1}(\mathcal{G}) = 50$ given our experimental parameters. Together, these calculations verify equation (A2).

That leaves two assumptions. Both of these are divided into two parts, one for upstream and one for downstream firms. By equation (A1), A.SCDMC and A.LNEXT are trivially satisfied for the upstream firm because they reduce to the inequality $0 \geq 0$. We need only verify A.SCDMC and A.LNEXT hold for downstream firms.

Translated into our setting, A.SCDMC holds for a representative downstream firm, say D_1, if

$$\pi_{D_1}(\mathcal{A} \cup \mathcal{B} \cup \{1\}) - \pi_{D_1}(\mathcal{A}' \cup \mathcal{B}) \geq \Delta\pi_{D_1}(\mathcal{G}, \{1\}) \tag{A3}$$

for all $\mathcal{B} \subseteq \mathcal{G}_{-U}$ and $\mathcal{A}, \mathcal{A}' \subseteq \mathcal{G}_U \setminus \{1\}$, where $\mathcal{G}_{-U}$ is the set of agreements that can be made with upstream firms besides U and $\mathcal{G}_U$ is the set of agreements that can be made with U. Considering the first term on the left-hand side of (A3), for all $\mathcal{B} \subseteq \mathcal{G}_{-U}$ and $\mathcal{A}, \mathcal{A}' \subseteq \mathcal{G}_U \setminus \{1\}$, we have

$$\begin{aligned} \pi_{D_1}(\mathcal{A} \cup \mathcal{B} \cup \{1\}) &= \pi_{D_1}(\mathcal{A} \cup \{1\}) && \text{(A4)} \\ &\geq \pi_{D_1}(\{2\} \cup \{1\}) && \text{(A5)} \\ &= \pi_{D_1}(\mathcal{G}). && \text{(A6)} \end{aligned}$$

To see (A4), noting that U is the only upstream firm, we have $\mathcal{G}_{-U} = \varnothing$, implying that $\mathcal{B} = \varnothing$. To see (A5), noting again that U is the only upstream firm, $\mathcal{G}_U = \mathcal{G}$, implying $\mathcal{G}_U \setminus \{1\} = \{2\}$. Hence $\mathcal{A}$ must be either $\varnothing$ or $\{2\}$. D_1's lowest payoff is generated by $\mathcal{A} = \{2\}$. Considering the second term on the left-hand side of (A3), for all $\mathcal{B} \subseteq \mathcal{G}_{-U}$ and $\mathcal{A}, \mathcal{A}' \subseteq \mathcal{G}_U \setminus \{1\}$, we have $\pi_{D_1}(\mathcal{A}' \cup \mathcal{B}) = \pi_{D_1}(\mathcal{A}') = 0$, where the first equality follows from $\mathcal{B} = \varnothing$ and the second from the fact that $1 \notin \mathcal{A}' \in \mathcal{G}_U \setminus \{1\}$. Hence the left-hand side of (A3) is at least $\pi_{D_1}(\mathcal{G})$. The right-hand side is $\Delta\pi_{D_1}(\mathcal{G}, \{1\}) = \pi_{D_1}(\mathcal{G}) - \pi_{D_1}(\mathcal{G} \setminus \{1\}) = \pi_{D_1}(\mathcal{G})$ because $\pi_{D_1}(\mathcal{G} \setminus \{1\}) = 0$. This completes the proof that A.SCDMC holds.

It remains to verify A.LNEXT. Translated into our setting, A.LNEXT holds if, for all nonempty $\mathcal{C} \subseteq \mathcal{G}$, there exists $j \in \mathcal{C}$ such that

$$\Delta\pi_{D_j}(\mathcal{G}, \mathcal{C}) \geq \sum_{j \in \mathcal{C}_{D_j}} \Delta\pi_{D_j}(\mathcal{G}, \{j\}), \tag{A7}$$

where $\mathcal{C}_{D_j} = \mathcal{C} \cap \{j\}$. Consider each of the three possibilities for $\mathcal{C}$ in turn, namely, $\{1\}$, $\{2\}$, and $\mathcal{G}$. First suppose $\mathcal{C} = \{1\}$. Taking $j = 1$, the left-hand side of (A7) becomes $\Delta\pi_{D_1}(\mathcal{G}, \{1\}) = \pi_{D_1}(\mathcal{G}) - \pi_{D_1}(\mathcal{G} \setminus \{1\}) = \pi_{D_1}(\mathcal{G})$. The right-hand side of (A7) can be simplified by noting $\mathcal{C}_{D_1} = \mathcal{C} \cap \{1\} = \{1\} \cap \{1\} = \{1\}$. Hence the summation reduces to the single term $\Delta\pi_{D_1}(\mathcal{G}, \{1\}) = \pi_{D_1}(\mathcal{G}) - \pi_{D_1}(\mathcal{G} \setminus \{1\}) = \pi_{D_1}(\mathcal{G})$. This proves that (A7) holds for $\mathcal{C} = \{1\}$. The proof that (A7) holds for $\mathcal{C} = \{2\}$ is identical. That leaves $\mathcal{C} = \mathcal{G}$. Taking $j = 1$, the left-hand side of (A7) then is $\Delta\pi_{D_1}(\mathcal{G}, \mathcal{G}) = \pi_{D_1}(\mathcal{G})$. The right-hand side can again be shown to involve a single term in the summation because $\mathcal{C}_{D_1} = \mathcal{G} \cap \{1\} = \{1\}$. This sum can again be shown to reduce to $\pi_{D_1}(\mathcal{G})$, proving the left- and right-hand sides of (A7) are equal in this case. This completes the proof that A.LNEXT holds.

A2. Generalizing Shapley Value

In this section of the appendix, we present 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

$$\alpha = \frac{\lambda_D}{\lambda_D + \lambda_U}. \tag{A8}$$

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

$$\begin{aligned} \Pi^m \Pr\left(A_U > \min_{i\in\{1,\ldots,d\}}\{A_i\}\right) &= \Pi^m \left[1 - \Pr\left(A_U < \min_{i\in\{1,\ldots,d\}}\{A_i\}\right)\right] && \text{(A9)} \\ &= \Pi^m \left(\frac{d\lambda_D}{\lambda_U + d\lambda_D}\right) && \text{(A10)} \\ &= \Pi^m \left(\frac{\alpha d}{1-\alpha+\alpha d}\right), && \text{(A11)} \end{aligned}$$

where (A10) follows from standard results for exponential distributions and (A11) from (A8).

The tariff implementing the equilibrium surplus share in (A11) is

$$T_i^* = \frac{\Pi^m}{2}\left(\frac{\alpha d}{1-\alpha+\alpha d}\right). \tag{A12}$$

This equation provides the fitted tariff values for the rows in Table 2.4 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 (A11) and (A12) are increasing in d. This provides a contrasting comparative-static result to that derived in the text for the Nash-in-Nash solution with general asymmetric bargaining weights.

Appendix B: Instructions

Basic Instructions

Welcome to our experiment! In the next hour you will make decisions at a computer. One thing is important right from the start: please be quiet during the entire experiment and please do not talk to your neighbors. The experiment runs over 15 rounds.

In the experiment we will use a fictitious currency called ECU. In the beginning you will get a starting capital in ECU. During the experiment you can earn some real money, but losses are also possible.

After the last round, you will be paid 1 euro for every 40 ECU you earned during the experiment. Concerning the payment, there is strict anonymity with respect to the other participants as well as with respect to us. We will record no data in connection with your name.

What is the experiment about? The experiment is about decision making in a market with one producer and two retailers. Some of you will make decisions for a producer, others for a retailer. You will be a producer or a retailer for all 15 rounds of the experiment. Consumers in the market are simulated by the computer program. You will be told whether you are a producer or a retailer at the beginning of the experiment. Currently, you are all reading the same instructions.

Note that in every round the producer-retailer groups change.

The basic market structure is the following, the producer produces a product which he sells to the retailers. The retailers resell the product to the final consumers in their stores.

What are you supposed to do as a producer or retailer? A producer has to decide how many units of the product he wants to sell at which price to the two retailers. This decision has the form of an offer to the retailers: each retailer is offered a specified quantity (integer) of the product at a specified total price. The producer may also decide to offer a quantity of zero of the product to one or both retailers. The starting capital for the producer is 60 ECU.

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

If a retailer receives an offer, he has to decide either to accept the offer or to reject it. If he accepts the offer, he receives the number of units of the product specified in the offer and has to pay the total price. If he rejects the offer, he does not receive the product and does pay anything to the producer. The starting capital for the retailer is 200 ECU.

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

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

Retailers' revenues are the number of units supplied (that is, bought from the producer) multiplied by the market price. In the example, retailer 1 has revenues of $2 \cdot 30$ ECU $= 60$ ECU, while retailer 1 has revenues of $1 \cdot 30$ ECU $= 30$ ECU.

Retailers' stores are run without cost. The profit of a retailer is thus the revenues minus the payment to the producer.

Suppose that, in the example, retailer 2 agreed to pay 35 ECU for the 1 unit he received. Then he would actually make a loss of 5 ECU. If he agreed to pay only 5 ECU, a profit of 25 ECU would result.

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

Which information do you get? Each retailer knows only his own offer but not the offer of the other retailer. Each retailer is told his own profit at the end of each round. The producer is told whether or not the retailers accepted the offers at the end of each round. The producer is informed about his own profit and the profit of the two retailers at the end of each round.

Additional Instructions for *Two Chat*

Before the producer and the retailers make their decisions, you have the possibility to communicate. The producer may write messages via a chat window to retailer 1 and to retailer 2. Both retailers may communicate with the producer , but they cannot write to each other; the retailers cannot observe what the producer and the other retailer talked about.

You are able to communicate at the beginning of every round and the time is restricted to 90 seconds in the first 5 rounds and 60 seconds after that.

Additional Instructions for *Three Chat*

Before the producer and the retailers make their decisions, you have the possibility to communicate. All of the three market participants may write messages in a chat window. Both the producer and the retailers may send messages and the other two market participants can read these messages. If, for example, retailer 1 sends a message, retailer 2 as well as the producer can read and answer it. You cannot write only to one of the other two market participants.

You are able to communicate at the beginning of every round for 60 seconds.

Additional Instructions for *Choose Chat*

Before the producer and the retailers make their decisions, you have the possibility to communicate. On the one hand communication between the producer and one of the retailers is possible whereas the other retailer cannot observe the conversation. That means the manufacturer may write to retailer 1 as well as retailer 2 via a chat window on the left and on the right hand side on the screen, respectively. In these chat windows both retailers may communicate with the producer, but they cannot write to each other; the retailers cannot observe what the producer and the other retailer talked about.

On the other hand it is possible to communicate with all of the three market participants. Both the producer and the retailers may send messages and the other two market participants can read these messages. If, for example, retailer 1 sends a message, retailer 2 as well as the producer can read and answer it. In this chat window you cannot communicate with only one of the other two market participants.

The producer may communicate via three different chat windows. The conversation among all three market participants takes place in the chat window in the middle of the screen. Via the chat window on the left- and right-hand side the manufacturer can communicate separately with retailer 1 and retailer 2, respectively. Retailers have two different chat windows. They can communicate either in a threesome in the middle chat window or separately with the manufacturer on the left and right hand side of the screen. (You can see the three different variants below.)

You are able to communicate at the beginning of every round for 90 seconds.

Chapter 3

Search Costs in Concentrated Markets – An Experimental Analysis

Co-authored with Torben Stühmeier and Tobias Wenzel

3.1 Introduction

In many markets consumers are only imperfectly informed about prices and face search cost to find the best deal. This is, in particular, true for many utilities such as energy or telecommunications. Another example are gasoline markets where, due to frequent price changes, consumers rarely know which firm offers the lowest price. A key characteristic of the above mentioned examples is that they are typically concentrated markets with only a few major firms dominating the market.

Policy makers are concerned about potential high prices in such markets, and there has been a recent interest in interventions to increase price transparency. For instance, the use of price comparison websites for utilities has been promoted in many countries. In gasoline markets, a number of countries (such as Australia, Canada, Germany or Austria) have backed interventions to reduce consumers' search cost. It is, however, a rather open question how effective such policies are.

In this paper, we present results of a laboratory experiment which helps answer this question. We consider a setting based on the search model by Stahl (1989) with two types of buyers. A share of buyers, called shoppers, is always informed about all sellers' prices while the remaining share, called non-shoppers, can only become informed at a cost. Search is sequential in the sense that the non-shoppers observe the price of one randomly drawn seller and then decide whether to invest in search. In a static framework, Stahl (1989) shows that sellers mix over prices and the price distribution is such that in equilibrium non-shoppers never search. Search costs, however, matter in the sense that the price distribution shifts downward as the search costs become smaller leading to lower average prices.

In our setting, we focus on concentrated markets with two sellers. We extend Stahl's model and consider a dynamic variant where the same two firms compete repeatedly over time. With an infinite time horizon cooperative equilibria, e.g. both sellers set the monopoly price, exist, if firms are sufficiently patient. Lowering search cost makes such cooperative equilibria more likely as punishment profits (that is, equilibrium profits in the static equilibrium) decrease when search cost become smaller. Thus, while search cost for buyers may be beneficial for sellers if they compete, the opposite effect may occur if firms can coordinate pricing decisions.

We consider experimental markets with two sellers and two buyers, i.e. both were represented by participants in the lab. We vary two treatment variables, the level of search cost is either high or low and on the other hand pre-play communication between sellers is either possible or not. Sellers set their price, upon observing the price, buyers choose whether they invest in search and from which seller to purchase the product. Buyers were asked for their entire search strategy, i.e. they decided for every possible price whether they want to invest in costly search. The search strategy was applied according to the observed price. Sellers were allowed to communicate via unrestricted written messages, the conversation could not be followed by the buyers. In order to reflect a dynamic setup, market constellations were held constant throughout the experiment and we implemented a random stopping rule.

Our experimental findings suggest that interventions which aim at decreasing consumer search cost are not an effective tool to lower prices in concentrated markets. While buyers search more frequently with lower search cost, average prices do not vary with the level of the search cost. In contrast to predictions, we find that prices are more dispersed when search cost is low. Moreover, because lower search costs go along with more frequent search, the consumer surplus is unaffected by changes in cost. Seller communication, on average, increases prices and reduces search incentives, however, lower search costs have no effect on prices and buyer surplus.

Our design allows us to shed light on buyers' search strategies. We find that, for any given price, a buyer is more likely to invest in search when search cost is low. More importantly, however, we also find that buyers are much less likely to search if sellers can communicate. This is consistent with the view that price coordination is higher with communication which reduces the incentives to search. Comparing the observed search strategy with the optimal one, we find that without communication, there is too much search at low prices, but too little search at high prices. This effect is more pronounced at lower search costs. In contrast, with communication there is always too much search.

A number of experimental studies investigate search markets. Closest to our paper is the study by Cason and Friedman (2003). Based on Burdett and Judd (1983) they investigate the impact of different shopper (each is matched with two sellers at the same time) and non-shopper (each is matched with one seller) ratios. In addition, they also vary whether buyers are computerized or active participants in the lab. Our contribution differs in several aspects from theirs. First, our experimental design allows us to observe the entire search strategy, i.e. we are able to identify why Cason and Friedman observe differences between computerized and participating buyers. Second, we focus on dynamic competition and potential collusion (with possible seller communication) which is not studied in their experiment.

Davis and Holt (1996) conducted an experiment on search markets which aims to test the validity of the Diamond paradox (Diamond, 1971). When all buyers have positive search cost, the prediction is monopoly pricing independent of the level of search cost. In contrast, Davis and Holt find a positive relationship between prices and the level of search cost. Abrams et al. (2000), also based on Burdett and Judd (1983), compare treatments where theory predicts either monopoly or marginal cost pricing. In contrast to the predictions they find that prices are biased towards the middle of these extremes. Cason and Datta (2006) report an experiment where sellers set prices and decide whether to advertise while buyer search is costly. They find that an increase in search cost raises equilibrium prices and increases sellers' advertising intensity.

Few studies consider collusion with search. To our knowledge, Orzen (2008) is one exception.[1] In this paper, buyers are simulated by a computer and consumer information, as measured by the share of informed buyers, is varied exogenously. In contrast, in our experiment buyers are participants in the lab and buyer information is an endogenous decision. Normann and Wenzel (2013) and Crosetto and Gaudeul (2016) also consider collusion in markets where consumers are not perfectly informed, but in their experiments it is sellers rather than buyers who can influence consumer information via obfuscation and confusion strategies.

The remainder of the paper is organized as follows. Section 3.2 outlines the theoretical background. Section 3.3 describes our experimental setup. In Section 3.4 we present the experimental findings. Finally, Section 3.5 concludes.

3.2 A sequential search model

3.2.1 Static competition

This section presents a sequential search model along the lines of Stahl (1989). We consider a duopoly version where a unit mass of consumers each demands one unit of a homogenous product up to a reservation value of r.

There are two groups of consumers, shoppers and non-shoppers. The mass of shoppers (non-shoppers) is μ $(1-\mu)$. Sampling of the first product is costless for both groups, and each consumer receives price information by one randomly chosen firm (with equal probability). For sampling the second product and to learn its price, consumers have to incur a search cost. Shoppers have zero search cost and, hence, always sample both products while non-shoppers have to incur a positive search cost of $c > 0$. Non-shoppers will only invest into search if the benefit of searching (i.e., the possibility of finding a lower price) outweighs the search cost.

The equilibrium of the static game, where firms compete only once, is characterized as follows:[2]

[1] The paper uses a similar experimental setting as in Morgan et al. (2006a) which focuses on static competition.

[2] The derivation of the equilibrium follows the one in Stahl (1989) and is therefore omitted. The only difference is that Stahl assumes downward sloping demand while we consider unit demand.

Lemma 1. i) Firms price according to

$$F(p) = 1 - \frac{1-\mu}{2\mu}\left[\frac{\bar{p}}{p} - 1\right] \tag{B1}$$

on $p \in [\underline{p}, \bar{p}]$, where

$$\bar{p} = \min\left\{r, \frac{c}{1 - \frac{1-\mu}{2\mu}\ln\left(\frac{1+\mu}{1-\mu}\right)}\right\} \tag{B2}$$

$$\underline{p} = \frac{1-\mu}{1+\mu}\bar{p}. \tag{B3}$$

Firms earn profits of $\Pi_c = \frac{1-\mu}{2}\bar{p}$.

ii) Non-shoppers employ a cut-off search strategy with cut-off price $\bar{p}$, but do not search in equilibrium.

The equilibrium structure is as follows. Equilibrium pricing is in mixed strategies due to the trade-off of charging a high price towards non-shoppers and competing for shoppers by offering a low price (Varian, 1980). In equilibrium, shoppers (with zero search cost) sample both firms, hence, buy from the firm that offers the lowest price. In contrast, non-shoppers (those consumers with a positive search cost c) employ a cut-off search strategy (with cut-off price $\bar{p}$), but in equilibrium they do not search and buy randomly from one of the firms. As a consequence, non-shoppers pay on average a higher price than shoppers.

It should be noted that the interval over which firms randomize prices depends on the level of the search cost. The upper bound is set such that non-shoppers only sample one firm (Stahl, 1989). The lower the search cost, the lower is the upper bound of the price distribution. If search cost is low, competition is high and the lower bound approaches marginal cost (here zero) for zero search cost. If, however, search cost is relatively high ($c \geq r\left[1 - \frac{1-\mu}{2\mu}\ln(\frac{1+\mu}{1-\mu})\right]$) the upper bound of the price distribution would be equal to reservation value of r and, hence, pricing would not depend on the search cost. In the following we will focus on cases where search cost matter, that is, $c < r\left[1 - \frac{1-\mu}{2\mu}\ln(\frac{1+\mu}{1-\mu})\right]$.

We are interested in studying the effects of a reduction in the search cost. A reduction of the search cost c shifts the price distribution $[\underline{p}, \bar{p}]$ downwards. Note also from Proposition 1 that equilibrium profits are proportional to the upper bound of the price distribution. As a reduction of the search cost shifts the upper bound downwards, profits are affected negatively. As demand is inelastic, consumer surplus is inversely related to industry profits so that with lower search cost consumer surplus increases.

This can also be seen by inspecting average market prices. The expected market price can be calculated as follows:

$$E(p) = \int_{\underline{p}}^{\bar{p}} p f(p) dp, \tag{B4}$$

where $f(p)$ is the density function associated with the equilibrium distribution function $F(p)$. Simplification then yields:

$$E(p) = \frac{c \ln(\frac{1+\mu}{1-\mu})}{\frac{2\mu}{1-\mu} - \ln(\frac{1+\mu}{1-\mu})}. \tag{B5}$$

It is straightforward to see that $\frac{\partial E(p)}{\partial c} > 0$. That is, a reduction in search cost c leads to a reduction of the average price in the market. We summarize the effects in the static game:

Proposition 1. In the static game, reducing search cost leads to lower expected market prices, reduced firm profits and higher consumer surplus.

3.2.2 Dynamic competition

We now study an infinitely repeated version of the static game presented in the preceding section. We focus on the case where firms may collude on the highest possible profits and analyze how a reduction in search cost influences firms' ability to collude.

The collusive price that maximizes joint industry profits is both firms choosing the price equal to consumers' reservation value, r. If both firms charge identical prices, non-shoppers have no incentive to incur costly search effort, no matter how small the search cost might be. Hence, non-shoppers buy randomly from the firm whose price they can observe. In contrast, shoppers who have no search cost still observe both prices and are indifferent between both firms' offers. Profits under this collusive strategy amount to:

$$\Pi_k = \frac{r}{2}. \tag{B6}$$

A firm considering to deviate can only reach shoppers (zero search cost) as only this group compares the prices of the two firms. The optimal deviation is then to slightly undercut the collusive price as to capture all shoppers. Profits of the deviating firm are then

$$\Pi_d = \left[\frac{(1-\mu)}{2} + \mu\right] r = \frac{r(1+\mu)}{2}. \tag{B7}$$

We consider trigger strategies. With this strategy, if all firms behave according to the collusive strategy, firms will continue to set the collusive price. However, as soon as any deviation occurs, collusion stops and firms revert to the static Nash equilibrium, earning profits of Π_c for the remaining periods. The critical discount factor, for collusion to be a subgame perfect equilibrium of the infinitely repeated game, is

$$\overline{\delta} = \frac{\Pi_d - \Pi_k}{\Pi_d - \Pi_c}, \tag{B8}$$

which, applied to the above framework, can be expressed as

$$\overline{\delta} = \frac{r\mu\left(2\mu + (1-\mu)\ln\left(\frac{1-\mu}{1+\mu}\right)\right)}{r(1-\mu^2)\ln\left(\frac{1-\mu}{1+\mu}\right) + 2\mu\left(2r\mu + (1-\mu)(r-c)\right)}. \tag{B9}$$

In the following, we analyze the impact of reduced search cost. Note that only Π_c depends on search cost, and profits under collusion and deviation are independent of the search cost. As $\frac{\partial \Pi_c}{\partial c} > 0$, it follows immediately that $\frac{\partial \overline{\delta}}{\partial c} > 0$. Hence, a reduction of search cost facilitates collusion as the punishments profits are lower if a deviation occurs. It is interesting to note that, on the one hand, a reduction in search cost tend to reduce prices in the static game, but on the other hand, tend to make collusive outcomes more likely in the dynamic game.

The following proposition summarizes the impact of search cost on the incentives to collude:

Proposition 2. A reduction of search cost stabilizes collusion.

3.3 Experimental design and procedures

We analyze the effect of a reduction of search cost in concentrated markets with repeated interaction. Therefore, we vary the level of search cost and whether or not seller communication is allowed.

Treatment	Communication	Search cost	participants	# groups
LowFix	no	$c = 0.5$	40	10
HighFix	no	$c = 2$	36	9
LowCom	yes	$c = 0.5$	20	5
HighCom	yes	$c = 2$	24	6

Table 3.1: Treatments

Experimental markets consist of 4 participants, 2 sellers and 2 buyers. Sellers produce at zero costs and set an integer price in $[1, 10]$. Buyers have an inelastic demand up to a reservation value of $r = 15$. Because there is no outside option, buyers always purchase one unit of the good.[3] There are two types of buyers in our experiment, shoppers and non-shoppers. With a probability of 25% a buyer is a shopper in a specific period, $\mu = 0.25$, which means the buyer observes prices of both sellers at zero costs. Buyers are non-shoppers with a probability of 75%, i.e. search is costly. Buyers search sequentially, i.e. upon observing the price of a randomly chosen seller they decide whether they want to invest search cost c in order to obtain price information of both sellers.

Each period is divided in three stages. In the first stage, sellers make their price decision and buyers determine their search strategy. We gave buyers a list of all possible prices and they had to decide for each of the prices whether they would invest in search if they observed that price.[4] In the second stage, buyers get informed if they are shoppers or non-shoppers and make their purchase decision. According to their search strategy from stage one, non-shoppers get either price information of one or both sellers. Shoppers observe prices of both sellers which is costless for them and independent of their search strategy. Buyers purchase one unit of the good from either of the sellers. Finally, sellers and buyers receive feedback on their earnings in the third stage. Sellers either sell 0, 1 or 2 units at the chosen price. Buyers obtain their valuation minus the price of the chosen seller. Additionally, non-shoppers have to pay the search cost c if they decided to invest in search.

At the beginning of every session, participants were randomly assigned to be a seller and buyer, respectively. The role of participants was kept constant during the entire experiment. Buyers and sellers remained in the same market throughout

[3]We chose a valuation above the maximum price to ensure positive payoffs for buyers even if the incurred search cost is high.

[4]We implemented the strategy method (Selten, 1967) to obtain the complete search strategy of buyers.

	low search cost	high search cost
Average price in static game	2.0	6.81
Price interval in static game	[2.0, 2.0]	[6.0, 8.0]
Critical discount factor	0.23	0.39

Table 3.2: Predictions

the experiment, i.e. we implemented a fixed matching protocol.[5] The experiment consisted of at least 20 periods and a random stopping rule of 50% was implemented afterwards. A priori, we randomly determined the number of periods to be 21 which was the same for all sessions.

In our four treatments we varied the level of search cost (high or low) and whether seller communication is allowed. This results in a 2 by 2 design summarized in Table 3.1. In the low cost treatments we set search cost $c = 0.5$ and we determined high search cost to equal $c = 2$. In the communication treatments, sellers were allowed to engage in private communication with their competitor before they make their price decisions. Sellers' communication was not structured and could not be observed by the buyers. The pre-play communication phase lasted for 60 seconds in the first five periods, 45 seconds afterwards and sellers could not exit this stage before time expired.

All sessions were run in the DICElab at the University of Düsseldorf and took about one hour. Participants were invited using ORSEE (Greiner, 2015) and the experiment was implemented using the software z-Tree (Fischbacher, 2007). The appendix contains an English translation of the instructions. Subjects received a show-up fee of 4 EUR and could earn additional amounts during the experiment. On average, participants received 14.50 EUR. In total, 120 subjects participated in our experiment. No subject participated in more than one session and none of the subjects had participated in a similar experiment before.

Table 3.2 provides the predictions given our parameter choices. The predictions also take into account that in the experiment the price is a discrete variable, ranging from 1 to 10.

Treatment	*LowFix*	*HighFix*	*LowCom*	*HighCom*
Price	5.37	5.25	8.96	8.70
(Std Dev)	(1.94)	(1.37)	(1.31)	(1.48)
Transaction price	4.60	4.90	8.84	8.59
(Std Dev)	(1.43)	(1.19)	(1.33)	(1.57)
Collusive intensity	1.69	-0.23	3.48	0.28
Search frequency	62.81%	16.67%	43.75%	23.96%
Search effort	0.32	0.31	0.22	0.45
CS	10.15	9.86	6.00	6.03
(Std Dev)	(1.49)	(1.36)	(1.37)	(1.65)

Table 3.3: Main results

3.4 Experimental results

This section presents the experimental results. Table 3.3 provides an overview of the main results. To account for learning effects, we exclude the first 5 periods. Throughout the paper we employ non-parametric tests where the number of independent observations corresponds to the number of matching groups.[6]

3.4.1 Price setting

We start by analyzing sellers' price choices. In Table 3.3 we report average posted prices and transaction prices.

Average posted and transaction prices

The first main finding is that varying the level of search cost does not lead to significant changes in average prices (independent of whether seller communication is allowed). As can be seen in Table 3.3, mean posted and mean transaction prices are slightly larger with lower search cost, but these differences are not significant (p-value of 0.71 and p-value of 0.46, for posted prices and p-value of 0.39 and p-value of 0.41 for transaction prices, Mann-Whitney U test).

[5]In addition, we ran 3 sessions implementing a random matching procedure where markets were randomly re-matched in each period. The results are by and large the same as obtained results with fixed matched markets.

[6]Tables 3.5 and 3.6 in the appendix provide an overview over all statistical tests.

Prices, however, vary with communication for both levels of search cost. Average prices are lower without communication. For instance, with low search cost the price rises from 5.37 to 8.96 when seller communication is allowed (p-value of < 0.01, Mann-Whitney U test). We observe the same pattern with high search cost and also when we consider transaction prices. The finding that seller communication leads to higher prices and better seller coordination is in line with existing studies (e.g. Fonseca and Normann, 2012, Cooper and Kühn, 2012, Harrington et al., 2016). We find this effect is also present in a search market environment where buyers are active participants in the experiment.[7]

Result 1. i) Reducing search cost has no effect on average posted and transaction prices. ii) Seller communication leads to higher posted and transaction prices in search markets.

Comparing the observed average price with the predictions of the static game, we calculate the intensity of collusion as $(p_{obs} - p_{static})/p_{static}$, which is reported in Table 3.3. Without communication we find that for high search cost observed prices are lower than theory predicts, yielding an intensity of collusion of -0.22. With low search cost, observed prices exceed the prediction raising the intensity of collusion to 1.69. We observe a similar result when sellers can communicate. With either level of search cost, the average price exceeds the competitive prediction, but again the collusive intensity is higher with lower search cost. Hence, reducing search cost seems to raise the incentives to collude when compared to the theoretical benchmark in the one-shot game.

Result 2. The collusive intensity is higher with lower search cost.

Price dispersion

So far we only looked at average prices. However, to understand buyers' incentives to invest in search, it is important to examine to which degree prices are dispersed.

Figure 3.1 shows the price distributions across treatments. The left panel displays price frequencies for treatments without communication and the right panel shows the case with seller communication. In line with the previous result, we observe that with seller communication the entire price distribution is shifted to the right. The modal price chosen by sellers is the monopoly price of 10 with communication and a price of 5 without the option to communicate (for both levels of search cost).

[7] We note that the aforementioned findings are stable over time. See Figure 3.5 in the appendix where we show posted prices over time for each treatment.

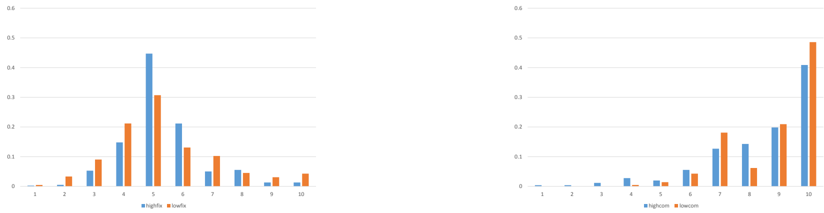

(a) No communication (b) Communication

Figure 3.1: Price distributions

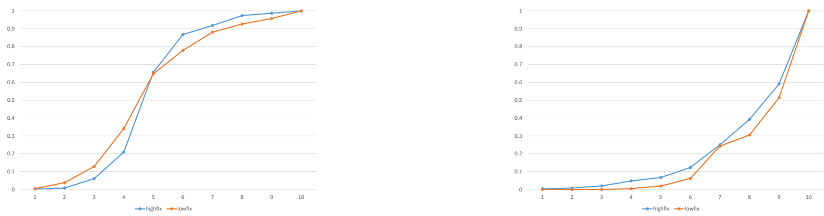

(a) No communication (b) Communication

Figure 3.2: Cumulative price distributions

Figure 3.2 provides an alternative representation by displaying the cumulative price distributions. Without communication we see that the cdfs with high and low search costs intersect implying more price dispersion with a lower search cost level. When sellers can communicate, the cdfs are close and do not intersect.

We conduct Kolmogorov-Smirnov tests in order to formally compare the price distributions between treatments.[8] The tests confirm communication to have a highly significant impact on the distribution of prices, irrespectively of the level of search cost (p-values of < 0.01, Kolmogorov-Smirnov test). In addition, the level of search cost affects the dispersion of prices if sellers cannot communicate. Without communication, lower search cost increase the dispersion of prices (p-values of < 0.01, Kolmogorov-Smirnov test). When communication is allowed, differences according to the level of search cost disappear (p-value of 0.30, Kolmogorov-Smirnov test).

[8]For a summary of the results see Table 3.6 presented in the Appendix.

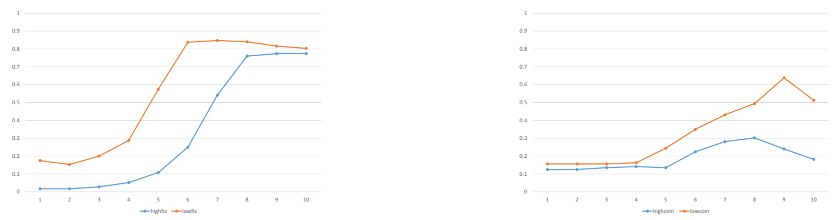

(a) No communication (b) Communication

Figure 3.3: Search strategies

Result 3. i) Without communication, prices are more dispersed with a lower level of search cost. ii) Seller communication decreases price dispersion.

3.4.2 Search strategies

With our experimental design we are able to observe the entire search strategy of buyers, i.e., we can observe for every possible price p whether or not the buyer would want to invest in search.

Figure 3.3 shows the buyers' search strategies for all treatments. It can be clearly seen that, for a given price, a buyer is more likely to search if search cost is lower. This accords with basic intuition and holds both for the treatments with and without seller communication.

We note that in treatments where sellers can communicate incentives to invest in search are lower.[9] In particular, the effect is strong for large prices (8 to 10). For instance, when search cost is high, the probability of searching when observing a price of 10 drops from around 80% without communication to less than 20% with communication. For the treatments with low search cost, there is also large drop in the search probability. Moreover, it is interesting to observe that, for high prices, the search probability is decreasing as the price rises, an effect that cannot be observed without communication. These findings are consistent with buyers expecting coordination on prices. If buyers expect sellers to successfully coordinate on high prices, there are clearly less benefits from search.

Thus, we summarize our findings.

[9]Buyers do invest slightly more in search for very low prices (1 to 3) when sellers can communicate, however, those prices are never chosen by sellers in the treatments with communication.

Treatment	*LowFix*	*HighFix*	*LowCom*	*HighCom*
Cut-off strategy	76.25%	86.46%	68.75%	64.06%
Cut-off price	3.84	6.55	7.55	9.02
(Std Dev)	(1.86)	(1.64)	(2.02)	(1.94)

Table 3.4: Cut-off strategies

Result 4. i) Lower search cost lead to larger search incentives. ii) There are smaller incentives to search with communication.

In the one-shot game, theory predicts that buyers apply a cut-off search strategy when sellers compete. That is, up to a cut-off price a buyer should not invest in search but should always invest if the price exceeds this threshold. Table 3.4 displays the share of searches that use a cut-off strategies in each treatment. Without communication there is a large share of cut-off strategies, but with communication the use of cut-off strategies is much lower. With low search cost the percentage is 76.25% compared to 68.75% when communication is allowed. With high search cost, the difference is more pronounced and decreases by more than 20%, from 86.46% in *HighFix* to 64.06% in *HighCom*. As discussed before, one potential explanation for the lower use of cut-off strategies might be successful price coordination when sellers talk. If buyers expect sellers to coordinate on high prices, there is little reason to search. However, lower prices might be interpreted as a break-down of price coordination increasing the incentives to search.

If cut-off strategies are used, we find that the cut-off price is significantly larger with communication (p-value of < 0.01, Mann-Whitney U test) which is consistent with our previous findings that incentives to search are smaller with communication.[10] Without communication, the cut-off price is larger with higher search cost (p-value of < 0.01, Mann-Whitney U test) and with communication we do not obtain significant differences (p-value of 0.36, Mann-Whitney U test).

Result 5. i) Communication leads to less frequent use of cut-off strategies and, if cut-off strategies are used, increases the cut-off price. ii) Without communication, lower search cost decrease the use of cut-off strategies and decrease the cut-off price

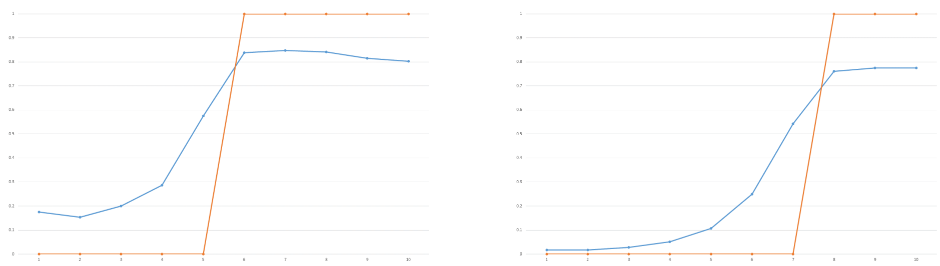

(a) *LowFix* (b) *HighFix*

Figure 3.4: Comparison of optimal and observed search strategies

3.4.3 Gains from search and consumer surplus

In this subsection, we consider the implications for consumer surplus (CS) which takes transaction prices and search efforts into account. Until now we have analyzed the search strategy but did not report to what extent search is actually taking place.

As can be seen in Table 3.3 lowering the search cost increases the frequency of search. Lower search cost lead to an increase in the search propensity from 16.67% to 62.81% without communication, this difference is highly significant (p-value of < 0.01, Mann-Whitney U test). With communication the fraction rises from 23.96% to 43.75% which is not significant (p-value of 0.93, Mann-Whitney U test). Communication decreases the search frequency with low costs, however, the difference turns out to be insignificant (p-value of 0.60, Mann-Whitney U test). In contrast, with high search cost the introduction of communication raises the frequency of search insignificantly (p-value of 0.23, Mann-Whitney U test).

Whether or not the investment in search pays out, depends on the price level and the distribution of prices. Our design allows to analyze whether the observed search strategies are indeed optimal given sellers' pricing strategies. Without seller communication, we calculate the average potential gain of search.[11] The resulting optimal search strategy is displayed in Figure 3.4, *LowFix* in the left panel and *HighFix* in the right panel. The figures reveal that for both levels of search cost a cut-off strategy is the optimal strategy for buyers. The optimal cut-off price is higher for high search cost. Whereas search pays from an observed price of 6 in the low cost treatment, buyers should search for price above 8 when search cost is high. While the observed search strategies (also displayed in Figure 3.4) is roughly in

[10] For a summary of the results see Table 3.5 presented in the Appendix.

[11] For each buyer in every period we calculated whether it would have been beneficial to invest in search or not. Taking the average gain of searching we derived the optimal search strategy.

line with optimal search, there are some differences. In particular, the comparison reveals that there is too much search at low prices and too little search for high prices. In particular, when search cost is low, there is too much search at low prices.

When seller communication is allowed we find that search is never optimal, given the observed pricing strategies, that is, at no price does the expected gain of search outweigh the cost. This finding holds for both levels of search cost. Thus, with seller communication we always observe excessive search.

Result 6. i) Without communication, reducing the level of search cost leads to more frequent search. With communication, there is no significant effect on the search frequency. ii) Without communication, there is too much search at low prices, but to little search at high prices. With communication, there is always too much search.

We can calculate the search effort incurred by non-shoppers by multiplying the search propensity (only for non-shoppers) with the search cost. As demand is inelastic in our setup, the search effort is also a measure of total welfare in this market. Our results show that without seller communication the search effort is independent of the search cost (p-value of 0.46, Mann-Whitney U test). That is, the benefits of lowering the search cost is entirely forgone by the higher search rate. With seller communication, we observe a decrease in search effort from 0.45 to 0.22, but this is not significant (p-value of 0.14, Mann-Whitney U test).

Examining the effect of lowering the search cost on consumer surplus we find that both with and without seller communication consumer surplus does not significantly increase with lower search cost (p-value of 0.22 and p-value of 0.86, Mann-Whitney U test). This is consistent with our finding that average prices do not significantly change. The change in search frequency seems to be outweighed by the level of search cost. Finally, as expected, communication with significantly higher price levels leads to a sharp decrease in consumer surplus (p-value of < 0.01, Mann-Whitney U test).

We summarize the effects of search cost reductions on consumer and total welfare.

Result 7. i) With and without communication, reducing the level of search cost has no effect on total search efforts. ii) With and without communication, reducing the level of search cost has no effect on consumer surplus.

Taken together, our findings suggest that policy interventions which aim at raising competition by lowering consumer search cost may not be effective in markets where industry concentration is high. We do not find evidence that consumers are better off. Neither do we observe higher total welfare as measured by search effort paid by buyers.

3.5 Conclusion

In many markets recent policies aim at reducing consumer search cost. Examples include telecommunication, banking, energy or gasoline markets. Via such interventions, policy makers hope to improve consumers' information and thereby also increase the competitiveness of markets. Many of the aforementioned industries are also characterized by a high market concentration. This paper experimentally analyzes the effects of search cost reductions in concentrated industries by focusing on duopoly markets.

The results of our experiments suggest that market interventions with the aim to decrease consumer search cost are not an effective tool to improve market performance. While we find that consumers do search more intensively when search cost is reduced, we do not observe lower prices. On average, prices with low and high search costs do not differ, but display a larger variation with lower search costs. As a result, consumer surplus does not differ across different search cost levels.

We are also studying markets where sellers are able to communicate. In line with existing studies, we find prices are on average higher, but again do not depend on the level of the search cost. Interestingly, in market where sellers can coordinate more easily, consumers are less likely to invest into search efforts and the search intensity does not depend on the search cost.

As many markets, where interventions to increase consumer information are discussed, are characterized by a high market concentration we focused on the two-seller case in this paper. It would be interesting for future research such interventions in markets with a larger number of sellers and higher intensity of competition.

References

Abrams, Eric, Martin Sefton, and Abdullah Yavas. (2000) "An experimental comparison of two search models," *Economic Theory*, 16: 735 – 749.

Burdett, Kenneth, and Kenneth L. Judd. (2000) "Equilibrium price dispersion," *Econometrica*, 51(4): 955 – 969.

Cason, Timothy N. and Shakun Datta. (2006) "An experimental study of price dispersion in an optimal search model with advertising," *International Journal of Industrial Organization*, 24(3): 639–665.

Cason, Timothy N. and Daniel Friedman. (2003) "Buyer search and price dispersion: A laboratory study," *Journal of Economic Theory*, 112: 232 – 260.

Cooper, David J. and Kai-Uwe Kühn. (2012) "Communication, renegotiation, and the scope for collusion," *American Economic Journal: Microeconomics*, 6(2):247 – 278.

Crosetto, Paolo, and Alexia Gaudeul. (2016) "Choosing whether to compete: Price and format competition with consumer confusion," Working Paper.

Davis, Douglas D., and Charles A. Holt. (1996) "Consumer search costs and market performance," *Economic Inquiry*, 34(1): 133 – 151.

Diamond, Peter A. (1971) "A model of price adjustment," *Journal of Economic Theory*, 3: 156 – 168.

Fischbacher, Urs. (2007) "ZTree: 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(8): 1759 – 1772.

Greiner, Ben. (2015) "Subject pool recruitment procedures: organizing experiments with ORSEE," *Journal of the Economic Science Association*, 1(1): 114–125.

Harrington Jr, Joseph E., Roberto Hernan Gonzalez, and Praveen Kujal. (2016) "The relative efficacy of price announcements and express communication for collusion: Experimental findings," *Journal of Economic Behavior & Organization*, 128: 251–264.

Morgan, John, Henrik Orzen, and Martin Sefton. (2006) "An experimental study of price dispersion," *Games and Economic Behavior*, 54(1): 134 – 158.

Normann, Hans-Theo and Tobias Wenzel. (2016) "Shrouding add-on information: an experimental study," Working Paper.

Orzen, Henrik (2008) "Counterintuitive number effects in experimental oligopolies," *Experimental Economics*, 11: 390 – 401.

Selten, Reinhard. (1967) "Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments," in H. Sauermann (ed.) *Beiträge zur experimentelle Wirtschaftsforschung*, Mohr, 136 – 168.

Stahl, Dale O. (1989) "Oligopolistic pricing with sequential consumer search," *American Economic Review*, 79: 700 – 712.

Varian, Hal R. (1980) "A model of sales," *American Economic Review*, 70: 651 – 659.

Appendix A

Further tables and figures

	Price	Transaction price	Cut-off price	CS
LowFix vs. *HighFix*	0.71	0.39	< 0.01	0.22
LowCom vs. *HighCom*	0.46	0.41	0.36	0.86
LowFix vs. *LowCom*	< 0.01	< 0.01	< 0.01	< 0.01
HighFix vs. *HighCom*	< 0.01	< 0.01	< 0.01	< 0.01

Table 3.5: p-values of pairwise comparisons with Mann-Whitney U test

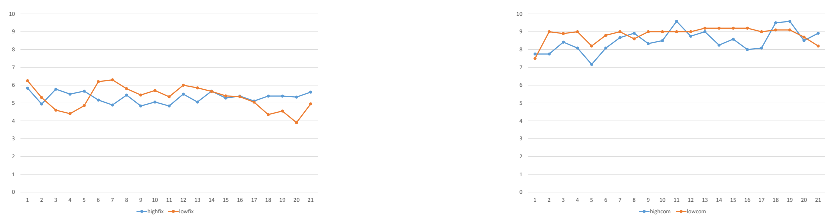

(a) Average prices over time with fixed matching

(b) Average prices over time with fixed matching and communication

Figure 3.5: Average prices over time

	no communication	vs.	communication
low search cost		< 0.01	
vs.	< 0.01		0.30
high search cost		< 0.01	

Table 3.6: Comparison of price distributions: p-vaules of the two-sample Kolmogorov-Smirnov test

Appendix B

Instructions

The following instructions were originally written in German.

Welcome to this experiment! Please read the instructions carefully. The experiment is fully anonymous, i.e., you do not know with whom of the participants you interact. We do not save any personal data in connection with your name.

Depending on your own and the other participants' decisions during the experiment, you can earn experimental currency units (ECU). At the end of the experiment, ECUs are converted into Euro at an exchange rate of 15 ECUs = 1 Euro and will be paid to you in cash. For participating in this experiment, you earn 4 EUR plus the payoff you earn during the experiment.

Please note that you are not allowed to talk to any other participants of the experiment. Please raise your hand if you have any questions regarding the experiment and we will come to your place.

Structure of the experiment

In this experiment, you either take the role of a seller or a buyer in a market. The role is randomly assigned at the beginning of the experiment and communicated to you. You keep your role over the entire experiment.

In the market, there are two sellers and two buyers. Before the experiment starts, two sellers and two buyers are randomly matched. Each constellation is fixed in all periods of the experiment, that is, in all periods you interact with the same participants.

Each seller intends to sell exactly one unit of the product to every buyer. Every buyer intends to buy exactly one unit of the product. Buyers neither have the possibility not to purchase the product nor to buy more than one unit of the product.

Every period of the experiment consists of three stages:

Stage 1:

Decision of sellers

In every period of the experiment, both sellers independently choose their price. The chosen price has to be an integer between 1 and 10 ECU.

[*The following additional instructions were only given in the treatment where communication between sellers was allowed*: Before setting a price, sellers have the opportunity to communicate via a chat window. At the beginning of each round, sellers is given time to chat. In the first five periods, communication is restricted to one minute, in the following rounds, communication is restricted to 45 seconds. Buyers neither can read the communication nor can they communicate with the sellers themselves.]

Decision of buyers

Buyers decide how well they wish to be informed about seller's prices. In stage 2, sellers will obtain the price of one randomly drawn sellers. However, sellers have already to decide in stage 1 for which posted prices they will invest in search. In case you will be assigned to the role of a buyer, the following screen will be shown in stage 1:

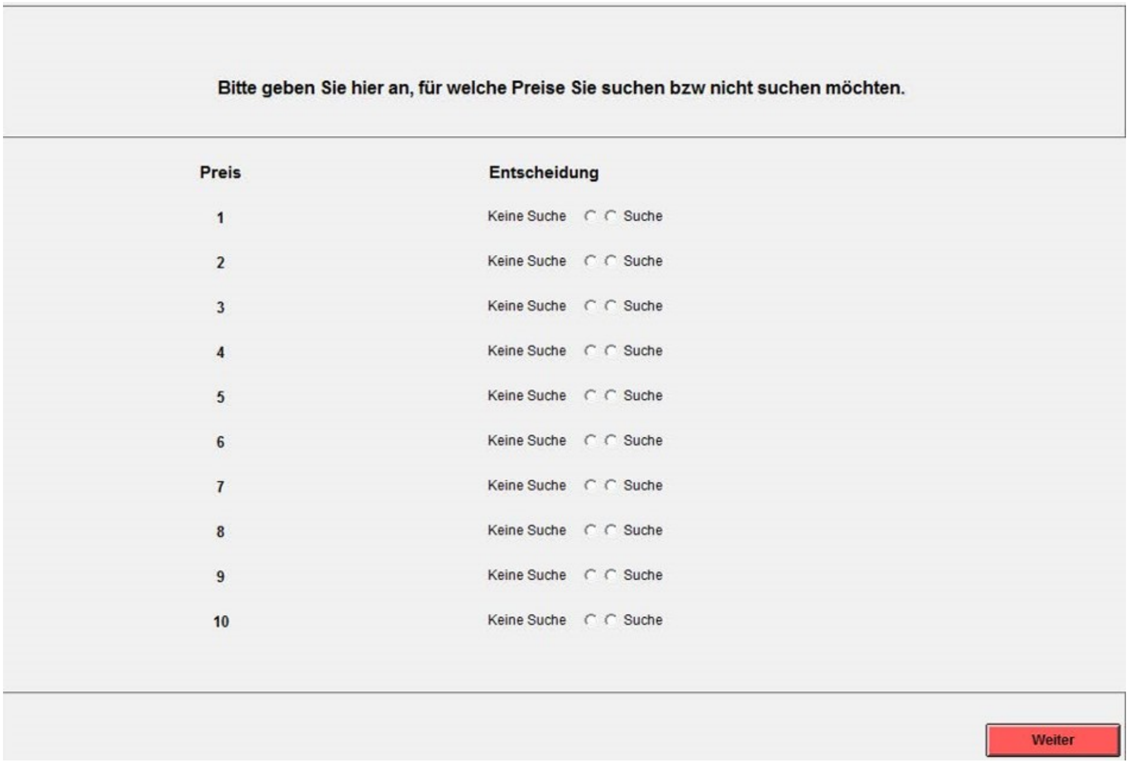

In the left column, all possible prices are displayed. In the right column, you can either decide in favor or against search for any price. You have to choose either "search" or "no search".

In stage 2, you are informed about the price of a randomly determined seller. If you decided to invest in search at the posted price in stage 1, you additionally obtain the price of the other seller and have to pay c ECU. In case you decided against search at that price, you do not receive any additional price information, but you also do not have to incur additional costs.

Stage 2:

In this stage only buyers make decisions.

With a probability of 25%, you obtain price information of both sellers free of costs and you can decide between both sellers' products. In this case, you have no additional costs, independent of your search decisions in stage 1.

With a probability of 75%, you only obtain price information of one seller without costs. Which of the sellers' prices you obtain is randomly determined. Your further options depend on your decision in stage 1. In case, in stage 1, you decided to obtain additional price information at the posted price, you also obtain the price of the other seller. You incur search cost of c. You can decide which of both products to purchase. In case, in stage 1, you decided not to invest in additional information at the posted price, you do not obtain additional price information. But also you do not incur any additional cost.

Stage 3:

At the end of each period you obtain information on your payoff.

Payoff of buyers:

The payoff of a buyer depends on the purchase decision, the purchase price, as well on the potential cost for receiving additional price information. The buyer earns the following payoff:

payment buyer=15-price-search cost

Payoff of sellers:

The payment of sellers depends on the chosen price and the purchase decision of buyers:

payment seller= price * sold quantity

If no buyer decided in favor of the sellers product, the seller would receive zero payoff in that period.

Example 1:

You are a seller and choose a price of 6 ECU. The randomly assigned buyer decided in stage 1 not to invest in search for a price of 6 ECU. The other buyer obtained randomly the information about prices of both sellers and decided to purchase your product. You sold 2 units of your product at a price of 6 ECU which sums up to a profit of 12 ECU for this period.

Example 2:

You are buyer and you have to decide for each price to invest in search or not. You decide not to search, if prices are between 1 and 5 ECU and for prices equal to or larger than 6, you decide to incur the search cost of c and obtain price information of both sellers.

In stage 2, you observe the price of seller 2 which is 7 ECU. Because you decided to search for a price equal to or above 6 ECU, you obtain information on the price of seller 1 as well. The price of seller 1 is 5 ECU. In stage 2, you decide to purchase the product of seller 1, buy one unit of the product, and pay a price of 5 ECU. In total, your payment for this period is $15 - 5 - c = 10 - c$ ECU.

End of the experiment

The experiment will be repeated for at least 20 times. At the end of period 20 (and in the possibly following periods) a random draw determines whether another period will follow. With a probability of 50%, another period follows, otherwise the experiment ends. As already stated above, in each round, you will interact with the same participants. At the end of the experiment, your earnings will be paid out to you in cash. Your earnings comprises the show-up fee and the points you have earned during the experiment.

Chapter 4

Reputation and Foreclosure with Vertical Integration – Experimental Evidence

4.1 Introduction

In 2008, the European Commission conducted an antitrust investigation against E.ON AG Düsseldorf. E.ON was accused of withholding capacity in the wholesale market for energy. In an official statement, the European Commission raised the concern that E.ON had been "deliberately not offering for sale the production of certain power stations which was available and economically rational, with a view to raising electricity prices to the detriment of consumers".[1] E.ON offered to divest energy generation capacity which was accepted by the European Commission.

Input foreclosure was first analyzed in Ordover, Saloner and Salop (1990, henceforth OSS), Hart and Tirole (1990) and Salinger (1988). With duopolies on either production level and one vertically integrated firm, OSS (1990) state that the integrated firm refrains from supplying the downstream rival. This strategy is profitable because the remaining upstream competitor gains monopoly power and therefore increases the price for the input good dramatically. On the downstream level, the non-integrated competitor suffers from increased input prices and the downstream division of the integrated firm profits through the raising-rivals'-costs effect. However, Hart and Tirole (1990) and Reiffen (1992) challenged the assumption of the capability to commit and argued that absent this assumption foreclosure is not a Nash equilibrium. Assuming Bertrand competition upstream, they argue that the integrated firm has an incentive to undercut the upstream rival. In addition to the still existent cost advantage downstream, the upstream division would gain almost monopoly profits.

The critique of Hart and Tirole (1990) and Reiffen (1992) leaves open the possibility that integrated firms may seek opportunities to commit whereas non-integrated firms would not. Indeed, in their reply, OSS (1992) argue that "The notion that vertically integrated firms behave differently from unintegrated ones in supplying inputs to downstream rivals would strike a businessperson, if not an economist, as common sense" (OSS, 1992, p. 698). Such differences in behavior may occur when chances to commit are present: what OSS (1992) show is that integrated firms have an incentive to jump at such opportunities whereas non-integrated firms have no such incentives.

In this paper I analyze the impact of the possibility to build a reputation on prices and upstream foreclosure if one firm is vertically integrated. Reputation is built by revealing the price history to the competitor. I address four questions: Do integrated

[1] See *Antitrust: Commission market tests commitments proposed by E.ON concerning German electricity markets*, Memo European Commission, 12 June 2008, available at `http://europa.eu/rapid/press-release_MEMO-08-396_en.htm?locale=en`

and non-integrated firms behave differently with respect to building a reputation? Does reputation building entail anticompetitive effects by raising the price for the input good? Does establishing a reputation enable the integrated firm to commit to a high price? Consequently, does foreclosure translate into monopolization of the input market?

In an experiment, I conducted three treatments.[2] The treatment *Choose_Rep* is structured in two stages. In the first stage, both upstream firms decide whether they want to build a reputation, hence, reveal their price history to the competitor. In the second stage, firms compete in the market for the input good in one-shot interactions. In treatments *U_1_Rep* and *No_Rep* the first stage changes and the second stage is the same as in *Choose_Rep*. In *U_1_Rep* one-sided reputation building of U_1 is imposed. That is, the non-integrated upstream firm learns all previous price choices of the opponent, while the integrated firm cannot observe the price history of the competitor. A setting without reputation building is studied in *No_Rep*. Thus, no upstream firm learns the previous price choices of the actual competitor.

While one-sided reputation building seems intuitive for seller – buyer relationships, it might be less plausible in a setting with two firms. Intuitively, one-sided reputation building of the integrated firm can be interpreted as the incentive of the non-integrated rival to behave like a "maverick". Maverick firms are tough competitors and attract attention from antitrust authorities as they are known to ensure effective competition. In the Non-Horizontal Merger Guidelines of the European Commission a maverick is defined as "a supplier that for its own reasons is unwilling to accept the co-ordinated outcome and thus maintains aggressive competition."[3].[4] By opting against reputation building, the non-integrated firm can commit to best respond to any action taken by the opponent. Accordingly, the integrated firm can use reputation building to pick an outcome on the best response function of the opponent. In my setting, a non-integrated maverick firm might even support the cooperative foreclosure outcome instead of destabilizing it.

In theory, one-sided reputation building turns out to be an equilibrium in the predictions of *Choose_Rep*, i.e. the integrated firm builds a reputation whereas the non-integrated firm does not. Why is that? Whether foreclosure and monopoliza-

[2]The experimental design is build on Normann (2011) who analyzed the potential anticompetitive effects of vertical integration as compared to no integration.

[3]See *Guidelines on the assessment of non-horizontal mergers under the Council Regulation on the control of concentrations between undertakings*, European Commission (85), October 18, 2008, available at `http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52008XC1018(03)&from=EN`

[4]Note that recent research has highlighted the role of maverick firms in antitrust cases such as horizontal mergers and collusion (see for example Gayle et al., 2009 and Marshall et al., 2016)

tion is an equilibrium does not depend on the non-integrated firm's decision about reputation building. In contrast, reputation building of the integrated firm is essential for foreclosure to be an equilibrium (Hart and Tirole, 1990, Reiffen, 1992). With two-sided reputation building, the grim trigger strategy[5] (Friedman, 1971) supports multiple outcomes in equilibrium, including foreclosure and monopolization. One-sided reputation building restricts the set of equilibria to the outcomes on the best response function of the non-integrated firm (Fudenberg, Maskin and Kreps, 1990). Obviously, equilibria beyond her best response function are strictly worse for the non-integrated firm. Moreover, with two-sided reputation, foreclosure and monopolization is not the payoff dominant[6] equilibrium in pure strategies anymore. Hence, the non-integrated firm might opt against reputation building to facilitate coordination on her favorite outcome and obtain monopoly profits.

Introducing uncertainty about the "type" of integrated firm who builds one-sided reputation leads to a unique prediction in pure strategies. The theoretical model of Fudenberg and Levine (1989) predicts the withdrawal of the integrated firm and monopoly prices in the input market after a finite number of periods. Results of Normann (2011) suggest that even in a static setting a small fraction of integrated firms are committed to foreclose the market. I will name these firm "Stackelberg"[7] types. The model of Fudenberg and Levine (1989) has the following intuition: Non-integrated firms have identical beliefs about types of integrated firms they face in the market. They believe that some integrated firms are committed to the Stackelberg strategy whereas others simply maximize their profits. With one-sided reputation of the integrated firm, profit maximizers start to imitate Stackelberg types. Why is that? By acting like the Stackelberg type and a sufficiently high discount factor, the profit maximizer can obtain almost Stackelberg profits. How? The integrated firm chooses the Stackelberg strategy in each period. The non-integrated firm observes the price history of her opponent and decides upon her own price. In the first periods she is still not convinced that the integrated firm will forgo upstream profits and sets a price strictly below the monopoly level. After a finite number of periods, the

[5]Grim trigger starts with cooperation and cooperates whenever the opponent cooperated in every previous period, otherwise the player applying grim trigger defects.

[6]Payoff dominance is a refinement of equilibria established by Harsanyi and Selten (1988). A payoff dominant Nash equilibrium is Pareto superior to all other equilibria.

[7]As other authors have done before, I name the total foreclosure outcome also as "Stackelberg outcome" with the corresponding "Stackelberg strategies". The reason is that the Stackelberg outcome would result with sequential price choices with the integrated firm being the first mover (Mouraviev and Rey, 2011). Because the raising-rivals'-costs effect is largest, the Stackelberg strategy is the pure strategy the integrated firm favors the most on the best-response function of the non-integrated firm. Consequently, the Stackelberg outcome involves withdrawal of the integrated firm and monopolization of the input market.

probability she attaches to total foreclosure of the integrated firm in the current period is sufficiently high and she best responses by setting the monopoly price. While she is not convinced that her opponent is actually the Stackelberg type, she believes that he will act as if he was. For the integrated firm his patience is beneficial as long as the future is "important enough".

My experimental evidence supports the predictions. The integrated firm chooses to build a reputation significantly more often compared to the non-integrated firm. In fact, one-sided reputation building seems to be empirically relevant. Furthermore, while reputation building of the non-integrated firm does not entail anticompetitive effects, reputation building of the integrated firm leads to substantially higher market prices and more foreclosure. This includes the price of the integrated firm, which increases on average by more than 50%, and the price of the non-integrated firm raises by more than 25%. Resulting in an increase of costs for the independent downstream firm by around 45%. Foreclosure occurs at least three times more often when the integrated firm builds a reputation. In *Choose_Rep* the integrated firm opts himself for reputation building which leads to an even larger difference: withdrawal of the integrated firm raises from around 6% to more than 60% of markets. Finally, foreclosure results in monopolization in less than 2.5% of observations without reputation building of the integrated firm and in more than 25% of observations with reputation building.

While it is clear from a theoretical perspective that reputation building introduces repeated-game effects to the static game and may result in (loosely speaking) collusive or foreclosure effects, there are many ways of how precisely collusive or foreclosure effects may occur. I tested the relevance of three different strategies for non-integrated firms based on the observed price history of the opponent. The described grim trigger strategy is applied for predictions in games with repeated interaction. It translates into choosing the monopoly price in the first period and in every subsequent period if the opponent always withdrew from the input market. Another strategy is based on a model by Fudenberg and Levine (1989). Here, the non-integrated firm needs to be convinced, that the integrated firm withdraws from the input market and monopolization occurs only after several periods. Finally, assuming that not all participants consider the whole price history, I tested a myopic best reply strategy[8]. This strategy assumes that the participant will always best reply to any action taken by the opponent in the previous period. I find evidence for the existence of all three strategies. Hence, the strategies are empirically relevant and were actually exerted by participants.

[8]The myopic best reply strategy is in line with the tit-for-tat strategy suggested in Axelrod (1984).

The paper is structured as follows. I will start with a short summary of the related literature, bringing together theoretical and experimental papers on vertical markets and reputation. In the subsequent section I describe the experimental design followed by the predictions. The results section studies anticompetitive effects, i.e. selling prices for the input good, proceeds with individual prices of upstream firms and partial foreclosure and closes with total foreclosure, monopolization and foreclosure strategies. Finally, I conclude.

4.2 Related literature

In this section, I summarize part of the literature on vertical integration and reputation, both, theoretical papers and related experiments.

Several papers contribute to the discussion of commitment and input foreclosure raised in OSS (1990). In Choi and Yi's (2000) framework upstream firms can either produce a generalized or a specified product. The generalized input good is similarly useful for both downstream firms while each of them would prefer an individually specialized intermediate product. Commitment in vertical integrated markets is realized via specialization of the input. Church and Gandal (2000) analyze a system product consisting of a software and hardware component. They show that integration and foreclosure can be an equilibrium outcome if the value depends on the software component. By making the software incompatible with the rival's hardware, commitment can be achieved. Allain, Chambolle and Rey (2011) show that the necessity of downstream firms to share sensitive information once they trade with an upstream firm might lead to input foreclosure. In a market with two upstream firms and vertical integration the non-integrated downstream firm might be reluctant to exchange information which cannot be protected by property rights. Deals between an integrated upstream supplier and non-integrated downstream firms might not occur due to the concern that information will be leached the downstream division. Allain, Chambolle and Rey (2016) show that vertical integration can create hold-up problems for competitors. If the integrated supplier can commit to be "greedy" or alternatively commits to offer a degraded input to the downstream competitor, hold-up problems occur. On the other hand they show that even without commitment, foreclosure emerges if the quality of the upstream product is non-verifiable.

Theoretical papers extend OSS's (1990) idea. Chen (2001) considered not only the change in incentives upstream but also in the downstream market in case of a vertical merger. He finds collusive effects but also efficiency gains and an ambiguous result for competitive effects in general. Nocke and White (2007) analyze vertical integration in a market with two-part tariffs upstream and repeated interaction. They show that a vertical merger facilitates upstream collusion. In a similar setting but with linear prices Normann (2009) shows that collusion is easier to sustain in a vertically integrated market. Related to reputation of being a "Stackelberg" type, Mouraviev and Rey (2011) show that price leadership can facilitate collusion. In a theoretical model they show that the choice of deciding simultaneously or sequentially about prices can sustain perfect collusion.

Normann (2011) was the first to analyze experimentally the effect of vertical integration on selling prices and market foreclosure. Although he finds a significant increase in the minimum price paid by the independent downstream firm, there is little evidence for total foreclosure. The integrated firm does not withdraw completely from the input market. However, partial foreclosure, i.e. the integrated firm sets a higher price than the non-integrated firm, indeed takes place. In an experiment, Allain et al. (2015) find support for the predictions in Allain, Chambolle and Rey (2016). Vertical integration creates hold-up problems, in particular, if commitment is possible. A related experimental study (Martin, Normann and Snyder, 2001) analyzed the commitment problem of an upstream monopolist to restrict the total quantity for downstream firms to the monopoly level. Public contracts between downstream firms and the upstream monopolist and, alternatively, vertical integration result regularly in monopolization of the input market. In contrast, if firms are independent and contracts are secret, beliefs of downstream firm about the contract offer to the rival determine the outcome. In this case, monopoly power cannot be sustained and market quantity is significantly above the monopoly level. Moellers, Normann and Snyder (2016) extend this study and analyze the impact of communication on the commitment problem. They find that open communication leads to monopolization whereas bilateral communication between the producer and retailers do not lead to the monopoly quantity downstream. Mason and Phillips (2000) analyze the double marginalization problem in a market with two upstream and two downstream firms. They find larger outputs and a higher consumer surplus with both firms vertically integrated as compared to no integration. Durham (2000) finds support for the double marginalization problem if upstream and downstream markets are monopolized, whereas competition downstream eliminates this problem.

By introducing the concept of sequential equilibria, reputation has been analyzed by Kreps and Wilson (1982), Milgrom and Roberts (1982). The sequential equilibrium supports the deterrence of entry in Selten's Chain Store Paradox by building a reputation of being "tough" even in a finitely repeated game. Fudenberg and Levine (1989) show that a long-run player who faces sequentially infinitely many (different) short-run opponents, can commit to the Stackelberg strategy in a simultaneous move game.

Camerer and Weigelt (1988) were the first to test whether the prediction of sequential equilibria holds in an experiment. In a lending game the player in the second stage can either pay back or renege. They implement uncertainty about the type by varying the preference of the borrower. In the majority of cases the player prefers to renege but there is a small exogenous probability that he will prefer to pay back. They find evidence in support of the reputation effects predicted by Kreps and Wilson (1982) and Milgrom and Roberts (1982). Neral and Ochs (1992) replicate the results of Camerer and Weigelt (1988) in an experiment but find deviations from theoretical predictions with different parameters. More recently and adding a pre-play stage which decides if reputation is potentially harmful or beneficial, Grosskopf and Sarin (2010) find that reputation is rarely harmful but it can be beneficial. While they find a positive effect, building a reputation was not as beneficial as predicted by theory.

Experimental studies have analyzed reputation building in a trust game, i.e. the effect of providing feedback on trustees' previous decisions. Several studies (Keser, 2002, Bohnet and Huck, 2004, Bolton, Katok and Ockenfels, 2004 as well as Bohnet, Huck, Harmgart and Tyran, 2005) show that one-sided feedback on previous decisions of trustees increases efficiency substantially. In addition, if trustors can observe histories of other trustors, Bohnet, Huck, Harmgart and Tyran (2005) document an additional positive impact on efficiency. If, on the other hand, trustors get information about all trustees' histories, this has no effect on efficiency as was shown by Huck, Lünser and Tyran (2012). However, when trustors can choose with whom they want to play, efficiency is above 80%. Also related to my work is Kartal, Müller and Tremewan's (2015) study on gradualism. In a setting with repeated interaction and hidden information they analyze the impact of reputation building on trust. Whereas the trustee knows his own type, either a low or a high discount factor, the trustor cannot observe the type of his trading partner. They find strong support for their gradualism theory, i.e. trustors start with a low level of trust and gradually raise the level of trust as long as the trustee returned.

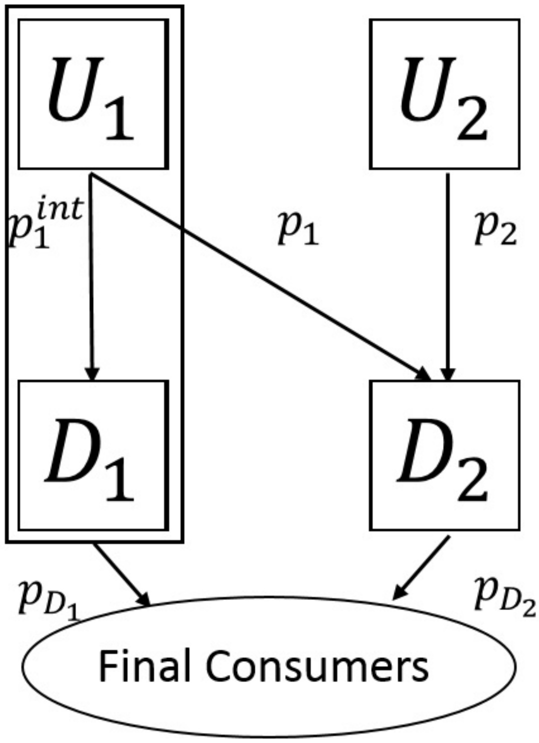

Figure 4.1: Market Structure

4.3 Experimental market

I build on the experimental design of Normann (2011) and use the theoretical model developed by OSS (1990). Figure 4.1 shows the underlying market structure of the experiment. Two competing upstream firms, U_1 and U_2, produce a homogeneous input good with constant marginal cost normalized to zero. Both simultaneously set a price p_i, $i \in \{1, 2\}$ subsequently downstream firm D_2 makes its purchase decision. Because of vertical integration downstream firm D_1 is assumed to purchase the input good internally from U_1 at a price equal to marginal costs ($p_1^{int} = 0$). Retailers produce with a constant returns to scale technology and transformation costs are assumed to be zero. Firms D_i offer the final good for a price of p_{D_i}, $i \in \{1, 2\}$. The demand of final consumers for heterogeneous retailer products is assumed to be $q_i(p_{D_1}, p_{D_2}) = a - bp_{D_i} + dp_{D_j}$ for $i \neq j$ and $i, j \in \{1, 2\}$.[9]

[9]Further specifications of the model as well as the derivation of profits can be found in Normann (2011).

Table 4.1: Payoff

p_1 \ p_2	1	2	3	4	5	6	7	8	9
1	85.5, 19.5	105, 0	105, 0	105, 0	105, 0	105, 0	105, 0	105, 0	105, 0
2	66, 39	101, 27	128, 0	128, 0	128, 0	128, 0	128, 0	128, 0	128, 0
3	66, 39	74, 54	118.5, 34.5	153, 0	153, 0	153, 0	153, 0	153, 0	153, 0
4	66, 39	74, 54	84, 69	136.5, 40.5	177, 0	177, 0	177, 0	177, 0	177, 0
5	66, 39	74, 54	84, 69	96, 81	150, 45	195, 0	195, 0	195, 0	195, 0
6	66, 39	74, 54	84, 69	96, 81	105, 90	181.5, 49.5	231, 0	231, 0	231, 0
7	66, 39	74, 54	84, 69	96, 81	105, 90	132, 99	204, 45	249, 0	249, 0
8	66, 39	74, 54	84, 69	96, 81	105, 90	132, 99	159, 90	216, 36	252, 0
9	66, 39	74, 54	84, 69	96, 81	105, 90	132, 99	159, 90	180, 72	223.5, 25.5

Note: Payoffs (Π_1, Π_2) for each combination of prices (p_1, p_2).

The crucial stage for the commitment problem is the price choice of the upstream firms. To keep the setting as simple as possible, downstream firms as well as final consumers are assumed to decide according to the Nash prediction. Hence, every market is represented by two participants in the laboratory. Both upstream firms simultaneously choose an integer price $p_i \in \{1, 2, ..., 9\}$.

The Nash prediction of the following stages, including the raising-rivals'-costs effect downstream, lead to payoffs in table 4.1. The upstream firm, which sets the lower price, obtains a positive profit in the market for the input good. If both upstream firms set the same price, i.e. $p_1 = p_2$, they will share the Bertrand profit equally. In addition, the integrated firm U_1 benefits from the cost advantage downstream. Depending on the input price $p_{min} := min\,(p_1, p_2)$ of the downstream rival D_2 there is a raising-rivals'-costs effect. The additional profit is positive and increasing in p_{min}.

In every period a participant is randomly matched with another participant in the lab. A random continuation rule of 90% was implemented (which can be interpreted as a discount factor of $\delta = 0.9$ as was done for example in Dal Bó (2005) before) and in total, four supergames were run.

I study three different treatments using the same market structure while varying the available information about competitors. In the baseline treatment, *No_Rep*, none of the firms gets information about previous prices of the opponent, thus, this setting represents a static game. In contrast, in treatment U_1*_Rep* firm U_2 observes the price history of the current opponent (although they may not have met before). However, the integrated firm U_1 does not learn the price history of his competitor. In *Choose_Rep*, participants choose themselves whether they build a reputation. In this treatment I add an additional stage in which both upstream firms, i.e. U_1 and U_2, get the opportunity to choose whether they want to disclose their price history to their competitor. They make this decision separately for each supergame, i.e. four times in total.

4.4 Predictions

Considering the three treatments, four different situations are possible. The matched pair of upstream firms both build a reputation, either U_1 or U_2 reveals previous prices one-sided or nobody does. Throughout the paper I will refer to a player with reputation building as "long-lived" and a player without reputation building as "short-lived" as was done for example by Fudenberg and Levine (1989).

4.4.1 No reputation building

Starting with the static game prediction, I analyze the best-response function of firm U_2 in a first step. Because of homogeneous products and Bertrand competition upstream, the non-integrated firm would like to undercut its rival. With discrete prices $p_i \in \{1, 2, \ldots, 9\}$ of firm U_i and production costs of zero, the best-response function reads (see table 4.1):

$$p_2^{BR}(p_1) = \min\left(\max(1, p_1 - 1), p^M\right)$$

whereas p^M is defined as the monopoly price in the upstream market. In my setting the monopoly price is equal to $p^M = 6$ (see table 4.1).

In a second step, consider the integrated firm U_1 with payoffs displayed in table 4.1. Despite the benefits from high input costs of D_2, the integrated firm has the incentive to undercut its rival on the input market. The gain from undercutting upstream outweigh the decrease of the raising-rivals'-costs effect as was shown by Hart and Tirole (1990) and Reiffen (1992).

$$p_1^{BR}(p_2) = \max(1, p_2 - 1)$$

Both reaction functions lead to the static Nash prediction of $\left(p_1^N, p_2^N\right)$ in equilibrium with $p_1^N := 1$ and $p_2^N := 1$.

4.4.2 Two-sided reputation building

The introduction of reputation building entails a dynamic component. In *Choose_Rep* both firms in the market potentially build a reputation and therefore play an infinitely repeated game. According to the folk theorem (Friedman, 1971) many outcomes can be supported in equilibrium with a grim trigger strategy. It implies cooperation in the first period and in every following period as long as the opponent always cooperated in the past. Once the opponent deviated, the static Nash prediction will be played forever.

Define Π_i^C as the coordination payoff of player i, Π_i^D as the deviation payoff and the payoff in the static game as $\Pi_i^N := \Pi_i(p_1^N, p_2^N)$. Deviation is assumed to occur in the first period because future periods are discounted, hence, the critical discount factor δ_i^{min} can be obtained:

$$\delta_i^{min} = \frac{\Pi_i^D - \Pi_i^C}{\Pi_i^D - \Pi_i^N}$$

The critical discount factor is increasing in the deviation profit, decreasing in coordination profit and increasing in the static game payoff.

In my setup and with a discount factor of $\delta = 0.9$, the set of equilibrium outcomes equals $(p_1, p_2) \in \ddot{S}^2$ with

$$\begin{aligned}\ddot{S}^2 := &\{(p_1, p_2) \,| p_1 = p_2, p_1 < 9\} \cup \left\{(p_1, p_2) \,| p_2 = p_2^{BR}(p_1), p_1 > 4\right\} \\ &\cup \{(8,7), (9,7), (9,8)\}.\end{aligned}$$

In equilibrium, the minimum price p_{min} is in the set $p_{min} \in \{1, 2, 3, 4, 5, 6, 7, 8\}$.

4.4.3 One-sided reputation building

Throughout this section, I assume that the long-lived player can only choose from a finite set of pure strategies. The second assumption I make is that the short-lived player $j \in \{1, 2\}$ always chooses her best reply, i.e. only outcomes for firm j and opponent $i \neq j$ on $p_j^{BR}(p_i)$ are possible. Fudenberg, Maskin and Kreps (1990) show that with these assumptions in games with one long-lived and one short-lived player a variant of the Folk theorem holds. The restriction to the best-response function of the short-lived player reduces the set of equilibrium outcomes compared to settings with two long-lived players.

One-sided reputation building of the non-integrated firm uniquely results in the static game prediction $\left(p_1^N, p_2^N\right)$. The reason is that best responses of firm U_1 always lead to zero profit for U_2 except if both choose a price of $p_1 = p_2 = 1$, i.e. $\Pi_2\left(p_1^{BR}(p_2), p_2\right) \neq 0 \Leftrightarrow p_1^{BR}(p_2) = p_2 = 1$.

In contrast, one-sided reputation building of U_1 leaves us with several equilibria. The set of equilibrium outcomes $\dot{S}^2$ equals

$$\dot{S}^2 := \{(1, 1)\} \cup \left\{(p_1, p_2) \,| p_2 = p_2^{BR}(p_1), p_1 > 4\right\}$$

with $\dot{S}^2 \subset \ddot{S}^2$. Hence, the set of possible equilibria lies within the set of equilibria with two long-lived players but is strictly smaller in my setting. Market prices $p_{min} \in \{1, 4, 5, 6\}$ are supported in equilibrium.

Intuitively one would expect the long-lived player U_1 to coordinate on the equilibrium he likes the most, i.e. maximize its profit Π_1 restricted to the best response function of U_2:

$$\max_{p_1} \Pi_1(p_1, p_2^{BR}(p_1)).$$

In line with the intuition of OSS (1990) this optimization program leads to complete withdrawal of the integrated firm U_1 (Stackelberg outcome). For the sake of convenience I denote the solution to the optimization program above with $\tilde{p}_1 \in \left\{p_1 | p_1 > p^M\right\} =: \tilde{S}_1$ and let $\tilde{p}_2 := p_2^{BR}(\tilde{p}_1) = p^M$.

Introducing uncertainty about the type of the long-lived player leads to the required restriction on the set of equilibria as was shown by Fudenberg and Levine (1989).[10] Following their line of reasoning I assume that there is a certain fraction of long-lived players whose preferences are such that the choice of $\tilde{p}_1 \in \tilde{S}_1$ is strictly favored in the repeated game. I define these long-lived players as Stackelberg type ω^*. Let type ω_0 be a long-lived player who prefers to undercut his rival. In addition to these two types there might be other types, for example type ω_l who strictly prefers to choose price $l \in \{1, \ldots, 6\}$ in the repeated game. Whereas the long-lived player knows his own type, the short-lived players have identical beliefs $\mu(\omega)$ about each type $\omega \in \Omega$. I assume that the short-lived players believe the probabilities of types ω^* and ω_0 are strictly positive, i.e. $\mu(\omega^*) > 0$ and $\mu(\omega_0) > 0$.

The idea of Fudenberg and Levine (1989) is the following: Suppose the short-lived players believe that some of the long-lived players, say a fraction of $\mu^* := \mu(\omega^*) > 0$, is initially committed to play the Stackelberg strategy $\tilde{p}_1 \in \tilde{S}_1$. For a sufficiently large discount factor, long-lived players will imitate the Stackelberg types in order to obtain profits close to $\Pi_1(\tilde{p}_1, \tilde{p}_2)$. If the long-lived player chose $\tilde{p}_1 \in \tilde{S}_1$ in every previous period, the short-lived player would become convinced after some time that he will set $\tilde{p}_1 \in \tilde{S}_1$ in the current period as well. After k periods the short-lived player will choose $\tilde{p}_2 = p_2^{BR}(\tilde{p}_1) = p^M$ because the probability she attaches to the price $\tilde{p}_1 \in \tilde{S}_1$ exceeds the required threshold. However, this does not necessarily mean that the non-integrated firm will change her belief about the type $\omega \in \Omega$ of her opponent.

Let me start with calculating the required number of periods k which are needed to convince the short-lived player to set $\tilde{p}_2$. First, it depends on the initial belief μ^*; the smaller μ^* the larger k. As I do not know anything about initial beliefs I take results from previous experiments. Normann (2011) found that 1 out of 20 participants seemed to be committed to $\tilde{p}_1 \in \tilde{S}_1$ in a treatment similar to *No_Rep*.[11] Hence, I will define $\mu^* := 0.05$ in my setting. Second, k depends on the critical fraction $\bar{f}$ of long-lived players choosing $\tilde{p}_1 \in \tilde{S}_1$. If the short-lived player chooses $\tilde{p}_2$, she either gets the monopoly profit or, in case of deviation, she gets nothing. On the other hand, U_2 can at least secure the profit from the static game prediction. With a price choice of p_2^N, she sets the lower price with a probability of at least μ^*,

[10] In my setting the type is not reflected in the actual payoffs, i.e. the payoff function Π_i only depends on i. I focus in my paper on heterogeneity of preferences for an equilibrium.

[11] In contrast to my setting Normann did not implement a random stopping rule but rather has a fixed number of 15 periods in his treatment. However, in *No_Rep* I find a similar fraction of Stackelberg types, 1 out of 16 participants chose without any exception a price of $\tilde{p}_1$.

i.e. 39 ECU, and has to share upstream profits with a probability of at most $1-\mu^*$. A lower bound on the fraction $\bar{f}$ can be obtained:[12]

$$\begin{aligned} 99\bar{f} + 0\left(1-\bar{f}\right) &= 19.5\left(1-\mu^*\right) + 39\mu^* \\ \Rightarrow \quad \bar{f} &\approx 0.21 \end{aligned}$$

Obviously, there is a positive number $k > 0$ since the initial belief μ^* is strictly smaller than the required fraction $\bar{f}$ of long-lived players choosing $\tilde{p}_1$, i.e. $\mu^* < 0.21$. Fudenberg and Levine's (1989) model implies that:

$$\begin{aligned} k &= \frac{\log(\mu(\omega^*))}{\log(\bar{f})} \\ &= \frac{\log(0.05)}{\log(0.21)} \\ &\approx 1.90. \end{aligned}$$

Rounding up leads to the conjecture that the number of periods equals $k = 2$. After 2 periods of choosing $\tilde{p}_1$ the short-run player will play her best response $\tilde{p}_2$. Therefore, the long-run player can assure himself at least a payoff of:

$$\begin{aligned} \Pi_1^{min} &= 66 + 66\delta + 132\frac{\delta^2}{1-\delta} \\ &= 1194.6. \end{aligned}$$

The normalized present value is $(1-\delta)\,\Pi_1^{min}$ which equals 119.46. This threshold for the payoff of the integrated firm cannot be reached by committing to any other pure strategy.[13] Therefore, commitment to the Stackelberg outcome $(\tilde{p}_1, \tilde{p}_2)$ is the unique prediction in my experimental market.

[12] Because prices are not as competitive as predicted (Normann, 2011), I recalculated with actual obtained payoffs of U_2 in *No_Rep*. With an average payoff of $\Pi_2 = 37.92$, results are $\bar{f} \approx 0.38$ and $k \approx 3.12$. Hence, 3 or 4 periods might be a more realistic bound for the time needed to convince non-integrated firms. The normalized present value for integrated firms equals $(1-\delta)\,\Pi_1^{min} = 109.3$ with $k = 4$ which is still larger than the profit obtained with any other pure strategy and restriction on p_2^{BR} (compare table 4.1).

[13] It might be possible to reach a payoff of 119.46 with a mixed strategy of the long-lived player but I do not consider mixed strategies here. In addition, as Fudenberg, Maskin and Kreps (1990) have shown, for equilibria with unobservable mixed strategies, observed actions and one long- and one short-lived player the Folk theorem does not hold.

4.5 Hypotheses

In this section I will state my hypotheses based on the predictions. Before I hypothesize outcomes of all three treatments, I make some definitions for the sake of clarity.

Definition. Let anticompetitive effects be defined as a comparative static change which raises p_{min} significantly.

As p_{min} equals input costs for the independent downstream firm, it determines the price setting downstream and therefore the consumer surplus.

Definition. Partial input foreclosure (alternatively partial foreclosure) occurs when the integrated firm sets on average larger prices than the non-integrated firm, i.e. the fraction of $p_1 > p_2$ is larger than the percentage of $p_1 < p_2$.[14]

Definition. Total input foreclosure (alternatively total foreclosure) occurs when the price of the integrated firm is above the monopoly price, i.e. $\tilde{p}_1 \in \tilde{S}_1$. That is, the integrated firm withdraws from the market.

The decision about reputation building in *Choose_Rep* is relevant for the predictions in the pricing stage, consequently, I begin with the hypothesis about the choice of reputation building.

The non-integrated firm can meet an integrated firm with or without reputation building. If the integrated firm does not build a reputation, the prediction for U_2 would not depend on whether she builds a reputation; in both cases, the static Nash prediction is the unique equilibrium. If, on the other hand, the integrated firm builds a reputation, the Stackelberg outcome, implying monopoly profits, would be an equilibrium in either case. Reputation building of U_2 is not necessary, it may even harm the achievement of the most favored equilibrium as predictions are less distinct. In any case, there is no incentive for U_2 to choose reputation building.

In contrast, if U_2 opts against reputation building, the integrated firm would have an incentive to build a reputation. As shown in the previous subsection, one-sided reputation building of the integrated firm can lead to substantially higher payoffs. Also, if the non-integrated firm decides to reveal previous prices, there would be an incentive to show the price history as well. While one-sided reputation building of U_2 leads to the static Nash prediction, equilibria with two long-lived players are by definition strictly favorable to the static game outcome (except $\left(p_1^N, p_2^N\right)$ itself).

[14] Because only U_1 has an incentive to foreclose the market, I focus on foreclosure of the integrated firm.

Hypothesis 1. In *Choose_Rep* the integrated firm decides in favor of reputation building whereas the non-integrated firm opts against it.

Hypotheses do not differ between U_1*_Rep* and *Choose_Rep* because I do not expect differences between imposed reputation building of U_1 and the outcome of reputation building decisions, i.e. U_1 opts for reputation building whereas U_2 decides against it. Concerning anticompetitive effects I hypothesize:

Hypothesis 2. In *No_Rep* the selling price is $p_{min} = 1$.[15]

Hypothesis 3. U_1*_Rep* entails anticompetitive effects compared to *No_Rep*.

Hypothesis 4. *Choose_Rep* entails anticompetitive effects compared to *No_Rep*.

In general, total foreclosure implies partial foreclosure. Theoretically, one-sided reputation building of U_1 restricts the set of equilibria to the total foreclosure outcome.[16] Therefore, I hypothesize about foreclosure:

Hypothesis 5. In *No_Rep* neither partial nor total foreclosure occurs.

Hypothesis 6. In U_1*_Rep* total foreclosure occurs more often than in *No_Rep*.

Hypothesis 7. In *Choose_Rep* total foreclosure occurs more often than in *No_Rep*.

The response of the non-integrated upstream firm to total foreclosure is relevant for the input prices of the independent downstream firm D_2, therefore, also for prices downstream and consumer surplus. According to the prediction, I expect that U_2 chooses a price $p_2 < \tilde{p}_2$ in the first k periods and afterwards, provided that U_1 totally foreclosed the market in every previous period, sets a price $p_2 = \tilde{p}_2$. Consequently, I hypothesize:

Hypothesis 8. Total foreclosure of U_1 in U_1*_Rep* and *Choose_Rep* leads to monopoly prices after k periods.

[15]Note that from earlier experiments (Normann, 2011) it is known that even with random matching and finitely repeated interaction, partial foreclosure, i.e. $p_1 > p_2$, and selling prices above Nash occur.

[16]However, without the introduction of different types ω, several outcomes were supported in equilibrium. Except the static Nash prediction all of them implied $p_1 > p_2$ but not necessarily $p_1 \in \tilde{S}_1$.

4.6 Procedures

Participants were invited via ORSEE (Greiner, 2015). Upon arrival in the laboratory, subjects were assigned the role of U_1 or U_2 which stayed the same during the whole session. After reading the instructions and having the opportunity to ask questions privately, the experiment proceeded. The experiments were programmed using zTree (Fischbacher, 2007). The number of periods for each of the four supergames were randomly predetermined to be 16, 6, 10 and 7.[17] Every subject was randomly matched with another subject in every period of the session both within and between supergames.

The experiments were conducted in the DICE laboratory at the University of Düsseldorf in June and July 2015. In each of 8 sessions between 16 and 18 subjects participated, the total number of subjects was 136. The three treatments as well as the number of subjects per treatment are summarized in table 4.2. Sessions took about one hour and at the end 300 ECUs (Experimental Currency Units) were exchanged for 1 Euro. Earnings were on average 15.13 Euro.

Table 4.2: Treatments

	No_Rep	*Choose_Rep*	*U_1_Rep*
random matching	yes	yes	yes
reputation building U_1	no	optional	yes
reputation building U_2	no	optional	no
number of subjects	32	70	34

4.7 Results

In the first subsection I present the results of the choice of reputation building in *Choose_Rep*. I proceed with anticompetitve effects, partial foreclosure and finally analyze total foreclosure. I distinguish four different outcomes in the *Choose_Rep* treatment; in *NoRep* two randomly matched firms both choose not to show their price history, in *U_iRep* solely firm $i \in \{1, 2\}$ decided to build a reputation whereas in *BothRep* both firms reveal previous price choices.

[17]Note that the expected number of periods with a random continuation rule of $\delta = 0.9$ is 10.

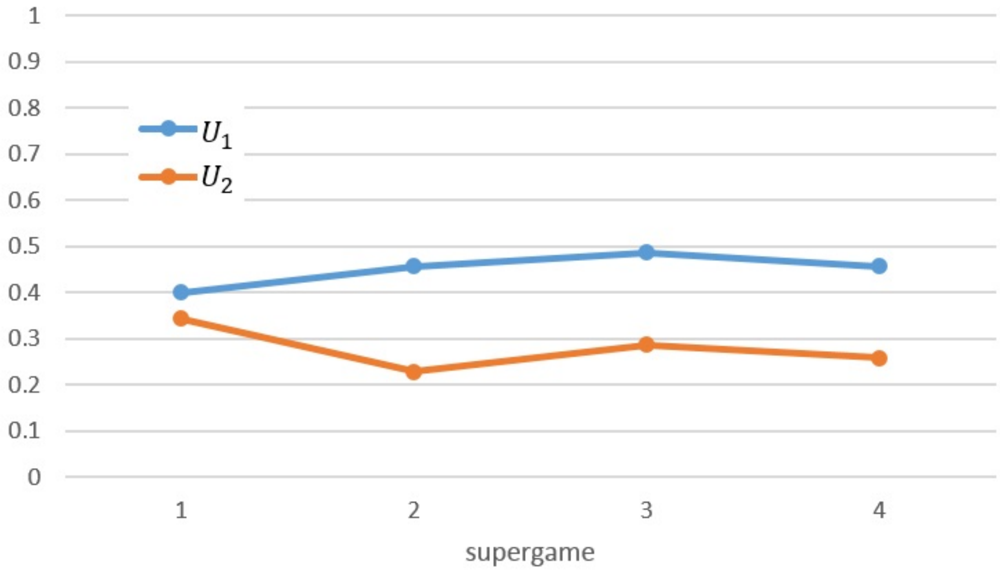

Figure 4.2: Fraction of reputation building

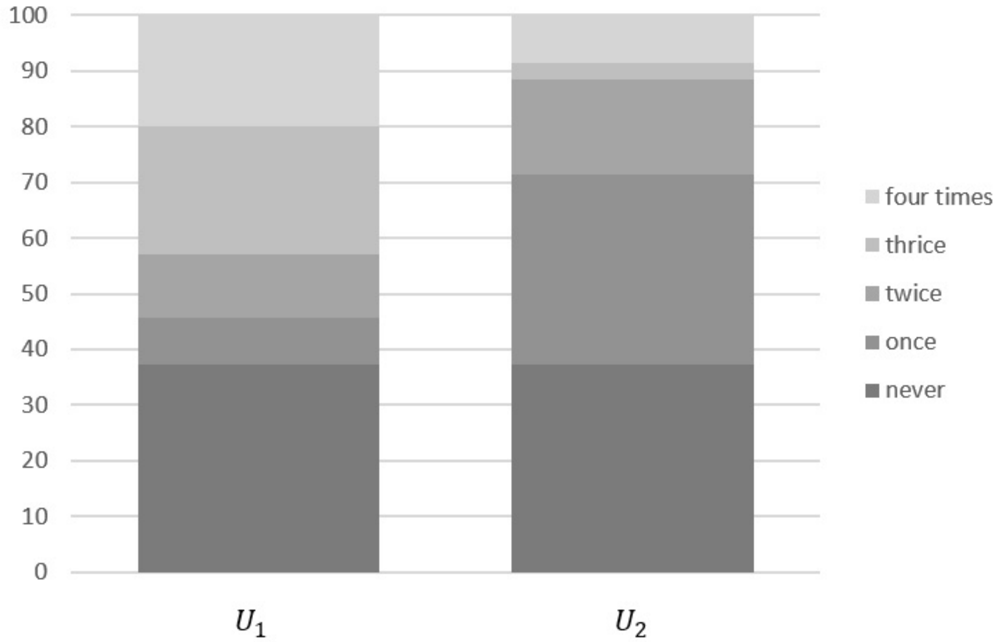

Figure 4.3: Frequency of reputation building

4.7.1 Choice of reputation building

Figure 4.2 depicts the fraction of U_1 and U_2 choosing to build a reputation in *Choose_Rep* for each supergame. On average, the integrated firm U_1 chooses more often to build a reputation than the non-integrated firm U_2 in all four supergames. However, in contrast to the conjecture, a substantial fraction (more than 50%) of integrated firms opts against reputation building in each supergame and around 25% of U_2's build a reputation.

Turning to the frequency of the individual choice to build a reputation (figure 4.3), exactly the same percentage of U_1 and U_2 (37.14%) never opts for it. The remaining 62.86% of all participants differ in their behavior depending on the firm type. The majority of integrated firms choose to build a reputation three or four times (22.86% and 20%) whereas the majority of non-integrated firms build a reputation once or twice (34.29% and 17.14%). This might be interpreted as learning effects.

The difference between firm types is confirmed to be significant at the 1% level (regression (1) in table 4.4[18]). While the comparative static result holds, differences are less pronounced than expected.

Result 1. Hypothesis 1 cannot be rejected, integrated firms opt for reputation building significantly more often than non-integrated firms.

[18]Throughout the analysis, I define dummy 1_E to be equal to 1 if statement E, i.e. equality or inequality E, holds and 0 otherwise.

Table 4.3: Descriptive statistics

	No_Rep	*Choose_Rep*				*U_1_Rep*
		NoRep	*U_2Rep*	*BothRep*	*U_1Rep*	
p_{min}	2.87	2.91	3.26	4.72	4.32	4.18
	(1.35)	(1.16)	(1.44)	(1.82)	(1.78)	(1.73)
p_{min}	2.87	3.65				4.18
	(1.35)	(1.68)				(1.73)
$p_1 > p_2$	48.71%	42.29%	33.33%	67.44%	75.12%	68.78%
$p_1 < p_2$	30.61%	31.2%	40.69%	18.02%	12.33%	18.25%
$p_1 = p_2$	20.67%	26.5%	25.97%	14.53%	12.55%	12.97%
$\tilde{p}_1$	18.11%	5.64%	5.63%	62.21%	60.7%	48.72%
$(\tilde{p}_1, \tilde{p}_2)$	1.12%	0.75%	2.16%	30.81%	29.07%	26.4%
Obs.	624	532	231	172	430	663
$\tilde{p}_1$ whole session	6.25%	8.57%				5.88%
Obs.	16	35				17
$\tilde{p}_1$ whole supergame	9.38%	1.3%		42.86%		22.06%
Obs.	64	77		63		68

Notes: In the *Choose_Rep* treatment I distinguish four different outcomes; in *NoRep* two randomly matched firms both choose not to show their price history, in *U_iRep* only firm $i \in \{1, 2\}$ decided to build a reputation whereas in *BothRep* both firms reveal previous price choices. Note that these outcomes are not independent, even within one supergame decisions of one firm are probably present in two groups. I define p_{min} as the selling price upstream, p_i as price of U_i. $\tilde{p}_1$ denotes a price above monopoly level, i.e. total foreclosure of U_1, and $(\tilde{p}_1, \tilde{p}_2)$ total foreclosure of U_1 and the monopoly price set by U_2 which results in monopolization of the input market. I have three different levels for total foreclosure, $\tilde{p}_1$ counts each period as single observation, $\tilde{p}_1$ "whole session" means that this participant chose $\tilde{p}_1$ in each and every period of the whole session and finally, $\tilde{p}_1$ "whole supergame" is the fraction of total foreclosure in each and every period of one supergame. Standard deviations are reported in parentheses. Note that I normalized prices above 7 to to 7 for both firms.

Table 4.4: Choice of reputation building and anticompetitive effects

	Dependent variable			
	(1)	(2)	(3)	(4)
	1_{Rep}	p_{min}	p_{min}	p_{min}
1_{U_1}	0.46***			
	(0.11)			
$1_{Rep.\ U_1}$				1.43***
				(0.16)
$1_{Rep.\ U_2}$				0.37
				(0.22)
$1_{Both\ Rep.}$				0.07
				(0.04)
1_{No_Rep}		-0.78	-1.04**	-0.03
		(0.44)	(0.43)	(0.4)
$1_{U_1\_Rep}$		0.54	0.6	-0.14
		(0.55)	(0.56)	(0.59)
$1_{Per.\ 6-10}$		-0.32***	-0.3***	-0.32***
		(0.06)	(0.06)	(0.06)
$1_{Per.\ 11-16}$		-0.42**		-0.42**
		(0.12)		(0.12)
$1_{SG\ 2}$		0.04		0.03
		(0.09)		(0.12)
$1_{SG\ 3}$		-0.05	-0.1	-0.1
		(0.15)	(0.12)	(0.15)
$1_{SG\ 4}$		-0.11	-0.16	-0.14
		(0.2)	(0.16)	(0.18)
Constant	-0.59***	3.85***	3.94***	3.13***
	(0.08)	(0.38)	(0.36)	(0.33)
Obs.	280	2652	1564	2652
R^2		0.09	0.12	0.19
Pseudo R^2	0.02			

Notes: Column (1) shows results of a probit regression of the reputation building on firm type clustered at session level. Columns (2)-(4) represent an ordinary least squares regression clustered at session level. Except in regression (2) all periods are included. Minimum prices are regressed upon dummy variables for reputation, imposed reputation and no reputation building. I include dummies for different phases in the game, e.g. $1_{Per.\ 6-10}$ for periods 6 - 10, and the number of the supergame, e.g. $1_{SG\ 2}$ for supergame 2. Dummies for supergames are included as Selten and Stöcker (1986) find learning effects between supergames in a finitely repeated prisoners' dilemma. Standard deviations are reported in parentheses. Note that I normalized prices above 7 to 7 for both firms. Significantly different from 0 in a two-tailed test at the *10% level, **5% level, ***1% level.

Table 4.5: Partial and total foreclosure

	Dependent variable					
	(5)	(6)	(7)	(8)	(9)	(10)
	p_1	p_2	$1_{p_1>p_2}$	$1_{p_1>p_2}$	$1_{\tilde{p}_1}$	$1_{\tilde{p}_1}$
$1_{Rep.\ U_1}$	2.06***	1.06***		0.88***		1.87***
	(0.28)	(0.11)		(0.21)		(0.15)
$1_{Rep.\ U_2}$	0.19	0.55*		-0.23*		0.05
	(0.13)	(0.27)		(0.13)		(0.14)
$1_{Both\ Rep.}$	-0.17	-0.004		0.00		0.00
	(0.11)	(0.12)		(0.16)		(0.16)
1_{No_Rep}	0.44	-0.01	-0.14	0.16	-0.39***	0.69***
	(0.38)	(0.41)	(0.11)	(0.11)	(0.15)	(0.18)
$1_{U_1\_Rep}$	-0.3	-0.04	0.38**	-0.19	0.49**	-0,3
	(0.61)	(0.45)	(0.16)	(0.26)	(0.24)	(0.31)
$1_{Per.\ 6-10}$	-0.46***	-0.31***	-0.13***	-0.14***	-0.14***	-0.16***
	(0.09)	(0.08)	(0.05)	(0.05)	(0.04)	(0.05)
$1_{Per.\ 11-16}$	-0.48***	-0.52***	-0.00	-0.00	0.23	-0.03
	(0.12)	(0.11)	(0.07)	(0.07)	(0.14)	(0.05)
$1_{SG\ 2}$	-0.02	-0.05	0.00	-0.04	0.27**	0.27**
	(0.13)	(0.09)	(0.1)	(0.1)	(0.11)	(0.12)
$1_{SG\ 3}$	-0.06	-0.16	0.1*	0.06	0.27**	0.23*
	(0.13)	(0.13)	(0.05)	(0.04)	(0.13)	(0.13)
$1_{SG\ 4}$	-0.06	-0.23	0.09	0.05	0.23	0.21
	(0.17)	(0.2)	(0.11)	(0.11)	(0.14)	(0.15)
Constant	3.98***	3.74***	0.11	-0.17*	-0.63***	-1,68***
	(0.32)	(0.28)	(0.07)	(0.1)	(0.12)	(0.13)
Obs.	2652	2652	2652	2652	2652	2652
R^2	0.21	0.12				
Pseudo R^2			0.02	0.06	0.05	0.21

Notes: Columns (5)-(6) represent ordinary least squares regressions and (7)-(10) are probit regressions clustered at session level. All periods are included. Price choices of U_1 and U_2 as well as partial foreclosure, i.e. $1_{p_1>p_2}$, and total foreclosure $1_{\tilde{p}_1}$ are regressed upon dummy variables for reputation building $1_{Rep.\ U_i}$, $1_{Both\ Rep.}$, imposed reputation building $1_{U_1\_Rep}$ and no reputation building 1_{No_Rep}. I include dummies for different phases in the game, e.g. $1_{Per.\ 6-10}$ for periods 6 - 10, and the number of the supergame, e.g. $1_{SG\ 2}$ for supergame 2. Dummies for supergames are included as Selten and Stöcker (1986) find learning effects between supergames in a finitely repeated prisoners' dilemma. Standard deviations are reported in parentheses. Note that I normalized prices above 7 to 7 for both firms. Significantly different from 0 in a two-tailed test at the *10% level, **5% level, ***1% level.

4.7.2 Anticompetitive effects

Table 4.3 summarizes outcomes in each of the treatments. The lowest minimum prices are obtained in *No_Rep* while averages are significantly larger than 1 (at the 1% level, confirmed in regression (2) table 4.4).[19] Imposed reputation building of U_1 in U_1*_Rep* leads to a price increase of 45.64% compared to *No_Rep*, the average price paid by D_2 is 4.18. In column (2) of table 4.4 the effect is confirmed to be significant at the 5% level.

Result 2. Hypothesis 2 can be rejected, average minimum prices in *No_Rep* are significantly (1% level) larger than $p_{min} = 1$.

Result 3. Hypothesis 3 cannot be rejected, average minimum prices in U_1*_Rep* are significantly (5% level) larger than average prices in *No_Rep*.

I distinguish four different outcomes in *Choose_Rep*: *NoRep*, U_2*Rep*, U_1*Rep* and *BothRep*. The averages in table 4.3 (2.91, 3.26, 4.72 and 4.32) suggest that, while the decision of U_1 to build a reputation has an impact on minimum prices, reputation building of firm U_2 does not play a role. In addition, comparing treatments *No_Rep* and U_1*_Rep* with their corresponding outcomes in *Choose_Rep* give similar results, i.e. the choice whether to build a reputation does not affect outcomes. These observations are confirmed in regression (4) in table 4.4, the impact of U_1's reputation building is highly significant (1% level) whereas reputation building of firm U_2 as well as treatment variables are not significant.

Considering all periods without the distinction between outcomes, treatment *Choose_Rep* is not significantly different from both treatments (regression (2), table 4.4). However, the average of 3.65 is closer to results obtained in U_1*_Rep*. And indeed, differences in market prices between *No_Rep* and *Choose_Rep* turn out to be significant (5% level) considering only supergames 2 - 4 (column (3) table 4.4).

Taking learning effects into account, I conclude:

Result 4. Hypothesis 4 cannot be rejected, the possibility of building a reputation has anticompetitive effects.

[19]The outcome is remarkably similar to results obtained by Normann (2011), 2.83 vs. 2.87 in *No_Rep*.

Table 4.6: Average prices in the first period of the first supergame

Treatment	1	2a	2b	3
	No_Rep	*Choose_Rep*		*U_1_Rep*
		NoRep	*Rep*	
p_1	4.81	4.52	6.14	5.41
	(1.56)	(1.6)	(1.29)	(1.77)
Obs.	16	21	14	17
p_2	5	4	4.67	4.47
	(0.97)	(1.65)	(1.23)	(1.23)
Obs.	16	23	12	17
H_0: $p_1 \equiv p_2$	0.73	0.23	0.01	0.03

Notes: Because it is the first period *Choose_Rep* is only divided in two outcomes depending on the own reputation building decision, i.e. *NoRep* if the firm decided not to show previous prices and *Rep* if she reveals the price history. Standard deviations are reported in parentheses. Note also that I normalized prices above 7 to 7. I performed a Wilxocon rank-sum test and reported p-values for the null hypothesis H_0: $p_1 \equiv p_2$ for each treatment and outcome in *Choose_Rep*.

Table 4.7: Mann-Whitney U test for treatment differences in the first period of the first supergame

Treatment comparisons	1 vs. 3	2a vs. 2b	1 vs. 2b	2a vs. 3	1 vs. 2a	2b vs. 3
p_1	0.2	0.00	0.02	0.06	0.68	0.19
p_2	0.2	0.18	0.4	0.34	0.03	0.61

Notes: p-values of Mann-Whitney U tests conducted within firm type and between treatments are reported. Treatment 1 is *No_Rep*, outcome 2a is *NoRep* and outcome 2b is *Rep* in *Choose_Rep*, finally, treatment 3 is *U_1_Rep*

4.7.3 Individual pricing decisions and partial foreclosure

In a first step, I analyze price setting in the first period of the first supergame. The results are reported in table 4.6. In *Choose_Rep* I only differentiate between participants who decided to show their own price history, i.e. *Rep*, or who do not reveal price choices in *NoRep*.

A Wilcoxon rank-sum tests confirms a positive effect of the own reputation building decision for p_1 in the first period. Compared to *No_Rep* price p_1 increases on average with reputation building in outcome *Rep* of *Choose_Rep* and U_1*_Rep* (table 4.6), although, differences are larger and turn out to be only significant in *Choose_Rep* (table 4.7). In contrast, neither the own reputation building nor reputation building of U_1 has an impact on p_2. In particular, U_2 does not seem to be a cooperator from the first period. In contrast to p_1, prices p_2 are not affected by imposed reputation of U_1. This leads to the conjecture that non-integrated firms need to be "convinced" that their opponent is a cooperator which is in line with the predictions based on Fudenberg and Levine (1989).

Comparing prices of the integrated vs. non-integrated firm, I obtain highly significant differences between prices p_1 and p_2 in U_1*_Rep* and outcome *Rep* of *Choose_Rep* (at the 5% and 1% level, respectively; table 4.6). In contrast, without reputation building there are no differences between prices of the integrated and non-integrated firms. These results suggest that partial foreclosure is related to reputation building of U_1. In *Rep* of *Choose_Rep* a substantial fraction of U_1 seem to withdraw completely from the input market, even average prices are above the monopoly price (compare table 4.6).

Figure 4.4 shows average session prices p_1 and p_2 for each treatment using all observations. Reputation building of U_1 leads to substantially higher prices for both firms, confirmed to be significant at the 1% level (table 4.5, regressions (5) and (6)). Although the average choice of U_1 seems to be slightly larger when firms opt themselves for reputation building, the difference between imposed and non-imposed reputation effects turns out to be insignificant. Reputation building of U_2 does not seem to have an impact on p_1 whereas there is a weakly significant positive effect on p_2 (10% level). I do not find any differences between supergames but there is a slight downward trend for both pricing decisions after five periods.

In each session, averages of p_1 are larger than p_2, with the only exceptions in outcomes U_2*Rep* and *BothRep* (compare figure 4.4). The relative difference between p_1 and p_2 increases with U_1's reputation building and is even more pronounced if the decision to build a reputation is made by themselves. Average session prices in U_1*Rep* and *BothRep* for the integrated firm are several times above the monopoly level which indicates that total foreclosure takes place.

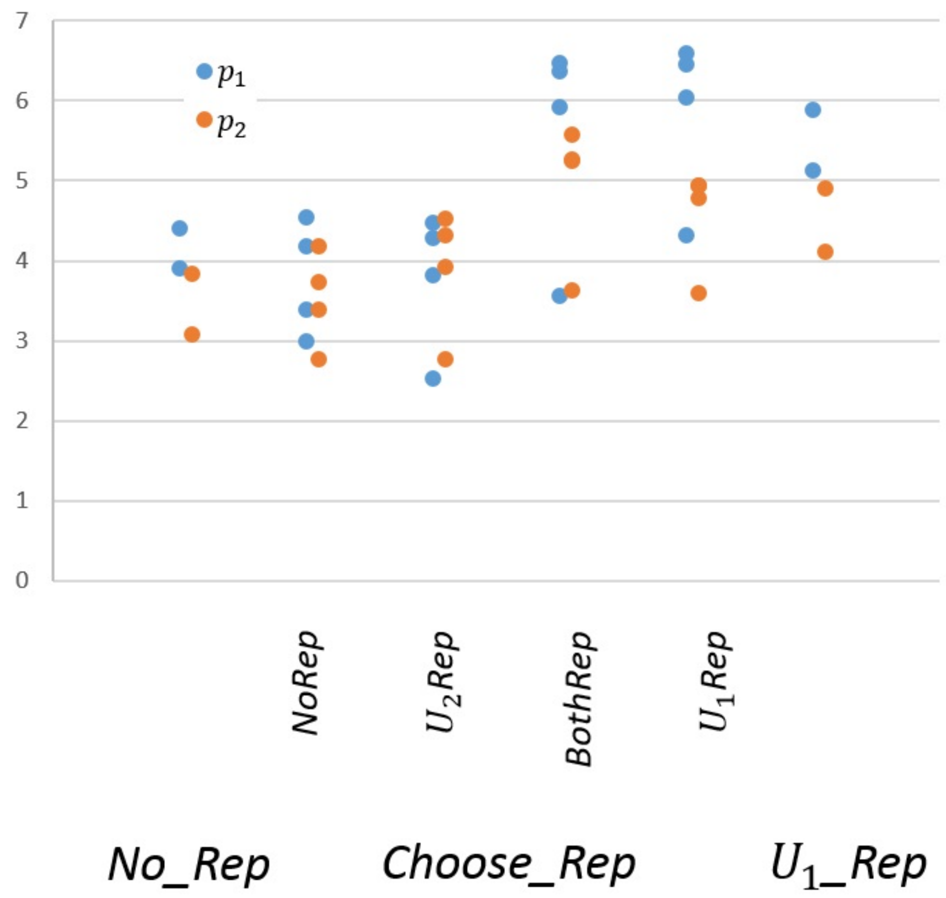

Figure 4.4: Average session prices

Results for partial input foreclosure are summarized in table 4.3. The fraction of prices $p_1 > p_2$ is always larger than fractions of $p_2 > p_1$ except in the outcome *U_2Rep* of treatment *Choose_Rep*. In the outcome *U_1Rep* of treatment *Choose_Rep*, partial foreclosure occurs in more than 75% of all observations. Again, reputation building of firm U_1 positively and significantly (regression (8), table 4.5) affects the percentage of $p_1 > p_2$. With self-imposed reputation building the difference is even more pronounced, however, the results of probit regression (8) of table 4.5 prove the difference between imposed and non-imposed reputation building to be insignificant. When firm U_2 builds a reputation, it leads to more undercutting of firm U_1. The effect is weakly significant at the 10% level (table 4.5, column (8)).

Overall, partial foreclosure occurs in 54.29% of observations in *Choose_Rep* (without differentiating outcomes). The fraction is closer to *No_Rep* (in which partial foreclosure is 5.58% less common) as compared to U_1_*Rep* (14.49% larger). Regression (7) in table 4.5 confirms partial foreclosure to be significantly less common in *Choose_Rep* than in U_1_*Rep* whereas the gap to *No_Rep* is insignificant. In *Choose_Rep* the positive effect of U_1's reputation building on $1_{p_1>p_2}$ is eliminated by the negative effect of U_2's reputation building. The mere possibility for both firms to build a reputation does not have a significant impact on partial foreclosure.

Total foreclosure might be the explanation for anticompetitive effects (supergames 2-4) without significantly more partial foreclosure in *Choose_Rep*.

4.7.4 Commitment and total foreclosure

In section 4.4 I used the conjecture that Stackelberg types ω^* exist, in order to refine the set of equilibria to the total foreclosure outcome $(\tilde{p}_1, \tilde{p}_2)$, $\tilde{p}_1 \in \tilde{S}_1$. In the data I find evidence in favor of the existence of Stackelberg types. In *No_Rep* 6.25% of the participants set $\tilde{p}_1$ in all 39 periods over the whole session (table 4.3). Without repeated interaction this behavior seems to contradict monetary incentives. However, it supports the approach in the predictions.[20]

The results suggest that the Stackelberg type is similarly common in each treatment (around 6%, table 4.3). However, in *Choose_Rep* the percentage is slightly larger (8.57%, table 4.3) and increasing to 11.43% if the first five periods of the first supergame are not considered. Surprisingly, the fraction of Stackelberg types is lowest in U_1_*Rep* (5.88%) whereas theory predicted that types $\omega \neq \omega^*$ would imitate the Stackelberg type in order to obtain almost Stackelberg payoffs.

One possible reason for missing treatment differences is that participants learn over time to mimic the Stackelberg type. The fraction of integrated firms which totally foreclose the input market during one supergame (in contrast to the whole session) is reported in table 4.3. Indeed, differences between treatments become apparent. In *No_Rep* the fraction is 9.38%, in U_1_*Rep* 22.06% and jointly for both outcomes in *Choose_Rep* the percentage equals 20%. Separating groups of U_1 firms with and without reputation building in *Choose_Rep* leads to fractions of 42.86% and 1.3%, respectively. The fractions are in any case larger than Stackelberg types who set $\tilde{p}_1$ the whole session and the degree of the increase crucially depends on the reputation building. In addition, the choice of reputation building of U_1 leads to total foreclosure during one supergame compared to U_1_*Rep* twice as often. It

[20] Several experimental studies which test reputation building in the lab change the payoff structure of Stackelberg types (Camerer and Weigelt, 1988, Neral and Ochs, 1992, Grosskopf and Sarin, 2010). On the other hand, experimental studies on reputation building in the trust game (for example Keser, 2002, Bohnet and Huck, 2004, Bolton, Katok and Ockenfels, 2004 as well as Bohnet, Huck, Harmgart and Tyran, 2005) do not change incentives exogenously.

seems that a substantial fraction imitates the Stackelberg type but not all of the participants follow this strategy throughout a whole supergame. Anecdotal evidence from a post-experimental survey suggests that some participants in the role of U_1 chose $\tilde{p}_1$ in several periods to gain trust followed by a "surprising" price cut.

Overall, the frequency of $\tilde{p}_1 \in \tilde{S}_1$ differs substantially between treatments (table 4.3). In *No_Rep* 18.11% of U_1s' price decisions equal $\tilde{p}_1$, this fraction almost triples in *U_1_Rep*, increasing by 30%. In *Choose_Rep* $\tilde{p}_1$ is observed in 5.64% and 5.63% of observations without vs. 62.21% and 60.7% with reputation building of firm U_1. In *Choose_Rep*, total foreclosure occurs ten times more often depending on reputation building of U_1.

Regression (10) in table 4.5 confirms the positive impact of U_1's reputation building (1% level), however, reputation building of firm U_2 has no significant effect on total input foreclosure. Whereas the coefficient of the treatment dummy *U_1_Rep* is insignificant, in *No_Rep* total foreclosure is significantly more common than in outcome *NoRep* (and *U_2Rep*) in *Choose_Rep* (1% level, column (10) of table 4.5). After five periods a slight downward trend is observed. Also, the probability of total foreclosure increases significantly in the second supergame compared to the first.

Concerning treatment effects, total foreclosure occurs overall in 30.11% of the observations in *Choose_Rep* which is somewhat in between the fractions observed in *No_Rep* and *U_1_Rep*. Differences turn out to be significant, the highest levels of total foreclosure are obtained in *U_1_Rep*, significantly lower fractions of $\tilde{p}_1 \in \tilde{S}_1$ in *Choose_Rep* (5% level, column (9), table 4.5) and least often in *No_Rep* (significant at the 1% level for both comparisons).

Result 5. Hypothesis 5 cannot be rejected. Although I find some evidence in favor of partial and total foreclosure, effects are insignificant in *No_Rep*.

Result 6. Hypothesis 6 cannot be rejected. Partial and total foreclosure in *U_1_Rep* are significantly (1% level) larger than in *No_Rep*.

Result 7. Hypothesis 7 cannot be fully rejected. Partial foreclosure in *Choose_Rep* is not significantly different from *No_Rep*. However, total foreclosure is significantly (1% level) more common as compared to *No_Rep* and occurs less frequently (significant at 5% level) than in *U_1_Rep*.

Having discussed total foreclosure of U_1, the response of U_2 is relevant for input prices. In table 4.8 the fractions of $\tilde{p}_2$ are displayed separately for periods 1-4. The difference in the first period between *No_Rep* and *U_1_Rep* are less pronounced than expected (9%, table 4.8). Whereas the fraction of $\tilde{p}_2$ in *No_Rep* decreases over time, the fractions in *U_1_Rep* increase. In period 3 the difference is more than 36%. In the first period of *Choose_Rep* the choice of reputation building determines whether fractions of $\tilde{p}_2$ are high (around 30%) or low (around 7%). With observable previous prices of the opponent, fractions increase whereas without reputation building of U_1 fractions decrease over time.

Table 4.8: Frequency of monopoly price $\tilde{p}_2$

$\tilde{p}_2$						
	No_Rep			*Choose_Rep*		*U_1_Rep*
Period		*NoRep*	*U_2Rep*	*BothRep*	*U_1Rep*	
1^{st}	20.31%	7.27%	36.36%	29.41%	6.50%	29.41%
2^{nd}	14.06%	5.17%	36.84%	30.00%	30.23%	33.82%
3^{rd}	7.81%	5.66%	37.50%	66.67%	35.42%	44.12%
4^{th}	7.81%	5.26%	25.00%	47.37%	56.82%	42.65%
Obs.	64			140		68
$\left(\tilde{p}_2 \vert p_1^i \in \tilde{S}_1,\ \forall i < t\right)$						
2^{nd}				33.33%	46.43%	53.66%
3^{rd}				63.64%	46.43%	85.71%
4^{th}				69.23%	87.50%	76.67%

Notes: In the upper part the frequency of $\tilde{p}_2$ in every treatment, periods 1-4, is displayed. In the part on the bottom, the sample is restricted to observations in which firm U_2 observes the price history of U_1. In addition, it is restricted to a subset of histories which only contain price $\tilde{p}_1 \in \tilde{S}_1$ in all previous periods $i < t$.

In table 4.8 I report fractions of $\tilde{p}_2$ restricted to a subset of observations. It contains only U_2 firms which observe previous prices of U_1 and, additionally, U_1 totally foreclosed the market, i.e. $\tilde{p}_1 \in \tilde{S}_1$, in every previous period. The fractions in the second period of the subset are not yet much different from the overall frequency of $\tilde{p}_2$ in the upper part. However, this clearly changes in period 4, where fractions are substantially larger.

Resulting from the large fraction of total market foreclosure with reputation building of firm U_1, monopolization changes. Without reputation building of U_1 downstream firm D_2 has to pay the monopoly price in less than 2.5% of all observations, with reputation building of U_1 more than 25% of all markets are monopolized (table 4.3).

To test how total foreclosure is achieved I consider the response of U_2 to particular histories of U_1. Using a fixed effects logit model I study the relevance of three different histories (the estimation is similar to Engle-Warnick and Slonim, 2006). The grim trigger strategy mentioned in the prediction might be an explanation for successful monopolization. Grim trigger begins with cooperation, i.e. U_2 chooses $\tilde{p}_2$ in the first period, and sets $\tilde{p}_2$ in every subsequent period whenever the history of U_1 contains only $\tilde{p}_1 \in \tilde{S}_1$. I define a variable $1_{tr} := 1_{t=1} + 1_{t>1} \prod_{i=1}^{t-1} 1_{p_1^i \in \tilde{S}_1}$ with t defined as the current period and p_1^i as price p_1 at period i. The second strategy included accounts for the concept in the prediction, i.e. every integrated firm needs to choose k times $\tilde{p}_1 \in \tilde{S}_1$ in order to convince U_2 that he will set $\tilde{p}_1 \in \tilde{S}_1$ in the current period. The equilibrium outcome equals $p_1 = \tilde{p}_1$ in each period and $p_2 \neq \tilde{p}_2$ the first k periods and in the following periods $p_2 = \tilde{p}_2$. A definition for this strategy is $1_{fl} := 1_{t>k}\ 1_{tr}$ for period t. Assuming that participants are not fully rational, a relevant strategy might be the myopic best response. In the first period U_2 sets $\tilde{p}_2$ and in the following she best responds to the action taken by the opponent in the previous period.[21] The corresponding variable is defined as $1_{mbr} := 1_{t=1} + 1_{t>1}\ 1_{p_1^{t-1} \in \tilde{S}_1}$.

Table 4.9 summarizes the results. Myopic best responses explain part of the observed behavior and is one of the strategies applied by the subjects (significant at 1% level). I interact strategy 1_{fl} with different phases in a supergame, i.e. periods $k+1$ to 5, periods 6 to 10 and periods 11 to 16. All interaction terms are positive and significant (at 5% level). However, I do not find a time trend after period k. That means, while the first k periods seem to have substantially lower fractions of $\tilde{p}_2$, later phases do not significantly differ from each other. For $k = 2$ basically all potential effects of 1_{tr} are captured by 1_{mbr} and one of the interaction terms of 1_{fl}. Consequently, the effect of 1_{tr} is insignificant with $k = 2$ but positive and significant at the 1% level for $k > 2$. I conclude that all three strategies explain choice of $\tilde{p}_2$ as a strategic reaction to a price history of U_1.

Result 8. Hypothesis 8 cannot be rejected. With total foreclosure of U_1 the strategy 1_{fl} has a highly significant impact on price choice of U_2. However, the well-known trigger strategy and myopic best responses are highly influential as well.

[21] The tit-for-tat strategy has a similar idea and turned out to be very successful in a prisoners' dilemma (Axelrod, 1984).

Table 4.9: Total foreclosure strategies

$1_{\tilde{p}_2}$	(11) $k=2$	(12) $k=3$	(13) $k=4$
1_{tr}	0.32	0.80***	1.13***
	(0.31)	(0.26)	(0.24)
1_{mbr}	1.94***	1.95***	1.95***
	(0.24)	(0.24)	(0.24)
1_{fl} $1_{Per.\ (k+1)-5}$	1.59***	1.42***	1.07**
	(0.36)	(0.37)	(0.48)
1_{fl} $1_{Per.\ 6-10}$	1.97***	1.49***	1.15***
	(0.41)	(0.37)	(0.35)
1_{fl} $1_{Per.\ 11-16}$	2.63***	2.19**	1.79**
	(0.91)	(0.90)	(0.88)
Obs.	1001	1001	1001

Notes: Each column (11)-(13) represents a fixed effects logit regression (subjects' choices during a whole supergame) on observations where the integrated firm builds a reputation. Note that because of the lack of variation 264 observations had to be deleted. These are, for example, observations of participants who never chose the monopoly price during a whole supergame. I include 1_{tr} as a dummy which is 1 in the first period and in every subsequent period if the history only contains prices above p^M, 1_{mbr} to control for myopic best responses, i.e. starts with 1 and equal to 1 if $1_{\tilde{p}_1}$ was chosen in the previous period. The parameter k varies between regressions (11)-(13) from $k = 2$, $k = 3$ to $k = 4$. Significantly different from 0 in a two-tailed test at the *10% level, **5% level, ***1% level.

4.8 Conclusion

This paper studies the impact of reputation building in a vertically related market with one integrated firm. OSS (1990) studied anticompetitive effects of vertical integration in a setting with two upstream and two downstream firms. Their assumption that the integrated firm can commit not to sell the input good to the downstream rival provoked a discussion among theorists. Bolton and Whinston (1991) summarize their concern: "There are two crucial steps in the OSS argument. The first is to show that as a result of a vertical merger, competition on the input market can be reduced. OSS establish this by assuming that the vertically integrated firm can commit to compete less fiercely on the input market. Exactly how this commitment is achieved is not explained. The second step is to show that by committing to compete less fiercely the integrated firm induces the other upstream firm to raise its input price and thus to raise the marginal cost of the unintegrated downstream sector."(p. 208).

In this paper, I address both concerns: First, I show theoretically and empirically that integrated firms can achieve commitment via reputation building. The integrated firm chooses significantly higher prices, undercuts the competitor significantly less often and withdraws from the input market at least three times more often when a reputation is built. Indeed, reputation building of the integrated firm seems to reduce competition upstream.

My results emphasize OSS's (1990) idea: they justify their assumption of commitment by stating that intuitively one would expect the integrated firm to behave differently. I find evidence in favor of this intuition. Indeed, when having the opportunity to build a reputation, the integrated firm decides more often to reveal previous price decisions. In addition, the mere opportunity to reveal previous prices leads to anticompetitive effects and significantly more total foreclosure.

Second, I find that in a substantial fraction of markets, the withdrawal of the integrated firm leads to monopoly prices. More than 25% of markets are monopolized as compared to not even 2.5% without reputation building of the integrated firm. In fact, I was able to identify three strategies which led to monopolization in my experiment. The non-integrated firms applied a myopic best response strategy, grim trigger and a strategy predicted by a theoretical model on reputation building by Fudenberg and Levine (1989).

The model of Fudenberg and Levine (1989) predicts that non-integrated firms need to be convinced via the reputation of the rival, that the integrated opponent will withdraw from the market. Strategies leading to cooperation in an infinitely repeated game often predict cooperation of both parties from the first period. I find evidence that the integrated firm has to invest in reputation building, i.e. total foreclosure which does not result in monopolization from the first period, in order to sustain the Stackelberg outcome.

I do not find much evidence for differences between imposed vs. non-imposed reputation building. However, for further research it might be insightful to test whether imposed reputation building of U_2 and two-sided imposed reputation in a infinitely repeated game support this conjecture. In addition, changing the default from no reputation to reputation building in *Choose_Rep* might support the hypothesis that the non-integrated firm would like to convey the impression of being a maverick. Also, more repetitions of supergames might manifest the impression that the integrated and non-integrated firms behave differently. Alternatively, one might obtain the same learning effect with less repetitions if histories of integrated firms would be revealed to other integrated firms as was done by Melis, Müller and Tremewan (2015).

References

Allain, Marie-Laure, Claire Chambolle, and Patrick Rey (2011) “Vertical Integration, Information and Foreclosure,” *Working Paper*.

Allain, Marie-Laure, Claire Chambolle, and Patrick Rey (2016) “Vertical Integration as a Source of Hold-up,” *The Review of Economic Studies*, 83(1): 1–25.

Allain, Marie-Laure, Claire Chambolle, and Patrick Rey (2015) “Vertical Integration as a Source of Hold-up: an Experiment,” *Working Paper EARIE 2016*.

Axelrod, Robert (1984) “The evolution of cooperation,” New York: Basic Books

Bohnet, Iris, Heike Harmgart, Steffen Huck and Jean-Robert Tyran (2005) “Learning Trust,” *Journal of the European Economic Association*, 3(2-3): 322–329.

Bohnet, Iris and Steffen Huck (2004) “Repetition and Reputation: Implications for Trust and Trustworthiness When Institutions Change,” *The American Economic Review*, 94(2): 362–366.

Bolton, Gary E., Elena Karok and Axel Ockenfels (2004) “How Effective Are Electronic Reputation Mechanisms? An Experimental Investigation,” *Management Science*, 50(11): 1587–1602.

Bolton, Patrick and Michael D. Whinston (1991) “The “Foreclosure” Effects of Vertical Mergers,” *Journal of Institutional and Theoretical Economics*, 147: 207 – 226.

Camerer, Colin and Keith Weigelt (1988) “Experimental Test of a Sequential Equilibrium Reputation Model,” *Econometrica*, 56: 1–36.

Chen, Yongmin (2001) “On Vertical Mergers and Their Competitive Effects,” *RAND Journal of Economics*, 32: 466–496.

Choi, Jay Pil, and Sang-Seung Yi (2000) “Vertical foreclosure with the choice of input specifications,” *RAND Journal of Economics*, 717–743.

Church, Jeffrey and Neil Gandal (2000) “System Competition, Vertical Merger and Foreclosure,” *Journal of Economics & Management Strategy*, 9(1): 25–51.

Dal Bó, Pedro (2005) “Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games,” *American Economic Review*, 95: 1591–1604

Dal Bó, Pedro and Guillaume R. Fréchette (2011) "The Evolution of Cooperation in Infinitely Repeated Fames: Experimental Evidence," *The American Economic Review*, 101(1): 411–429

Durham, Yvonne (2000) "An experimental examination of double marginalization and vertical relationships," *Journal of Economic Behavior and Organization*, 42(2): 207–229.

Engle-Warnick, Jim and Robert L. Slonim (2006) "Inferring repeated-game strategies form actions: evidence from trust game experiments," *Economic Theory*, 28: 603–632

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

Friedman, James W. (1971) "A Non-cooperative Equilibrium for Supergames," *The Review of Economic Studies*, 38(1): 1-12.

Fudenberg, Drew, David M. Kreps and Eric S. Maskin (1990) "Repeated Games with Long-run and Short-run Players," *Review of Economic Studies*, 57: 555–573.

Fudenberg, Drew and David K. Levine (1989) "Reputation and Equilibrium Selection in a Games with a Patient Player," *Econometrica*, 57: 759–778.

Gayle, Wayne-Roy, Robert C. Marshall, Leslie M. Marx, and Jean-François Richard (2011) "Coordinated Effects in the 2010 Horizontal Merger Guidelines," *Review of Industrial Organization* 39(1-2): 39–56.

Greiner, Ben (2015) "Subject pool recruitment procedures: organizing experiments with ORSEE," *Journal of the Economic Science Association*, 1(1): 114–125.

Grosskopf, Brit and Rajiv Sarin (2010) "Is Reputation Good or Bad?," *American Economic Review*, 100: 2187–2204.

Hart, Oliver and Jean Tirole (1990) "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity (Microeconomics)*, 205–276.

Huck, Steffen, Gabriele K. Lünser and Jean-Robert Tyran (2012) "Competition fosters trust" *Games and Economic Behavior*, 76: 195–209.

Kartal, Melis, Wieland Müller and James Tremewan (2015) "Gradualism in Repeated Games with Hidden Information: An Experimental Study," mimeo.

Keser, Claudia (2002) “Trust and reputation building in e-commerce,” Working paper, IBM, New York.

Kreps, David M. and Robert Wilson (1982) “Reputation and Imperfect Information,” *Journal of Economic Theory*, 27: 253–279.

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

Marshall, Robert C., Leslie Marx, Lily Samkharadze, and Maarten Pieter Schinkel (2016) “Monopolization conduct by cartels,” Working paper.

Mason, Charles F. and Owen R. Phillips (2000) “Vertical integration and collusive incentives: an experimental analysis” *International Journal of Industrial Organization*, 18: 471–496.

Milgrom, Paul and John Roberts (1982) “Predation, Reputation and Entry Deterence,” *Econometrica*, 50: 443–460.

Moellers, Claudia, Hans-Theo Normann and Christopher M. Snyder (2016) “Communication in Vertical Markets: Experimental Evidence,” NBER Working Paper No. 22219

Mouraviev, Igor and Patrick Rey (2011) “Collusion and Leadership,” *International Journal of Industrial Organization*, 29: 705–717.

Neral, John and Jack Ochs (1992) “The Sequential Equilibrium Theory of Reputation Building: A Further Test,” *Econometrica*, 60: 1151–1169.

Nocke, Volker and Lucy White (2007) “Do Vertical Mergers Facilitate Upstream Collusion?” *American Economic Review*, 97: 1321–1339.

Normann, Hans-Theo (2009) “Vertical integration, raising rivals’ costs and upstream collusion,” *European Economic Review*, 53(4): 461–480.

Normann, Hans-Theo (2011) “Vertical mergers, foreclosure and raising rivals’ costs - Experimental evidence,” *The Journal of Industrial Economics*, 59: 506 –527.

Ordover, Janusz A., Garth Saloner and Steven C. Salop (1990) “Equilibrium Vertical Foreclosure,” *American Economic Review*, 80: 127–142.

Ordover, Janusz A., Garth Saloner and Steven C. Salop (1992) "Equilibrium Vertical Foreclosure: Reply," *American Economic Review*, 82: 698–703.

Reiffen, David (1992) "Equilibrium Vertical Foreclosure: Comment," *American Economic Review*, 82: 694–697.

Salinger, Michael A. (1988) "Vertical mergers and market foreclosure," *The Quarterly Journal of Economics*, 345-356.

Selten, Reinhard (1978) "The chain store paradox," *Theory and Decision*, 9: 127–159.

Selten, Reinhard and Rolf Stoecker (1986) "End behavior in sequences of finite prisoner's dilemma supergames: A learning theory approach," *Journal of Economic Behavior and Organisation*, 7: 47–70.

Appendix: Instructions

Instructions in ***No_Rep***

In this experiment you will make decision in a fictitious market. Please take your time to read the instructions carefully. By understanding the instructions accurately and by well conceived decisions you can earn a substantial amount of money in the experiment. Your earnings will be paid out to you in cash at the end of the experiment.

Your role and your task in this experiment

Every participant will represent a firm in this experiment. There are two types of firms: firm 1 and firm 2. The computer will allocate half of the participants the role of firm 1 and the other participants the role of firm 2. Your role as firm 1 or firm 2 stays the same over the whole experiment. You will get to know at the beginning of the experiment whether you are firm 1 or firm 2.

One firm 1 and one firm 2 will meet in a market for a fictitious good. This market will be named market A. Firm 1 - but not firm 2 - operates in a second market, market B.

The computer will randomly determine in every period one firm 1 and one firm 2 who meet in market A. Which firms are selected is completely random, i. e. there is no connection between the participant with whom you were matched last period and the participant who is assigned to you by lot in this period.

Your task is the same in each period, irrespectively of whether you are firm 1 or firm 2: You have to decide upon a price. This price can be an integer between 1 and 9. The profit, which you can make via your price choice, can be calculated as follows.

Profit calculation

In market A holds that

- the firm with the lowest price makes the profit which you can learn from the following table (profits in points).
- the firm with the highest of both prices does not get any profit in this period.
- in case both choose the same price, the profit in the table will be separated

Here are two examples for market A.

i If you set a price of 7 and the other firm a price of 4, you will get a profit of zero and the other firm a profit of 81 points in market A.

Price	Profit in market A in points	Profit in market B in points (only firm 1)
1	39	66
2	54	75
3	69	84
4	81	96
5	90	105
6	99	132
7	90	159
8	72	180
9	51	198

ii If both of you choose a price of 3, then both firms get a profit of $69/2 = 34.5$ points in market A.

In market B only firm 1 gets a profit - the profit in the column with “market B”. Like already in market A the lower of both prices decides about the profit, regardless of whether firm 1 or firm 2 has chosen it. Here an example for market B. If firm 1 sets a price of 7 and firm 4 a price of 4, then firm 1 will get a profit of 96 in market B.

Let us consider both markets together, firm 1 gets a profit from market A plus the one from market B. In our example this would be 96 points. Firm 2 gets no profit in market B, but in market A and therefore 81 points.

Procedure and End of the experiment

In every period both firms choose a price. At the end of each period every firm gets the following feedback: the own price, the price of the other firm and the own profit gained in this period.

There will be 4 rounds. In each round it is randomly determined how many periods take place. Before the beginning of each period it is drawn by lot whether the period will take place. With a chance of 90% the period takes place, with a chance of 10% the round stops.

After 4 rounds you get your whole profit. For each 300 points you get 1 Euro cash. In addition, you get 4 Euros.

Additional Instructions for U_1_*Rep*

Information

As already mentioned, the firm you will meet in market A is randomly determined in each period. However, firm 2 gets information about firm 1. In contrast, firm 2 does not get the information about firm 2.

Firm 2 can observe prices of firm 1 in the previous periods of the current round while setting the price. That means, there is a table displayed on the screen with all previous periods and the corresponding price of firm 1. Firm 2 can only observe prices of the current market participant firm 1. (There is an example for the table presented below.)

In contrast, firm 1 only knows the own pricing decisions.

Additional Instructions for *Choose_Rep*

Information

As already mentioned, the firm you will meet in market A is randomly determined in each period. However, both firms might get additional information about the other firm.

At the beginning of each round you can decide whether the other firms can observe your previous prices. If you decide to reveal your price information, the other firm will observe prices from every previous period in the current round while setting the price. That means, there is a table displayed on the screen with all previous periods and your corresponding price choices. (There is an example for the table presented below.)

Chapter 5

Conclusion

In my thesis, I studied in three market experiments the impact of different parameters and conditions on competition. Various market structures and theoretical frameworks were considered and ceteris paribus changes were analyzed.

In the second chapter, we found that bilateral communication between an upstream monopolist and either of the two downstream firms does not entirely solve the commitment problem. The offered quantity is lower, however, the market quantity is unaffected. Communication between all the three firms, involving competitors downstream, clearly monopolized the market. The gains from the monopolization are exclusively obtained by the retailers. In a hybrid of both treatments where both channels are available, we obtain results in between the other treatments.

In the third chapter we analyzed the impact of a change in consumer search cost on market outcomes. In a dynamic interaction between two firms, we obtained that decreasing the search cost did not change average prices. On the other hand, prices vary more and consumers search more frequently with lower search cost. Because of more search with lower cost, the consumer surplus does not change significantly with the level of search cost. Communication between sellers leads to a higher price level which does not depend on the level of search cost.

In the fourth chapter, I studied a vertically related market with one integrated and one non-integrated firm. In one-shot interactions, upstream firms got the opportunity to build a reputation. I found that vertically integrated upstream firms build a reputation more often than non-integrated firms. Integrated firms use reputation building to overcome the commitment problem, i.e. withdrawal from the input market. Because the non-integrated firms gain monopoly power, the integrated firms gain downstream via the raising-rivals'-cost effect. In my experiment, I obtained input foreclosure resulting in monopolization ten times more often, when the integrated firm built a reputation.

Coming back to the example of the market for fuel in the first chapter. In chapter 3 we obtain the same outcome as the intervention "Markttransparenzstelle für Kraftstoffe" of the cartel office. Lowering search cost has no effect on the price level but increases the variation of prices. In chapter 2 we study a upstream monopolist facing downstream competition and how communication affects the commitment problem. In some local markets, non-integrated refineries might have market power (for example the refinery in Heide owned by the Klesch Group). Communication between refineries and gas stations cannot be avoided and might be used to overcome the commitment problem. Finally, in chapter 4 I analyze an integrated firm which faces a non-integrated upstream competitor. The integrated firm might search for opportunities to commit to foreclose the input market. Potentially, this might be a reason why independent gas stations complain about inflated input costs.

Eidesstattliche Versicherung

Ich, Claudia Möllers, versichere an Eides statt, dass die vorliegende Dissertation von mir selbstständig, und ohne unzulässige fremde Hilfe, unter Beachtung der "Grundsätze zur Sicherung guter wissenschaftlicher Praxis an der Heinrich-Heine-Universität Düsseldorf" erstellt worden ist.

Düsseldorf, 31. October 2016

Unterschrift