



Munich Personal RePEc Archive

The Impact of Forward Trading on Tacit Collusion: Experimental Evidence

Jens Schubert

University of Tennessee, Knoxville

1 January 2013

Online at <https://mpra.ub.uni-muenchen.de/43768/>

MPRA Paper No. 43768, posted 14 January 2013 01:52 UTC

The Impact of Forward Trading on Tacit Collusion: Experimental Evidence

Jens Schubert

January 2013

Department of Economics, University of Tennessee, Knoxville

Email: jens.schubert@utk.edu

Abstract

This article reports the results of a laboratory experiment that examines the strategic effect of forward contracts on market power in infinitely repeated duopolies. Two competing effects motivate the experimental design. Allaz and Vila (1993) argue that forward markets act like additional competitors in that they increase quantity competition among firms. Conversely, Liski and Montero (2006) argue that forward contracting can facilitate collusive outcomes by enabling firms to soften competition. The experiment provides a first simultaneous test of these rival effects. Contrary to previous experimental studies, the results do not support the quantity-competition effect. Further, the findings provide evidence in support of the collusive hypothesis.

JEL Classification: C72, C91, D43, L13, Q49

Keywords: Cournot oligopoly, Collusion, Experiments, Forward markets, Electricity markets

1 Introduction

Antitrust authorities and researchers have a profound interest in the factors that determine the likelihood of collusion. Extensive theoretical and empirical work focuses on the determinants of firms' coordinated efforts to achieve profits in excess of the competitive outcome. Most empirical studies are experimental as strategic field data is difficult to obtain and identification of specific factors can be challenging due to interactions and unobservables. Controlled laboratory experiments, however, allow targeted tests in market environments that satisfy the assumptions of the underlying model of interest. This article studies the effect of forward contracts on tacit collusion in duopolies with quantity-setting firms.

A forward contract is an agreement between two parties to buy or sell a fixed quantity at a specified time in the future at a price agreed upon today. Historically, forward contracts have played an important role in commodity markets and more recently in financial asset markets. Forward contracts have also become increasingly important in electricity wholesale markets. Forward trading is a prevalent instrument in hedging risk: forwards contracts allow buyers and sellers to potentially offset unfavorable price movements in the spot market by shifting risk to less risk-averse market participants.

However, Allaz (1992) and Allaz and Vila (1993) hypothesize that even in the absence of risk and uncertainty, forward markets can emerge and will lead to more market efficiency. The underlying intuition is that quantity-setting firms will sell some of their production forward to improve their position relative to competitors in the spot market. In the spot market, firms will then compete over the residual demand. Firms will find themselves in a prisoner's dilemma type situation: Although firms would be jointly better off by avoiding selling in advance, it is beneficial to an individual firm to do so (Stackelberg leadership advantage). As a result, each firm produces more than in the absence of forward markets, which reduces their market power. Following the Western U.S. energy crisis of 2000 and 2001, this pro-competitive prediction led to suggestions to remove restrictions on forward contracts with the goal of limiting the ability of electricity generators to exercise market power.

The pro-competitive hypothesis assumes that oligopolists only compete with each other a limited number of times. Competing theories relax this assumption and derive hypotheses that challenge the pro-competitive argument. According to the Friedman (Folk) theorem, there are multiple equilibria in an infinitely repeated setting: Ferreira (2003) shows that if firms are able to sell their production in infinitely

many forward phases prior to the spot market, forward contracts can have an anti-competitive effect. Mahenc and Salanié (2004) show that when firms compete over prices of slightly differentiated products, firms will take long positions in the forward market which will lead to higher prices and thus higher profits compared to oligopolistic markets without forward markets. Liski and Montero (2006) study the effect of forward contracts in an infinitely repeated oligopoly; the authors demonstrate that forward markets enable quantity-setting firms to soften competition more than they could in the absence of forward markets. In particular, when firms repeatedly interact both in forward and spot markets, the existence of forward markets yields a wider range of discount rates which allow for the collusive equilibrium. The gains from deviating from the collusive path are never greater than the gains in an infinitely repeated oligopoly without forward markets, and the profits from the ensuing sanctioning equilibrium (Allaz and Vila equilibrium) are less than the profits from the sanctioning equilibrium in the absence of forward contracts (standard Cournot equilibrium). The focus of this article is to investigate whether forward sales yield strategic effects in an infinitely repeated Cournot setting. In particular, we test the collusive hypothesis of Liski and Montero against the pro-competitive hypothesis of Allaz and Vila in a controlled laboratory experiment¹.

Previous experimental studies on the two-phase forward model of Allaz and Vila report results that support the pro-competitive prediction. In a finitely repeated two-phase Cournot setting with fixed matching, Le Coq and Orzen (2006) find that a forward trading phase leads to increases in market efficiency. However, the pro-competitive effect is less pronounced than predicted by theory. Van Koten and Ortman (2011) use a similar experimental design with producers' cost functions that resemble electricity generators more closely. Their findings also suggest that introduction of a forward market lowers market prices through increased aggregate output. Brandts et al. (2008) report that both in settings with quantity competition and supply-function competition, forward markets lead to reductions in market prices and thus yield greater market efficiency. The authors also use a finitely repeated protocol with fixed matching. Ferreira et al. (2010) test the strategic effects of forward markets in quantity competition settings with finitely and infinitely many forward trading phases with random re-matching of subjects after each round. In the finitely repeated treatments, their findings support the competitive hypothesis of Allaz and Vila. Further, the authors do not find evidence of collusive outcomes

¹Note that repeated play of the Allaz and Vila stage-game strategy is one of many subgame perfect equilibrium strategies in an infinitely repeated setting.

in the treatments with infinitely many forward markets.

Liski and Montero (2006) predict that several strategies can yield the collusive equilibrium in the presence of forward markets. However, previous experimental studies use a pricing rule that significantly reduces the set of possible collusive strategies: The forward pricing rule dictates a forward price that is less than or equal to the spot price - the forward price is equal to the spot price if and only if firms either play the pro-competitive strategy or jointly refrain from selling forward. This pricing rule also introduces uncertainty about price differences between forward and spot market phases which leads to interaction of strategic and risk hedging motives. To allow for multiple collusive equilibria and to eliminate price uncertainty effects, the experimental design in this article strictly imposes forward-spot price parity. We achieve this by restricting firms' quantity choices to a discrete choice set. The possible choices reflect different pure strategies in the quantity-setting stage-game. The set of limited strategies also increases the likelihood of collusive outcomes (see Holt (1995)).

We compare the market outcomes of a two-phase duopoly with forward trading to the results of a standard, one-phase duopoly. Specifically, we examine differences in collusive behavior between these two treatments. The collusive hypothesis predicts that multiple collusive equilibria can emerge in the two-phase duopoly. We investigate stage-game outcomes in the spot market phase (conditional on forward phase outcomes) to test for differences in forward trading between colluding and non-colluding firms. To compare the competitive effect of market entry to the effect of forward markets, we report the differences in market efficiency between a three-firm oligopoly and the two-phase duopoly.

The main result of this article is that, contrary to previous experimental findings, introducing a forward market in a duopoly may not increase market efficiency. The pro-competitive hypothesis predicts that the effect of a forward market is equivalent to squaring the number of firms. However, we find that one additional competitor significantly limits market power in a duopoly whereas a forward market does not. Further, we provide evidence that allowing firms to trade forward can facilitate collusion as predicted by Liski and Montero's collusive hypothesis.

The organization of the remainder of the article is as follows. Section 2 presents the predictions of the pro-competitive and collusive theories and derives the hypotheses which guide the experimental design. Section 3 describes the experimental design and procedures. Section 4 presents the results of the article, and Section 5 discusses the main findings.

2 Theoretical Framework and Hypotheses

We will first derive the pro-competitive predictions of the stage-game and then contrast them to the collusive predictions of the infinitely repeated game. Notice that, according to the Folk theorem, repeated play of the stage-game equilibrium strategy is a subgame perfect equilibrium strategy in the infinitely repeated game. In the following derivation, we only consider a single forward market opening prior to the spot market (for a detailed derivation with multiple forward market openings, see Allaz and Vila (1993); Ferreira (2003)).

Competitive Framework

Standard Cournot Game

First, consider a single phase Cournot game with J firms that compete over quantity. Without loss of generality, assume symmetric firms with zero production cost. For simplicity, let the inverse demand function be given by

$$p(q) = \alpha - \sum_{j=1}^J q_j \quad (2.1)$$

where q_j denotes firm j 's output. The single period, unique Nash equilibrium is given by

$$q_j^c = \frac{\alpha}{J+1}; \pi_j^c = \frac{\alpha^2}{(J+1)^2}; \forall j; p^c = \frac{\alpha}{J+1} \quad (2.2)$$

where π_j denotes firm j 's profits. Backward induction implies that the same one-shot game predictions hold in a finitely repeated game.

Two-Phase Cournot Game

Now consider a two-phase Cournot game in which a forward market is followed by a standard Cournot game spot market. The good is physically bought and sold in the spot market. In the first phase (forward market), firms can sell some or all of their production for delivery in the second phase (spot market). At the end of the first phase, firms observe the forward market outcome. In the second phase, firms compete in quantity over the residual demand. At the end of the second phase, firms observe the spot market production and total production of their competitor(s), the market price, p , and profit π_j . For a detailed derivation of the two-phase equilibrium, see Allaz and Vila (1993) and Bushnell (2007).

The existence of arbitrage traders in the market will yield forward-spot price parity, $p^f = p^s = p(q)$ (where p^f (p^s) denote the forward-phase (spot-phase) price, respectively). Arbitrage traders will compete in prices over firms' short forward positions and will try to sell them at a profit to buyers in the spot market. In equilibrium, any price differences between the two phases will disappear. Another way to think about forward-spot price parity is that buyers have perfect foresight and are therefore indifferent between buying in the forward or spot market.

The game can be solved using backward induction. Let $f(s)$ denote total units sold in the first (second) phase, respectively. Although the demand has perfect foresight, the theoretical model assumes that firms treat their first-phase profits as being unaffected by their second-phase production decisions. Given forward positions, firm j 's profit maximization problem in the spot market game can be written as

$$\max_{s_j} p(s_j, s_{-j}, f) s_j; \quad j = 1, \dots, J \quad (2.3)$$

with corresponding first order condition

$$0 = p(\cdot) + \frac{\partial p(\cdot)}{\partial s_j} s_j; \quad \forall j \quad (2.4)$$

With an inverse demand function as given in equation 2.1, the first order condition is

$$0 = \alpha - f - \sum_{k=1}^J s_k - s_j; \quad \forall j$$

Simultaneously solving the J best response functions gives firm j 's optimal second phase production:

$$s_j(f) = \frac{\alpha - f}{J + 1}; \quad \forall j \quad (2.5)$$

which is a best response to any arbitrary level of forward sales commitment. To obtain the first phase equilibrium, the second phase best response functions are nested in the first phase objective function:

$$\max_{f_j} p \left(f_j, \sum_{k \neq j} f_k, \sum_{k=1}^J s_k(f) \right) (f_j + s_j(f)); \quad \forall j \quad (2.6)$$

with corresponding first order condition

$$0 = p(\cdot) \left(1 + \frac{\partial s_j}{\partial f_j}\right) + \frac{\partial p}{\partial q} \left(1 + \sum_{k=1}^J \frac{\partial s_k}{\partial f_j}\right) (f_j + s_j); \quad \forall j \quad (2.7)$$

$$= \frac{J-1}{J+1} (\alpha - f) - f_j \quad (2.8)$$

Simultaneously solving the J first order conditions and imposing symmetry gives

$$f_j = \frac{J-1}{J^2+1} \alpha; \quad \forall j \quad (2.9)$$

The two-phase Cournot equilibrium can be summarized as

$$f_j^{fs} = \frac{J-1}{J^2+1} \alpha; \quad s_j^{fs} = \frac{1}{J^2+1} \alpha; \quad q_j^{fs} = \frac{J}{J^2+1} \alpha; \quad \pi_j^{fs} = \frac{J}{(J^2+1)^2} \alpha^2; \quad \forall j \quad (2.10)$$

with equilibrium price

$$p^{fs} = \frac{\alpha}{J^2+1} \quad (2.11)$$

Note that the Cournot equilibrium output of a J -firm, two-phase oligopoly equals the output of a J^2 -firm, single-phase oligopoly: $q^{fs}(J) = q^c(J^2)$. To summarize, in a finitely repeated setting, the existence of a single forward market increases quantity competition between firms which increases market efficiency. The following two hypotheses capture the predictions of the finitely repeated two-phase game:

Hypothesis 1. *Oligopoly markets with a forward market phase yield higher output (lower prices) on average than oligopoly markets with a spot market phase only.*

Hypothesis 2. *The market outcome (total output, price, and profit) of a J -firm, two-phase oligopoly is equivalent to the market outcome of a J^2 -firm, single phase oligopoly.*

Tacit Collusion

Next, consider an infinitely repeated Cournot game where the same firms compete repeatedly with each other. According to the Friedman theorem, all firms jointly producing the monopoly quantity is a subgame perfect equilibrium strategy for sufficiently high discount rates δ . We assume that when firms play the cooperative subgame strategy, they split the monopoly output equally. The stage-game collusive

outcome can be summarized as

$$q_j^{tc} = \frac{\alpha}{2J}; \pi_j^{tc} = \frac{\alpha^2}{4J}; p^{tc} = \frac{\alpha}{2}; \quad \forall j \quad (2.12)$$

Comparison of the different equilibrium profit predictions yields $\pi_j^{fs} < \pi_j^c < \pi_j^{tc}$.

Standard Cournot Game

In deriving the cooperative, subgame perfect equilibrium predictions, we assume that firms will cooperate as long as they observe the other firms playing the cooperation strategy. Once a firm cheats, firms will play the stage-game Nash equilibrium strategy thereafter.

In the single phase Cournot game, firm j 's one-period incentive to deviate from the collusive strategy (cheating) is

$$\max_{q_j} \left(\alpha - (J-1) \frac{\alpha}{2J} - q_j \right) q_j \quad (2.13)$$

Firm j 's production and profit and the resulting market price are:

$$q_j^d = \frac{(J+1)}{4J} \alpha; \pi_j^d = \frac{(J+1)^2}{16J^2} \alpha^2; p^d = \frac{J+1}{4J} \alpha \quad (2.14)$$

The cooperative strategy q^{tc} will be a subgame perfect equilibrium strategy, if the following condition holds

$$\pi_j^d + \frac{\delta}{(1-\delta)} \pi_j^c < \frac{1}{(1-\delta)} \pi_j^{tc}, \quad \delta \in [0, 1] \quad (2.15)$$

The implied critical discount factor for the existence of the subgame perfect equilibrium can be calculated as

$$\frac{(J^2 - 1)^2}{(J + 1)^4 - 16J^2} < \underline{\delta}(J) \quad (2.16)$$

Two-Phase Cournot Game

In the following derivation, we generalize Liski and Montero's framework to an oligopolistic setting with J firms. For simplicity, we restrict firms' positions in the forward market to short positions only (see Liski and Montero (2006) for details on firms' holding long positions in the forward market). In the two-phase Cournot game, several collusive strategies support the subgame perfect equilibrium. Assume

that, in the cooperative subgame, firm j sells $f_j^{tc} = \lambda_j \cdot q_j^{tc}$, $\lambda_j \in [0, 1]$ units in the first phase and $s_j^{tc} = (1 - \lambda_j) \cdot q_j^{tc}$ units in the second phase². The model assumes that firms treat their forward market profit as being unaffected by their production decision in the spot market. This implies that firms' incentives to deviate from the collusive path are smaller in the spot-phase stage-game if they have forward sales positions. Therefore, the gains of deviating from the collusive path will never be greater than the profit from deviating in the single phase stage-game. Further, deviation is more costly in the two-phase game as the sanctioning path is the two-phase stage-game Cournot equilibrium. These two effects result in a strictly lower critical discount factor that supports the collusive outcome. Firm j 's one-period incentive to deviate from the collusive strategy in the spot market is³

$$\max_{s_j} \left(\alpha - (J - 1) \frac{\alpha}{2J} - \lambda_j \frac{\alpha}{2J} - s_j \right) s_j \quad (2.17)$$

where $\lambda_j \alpha / 2J = \lambda_j q_j^{tc}$ denotes firm j 's forward sales expressed in terms of the collusive amount. Firm j 's production and profit and the resulting market price are:

$$f_j = \lambda_j \frac{\alpha}{2J}; \quad s_j^d = \frac{(J + 1 - \lambda_j)}{4J} \alpha; \quad \tilde{\pi}_j^d = \pi_j^d - \frac{\lambda_j^2}{16J^2} \alpha^2; \quad \tilde{p}^d = \frac{(J + 1 - \lambda_j)}{4J} \alpha \quad (2.18)$$

Note that the one period profit from cheating in the two-phase game is always less than or equal to the single phase deviating profit. Strategy $\{s_j^{tc}, f_j^{tc}\}$ denotes a subgame perfect equilibrium strategy if the following inequality is satisfied

$$\tilde{\pi}_j^d + \frac{\delta}{(1 - \delta)} \pi_j^{fs} < \frac{1}{(1 - \delta)} \pi_j^{tc} \quad (2.19)$$

The left-hand side in equation 2.19 is strictly less than the left-hand side in equation 2.15. The critical discount factor is therefore strictly lower than the critical discount factor in the single phase game:

$$\frac{\left[(J + 1)^2 - \lambda_j^2 - 4J \right] (J^2 + 1)^2}{\left((J + 1)^2 - \lambda_j^2 \right) (J^2 + 1)^2 - 16J^3} < \tilde{\delta}(\lambda_j, J) < \underline{\delta}(J), \quad \forall \lambda_j \in [0, 1].$$

Note that $\tilde{\delta}(\lambda_j, J)$ is decreasing in λ_j . Table 1 summarizes the subgame equilibria predictions.

²Firms' forward positions do not have to be symmetric ($\lambda_i \neq \lambda_j$) in order for the collusive subgame perfect equilibrium to exist.

³It is never profitable to cheat in the forward market (see Liski and Montero (2006)).

The following main hypotheses guide the experimental design. These hypotheses reflect the cooperative subgame perfect equilibrium predictions in the infinitely repeated, two-phase Cournot game.

Hypothesis 3. *In an infinitely repeated setting, two-phase oligopoly markets yield lower output (higher prices) on average than single phase oligopolies.*

Hypothesis 4. *Firms can sustain the cooperative subgame equilibrium across both phases (forward and spot market) in infinitely-repeated oligopolies.*

Hypothesis 5. *In infinitely repeated two-phase oligopolies, firms that sell forward are less likely to defect than firms that have no forward sales position.*

[insert Table 1 here]

3 Experimental Design

The objective of the experimental design is to test the strategic effect of forward sales in an infinitely repeated setting. In order to test for the existence of cooperative subgame equilibria, it is important to create a market environment in the laboratory that gives the predicted collusive equilibria the best chance of occurrence. The following main findings from previous oligopoly experiments contributed to our design (see Engel (2007) for a comprehensive meta-analysis of oligopoly experiments). First, the larger the number of firms, the smaller the observed degree of collusion (see also Huck et al. (2004)). Second, experienced subjects tend to collude more than inexperienced subjects, i.e. learning plays an important role (Huck et al. (1999)). Third, the better subjects are informed, the more likely they play a cooperative strategy. Lastly, if subjects play against human buyers, “collusion rates plummet” (Engel (2007)).

Our experiment compares a standard duopoly (C2 treatment) to a two-phase duopoly with a single forward and a single spot market phase (FS2 treatment). A third, standard three-firm oligopoly treatment (C3 treatment) allows us to analyze differences between the effect of adding one additional competitor to the effect of a single forward market. Adding one additional competitor serves as a lower bound on the effect of increased competition from additional firms.

Strategy Design

The main design challenge is to implement forward-spot price parity. The underlying theoretical models assume that demand has perfect foresight. However, in the

laboratory, it is impossible to perfectly predict the decisions that subjects make in a stage-game. Previous experimental studies that test the pro-competitive prediction (Le Coq and Orzen (2006); Ferreira et al. (2010); Van Koten and Ortmann (2011)) use a pricing rule which dictates the forward price to equal the spot price if and only if all firms play the pro-competitive strategy. This pricing rule introduces price uncertainty and it eliminates all cooperative subgame perfect strategies in the forward market as the calculated forward price is always less than the collusive price. Brandts et al. (2008) let human buyers compete over firms' forward market positions in a Bertrand game; however, this significantly reduces the likelihood of collusive outcomes.

Our design automates demand using a computer program. We implement forward-spot price parity by restricting subjects' quantity choices to a discrete choice set. The market price is not determined until after the end of the spot phase. This implies that subjects do not observe their forward profits before making their spot phase decisions⁴. Instead, the quantity choices in the spot phase of the stage-game are calculated as if the spot phase choices do not affect the profits in the forward market. The set of limited strategies also decreases unintended effects of inexperienced subjects and importantly increases the likelihood of collusive outcomes (Holt (1995)).

In the forward phase of the FS2 treatment, subjects have the following two choices: either selling zero units or selling the stage-game equilibrium forward quantity as predicted by the pro-competitive theory. Notice that the forward quantity is less than the collusive amount, which admits a collusive strategy across forward and spot phases. In the spot market (C2 and FS2 treatments), the possible choices are zero, collusive, Cournot, defecting, and punishing output, which reflect pure strategies. In the FS2 treatment, the quantity choices are calculated based on the residual demand (total demand less forward sales).

We provide subjects with a detailed payoff table that lists all possible outcomes. Subjects are knowledgeable of their own and their competitors' profit in any feasible stage-game outcome. (A copy of the instructions can be found in Appendix D.) Further, in all treatments, subjects can perfectly monitor the choices made by their competitor(s).

⁴Subjects only observe the forward quantity commitments.

Demand Specification

The demand side is automated and subjects have zero production costs ($\gamma = 0$). The inverse demand is given by

$$p_{m,t} = \max \{120 - q_{m,t}, 0\} \quad (3.1)$$

where $q_{m,t}$ denotes the total units sold in market m in round t . As stated above, we strictly impose forward-spot price parity in the FS2 treatment: $p_{m,t}^s = p_{m,t}^f = p_{m,t} = 120 - f_{m,t} - s_{m,t}$, where $f_{m,t}$ and $s_{m,t}$ respectively denote total units sold in the forward and spot phase. This assures that the conditions of the game in the experiment are as close to theory as possible without affecting the testable hypotheses. Importantly, subjects receive the same price for any units sold in either forward or spot phase. In each round, a subject's total profit is calculated as the product of their individual total production times the market price.

[insert Table 2 here]

Table 2 lists the different strategy choices by treatment. In both duopoly treatments, there are five output choices in the spot phase stage-game. In the C3 treatment however, the defecting and punishing output quantities are equivalent, $q_j = 40$. Therefore, subjects could only choose from a set of four different quantities in the C3 treatment. In the FS2 treatment, subjects can play the collusive strategy in two different ways: either selling zero units forward and 30 units in the spot phase or selling 24 units forward and 6 units in the spot phase, respectively. This yields four different collusive subgame perfect equilibria in the FS2 treatment. Table 3 contrasts the collusive, Cournot, and defecting outcome predictions for all three treatments. Notice that selling forward makes the defecting strategy less tempting in the spot phase of the stage-game in the FS2 treatment.

[insert Table 3 here]

The implied critical discount factors in the experiment are $\underline{\delta} = 9/17$ in the C2 treatment, $\underline{\delta}(\lambda_j = 0.8) = 1/9$, $\delta(\lambda_j = 0) = 25/97$ in the FS2 treatment, and $\underline{\delta} = 4/7$ in the C3 treatment. The punishing strategy in the stage-game allows subjects to play a more severe grim strategy than just the Nash-reverting strategy. This implies lower critical discount factors of $\underline{\delta} = 1/9$ in the C2 treatment, $\underline{\delta}(\lambda_j = 0.8) = 9/209$, $\delta(\lambda_j = 0) = 1/9$ in the FS2 treatment, and $\underline{\delta} = 1/4$ in the C3 treatment.

Termination Rule

Our design implements a repeated game with uncertain end, which, according to the Friedman theorem, allows for several subgame equilibria to exist (Friedman (1971)). Subjects compete with the same other subject(s) for many rounds (fixed matching), but they do not know the exact number of rounds until the end of the experimental session. Normann and Wallace (2012) show that the termination rule in prisoner dilemma games does not significantly affect cooperation but may influence how cooperation can be sustained over time and its influence on end of game effects (see also Selten and Stoecker (1986)). Further, the authors find that the number of rounds significantly increases cooperation rates. The two-phase duopoly game is a complicated market mechanism; therefore, we refrain from using a stochastic termination rule with continuation probability to avoid unnecessary confusion of subjects' comprehension of the mechanism. Initially, we considered two different termination rules: known-end (subjects learn the exact number of rounds at the beginning of the session) and unknown-end. Specifically, we employed the known-end termination rule in one C2 and one C3 session. In comparing outcomes, we find no statistically significant difference between the unknown-end and known-end C3 sessions. In testing for end of game effects, we find that, on average, subjects chose higher outputs (more competitive strategies) in the final round of the known-end C2 session. We therefore exclude the final round observations in the known-end C2 session from the analysis. Appendix A shows the statistical analysis of the termination rules and end of game effects in detail.

Procedures

The data was collected in seven experimental sessions at the University of Tennessee, Knoxville in the Spring and Summer semesters in 2012. A total of 144 undergraduate student subjects participated in the sessions. Each subject participated in one session only. Each session consisted of 27 rounds⁵ and lasted between one hour and one hour 30 minutes (the FS2 sessions lasted longer than the Cournot sessions due to the two-phase format). Subjects earned \$23 on average.

At the beginning of each session, subjects were randomly and anonymously matched with one (two) other subject(s). Subjects were informed that they will interact with the same other subject(s) for several rounds. A monitor read the experimental instructions and explained the computer program to participants. The

⁵The two known-end termination rule sessions consisted of 25 rounds each.

monitor thoroughly described the payoff table that accompanied the instructions. To verify that subjects understood how their earnings were calculated, the computer program asked each subject four practice questions before the start of the experiment. The computer program also displayed a payoff table in each decision round that listed all feasible sales combinations along with payoffs. In the second phase of the FS2 treatment, the computer program updated this payoff table conditional on the sales decisions in the first phase.

In each round, each participant had to choose an output amount from a list on the computer screen. After all participants submitted their sales decisions, the computer program determined the total sales units and price in each market. (At the end of the first phase in the FS2 treatment, subjects only observed the forward sales of their competitor and total forward sales in their market.) At the end of each round, each subject learned the total output of the other subject(s) in their market, the total market output, the resulting market price, and their profit for that round. The computer program summarized and updated the market outcomes from previous rounds in the form of a table that was displayed on the computer screen at the time subjects submitted their decisions. (Appendix E shows screen shots of the FS2 treatment.) All treatments were programmed in z-Tree (Fischbacher (2007)).

4 Experimental Results

[insert Figure C.1 here]

Market Efficiency

First, we analyze the results in terms of total output and market efficiency. Figure C.1 plots the average total output in each round by treatment. Horizontal lines at 60, 80, 90, and 96 denote respectively the collusive, standard duopoly stage-game equilibrium, three-firm stage-game equilibrium, and two-phase duopoly stage-game equilibrium output. The figure shows that the average two-phase duopoly output (black circles) is not different from the average standard duopoly output (light gray diamonds). Further, the average total output in the two-phase duopoly is far less than the predicted two-phase stage-game equilibrium quantity of 96 units. In both duopoly treatments, the average total output fluctuates at or below the standard stage-game equilibrium amount of 80 units. The aggregate three-firm output (gray triangles) oscillates around the stage-game equilibrium amount of 90 units. The graph also indicates that total output in both duopoly treatments is less than in the

three-firm treatment. Figures C.4, C.5, and C.6 in Appendix C show the total output by individual markets. These graphs indicate that outcomes are heterogeneous across markets. Some markets maintain either the collusive or the standard stage-game Cournot output for the majority of the rounds. In other markets, total output is characterized by high volatility.

[insert Table 4 here]

Table 4 lists the average total output (by phase), prices, seller profits, and market efficiency across all rounds by treatment. The average total output in the two-phase duopoly treatment is not statistically different from average total sales in the standard duopoly treatment. The total quantity in the three-firm treatment is larger on average than the average total output in either duopoly treatment. In all three treatments, average total output is significantly greater than the collusive output (60 units). In the two-phase duopoly treatment, average forward sales (20.70 units) are significantly less than 48 units and spot sales are significantly greater than 48 units. Subjects sell significantly more units in the spot phase than in the forward phase (see Table 5 for detailed test statistics).

[insert Table 5 here]

Observations are likely dependent upon each other within a single market (group of matched subjects) and across time. We account for these potential inter-dependencies using a standard OLS model with robust standard errors clustered at the market level. The model tests whether total market output and market efficiency differ across the three treatments. Table 6 presents the estimation results. Specification 2 allows for a cubic time trend. Specification 3 allows for the cubic time trend to differ between the two-firm treatments and for a quadratic time trend in the C3 treatment. In all three specifications, total output and efficiency in the three-firm treatment are significantly greater than in either two-firm treatment. However, the coefficient estimate on C2 is not significantly different from zero. Statistical significance of the coefficients on the time trend terms indicates that the chosen specifications capture the observed fluctuations across time well. In particular, both two-firm treatments exhibit oscillatory patterns. Market efficiency in the three-firm treatment is increasing at a decreasing rate over time.

[insert Table 6 here]

These findings indicate that there are no significant differences in total output and efficiency between the two duopoly treatments. We conclude that

Result 1. *In an infinitely repeated setting, market efficiency in two-phase duopolies is not different from market efficiency in single phase duopolies.*

The following two findings are possible explanations of this result. First, on average, neither firm committed to any forward sales in 38% of individual two-phase duopoly stage-games, which means that subjects faced the single phase Cournot stage-game in more than one third of individual stage-games. In 20 out of 24 markets, both firms avoided forward sales in at least one round. Both firms sold in the forward phase in only 24% of all individual market outcomes. Second, as outlined in section 2, several collusive equilibria can be sustained in the two-phase duopoly game. The following discussion examines the latter conjecture by analyzing strategy choices in the spot phase of the stage-game.

Strategy Choices

[insert Figure C.2 here]

Figure C.2 contrasts the distributions of chosen (stage-game) strategies in the duopoly treatments. Standard normality tests suggest that both distributions have a positive skew. Subjects chose the collusive and Cournot (stage-game) strategies most frequently in both treatments. Whereas the difference between the collusive and Cournot strategies is not significant in either treatment (Wilcoxon matched pairs, $z_{FS2} = 1.23$, $p_{FS2} = 0.22$, $z_{C2} = 0.74$, $p_{C2} = 0.46$), all other differences between strategies are significant at the 1% level within each treatment. In both two-firm treatments, comparing the frequency of chosen strategies results in the following order: collude, Cournot > defect > punish > zero. The chart in Figure C.2 also shows that subjects chose the collusive strategy more frequently in the FS2 treatment than in the C2 treatment. Further, sellers chose the defective and competitive strategies less frequently in the FS2 treatment than in the C2 treatment. However, these differences are not significant (all strategies jointly: Kruskal-Wallis test, $\chi^2 = 0.92$, $p = 0.34$; individual strategies: test of proportions with p-values ranging from 0.66 to 0.96). Note that decisions in the experimental markets are very heterogeneous. (Figures C.7, C.8, and C.9 in Appendix C show the distribution of chosen strategies by market.)

[insert Table 7 here]

As a robustness check, we jointly test whether there are differences in distribution of chosen (stage-game) strategies in a multinomial logit model with standard errors clustered at the market level. Table 7 reports the estimation results. The coefficient estimate on the C2 indicator variable is not significantly different from zero for all strategies, which confirms that there are no significant differences in distribution between the C2 and FS2 treatment. An interesting result is that sellers chose the defective strategy less often in later rounds relative to the collusive strategy. Also, in both treatments, sellers chose the zero output strategy (dominated strategy) less often in later rounds.

[insert Figure C.3 here]

Figure C.3 shows the distribution of strategies in the C3 treatment. Sellers chose the Cournot strategy most frequently. However, the difference between collusive and Cournot strategies is not significant (Wilcoxon matched pairs, $z = -1.58$, $p = 0.11$). Although subjects did not choose the collusive strategy significantly more often than the defective strategy ($z = 0.13$, $p = 0.90$), they chose the Cournot strategy significantly more often than the defective strategy ($z = 1.97$, $p = 0.05$). We do not test for differences in strategy distribution between the C3 and the duopoly treatments as the choice set in the C3 treatment consists of four choices only.

[insert Table 8 here]

Next, we focus on the two-phase duopoly treatment only. To analyze how forward sales affect the output decisions in the spot phase, we test for differences in chosen strategies in the spot phase of the stage-game conditional on the outcome in the forward phase. Table 8 reports the estimation results of a multinomial logit model that allows for a linear time trend. The binary variables 'Self Sold Forward', 'Competitor Sold Forward' and 'Self-Competitor' uniquely describe the four possible forward market outcomes. There are no observable differences between sellers choosing either the collusive or the Cournot strategy conditional on the forward market outcome. However importantly, subjects were less likely to choose the defective strategy if they sold in the forward market phase. Based on the marginal effect, firms that hold forward positions are 15.7% less likely to defect in the spot market relative to the collusive strategy. This finding indicates that a firm can commit to the collusive strategy more decidedly by selling forward. The following two results summarize the findings of the two-phase duopoly treatment:

Result 2. *In duopolies with a single forward market opening, the collusive outcome can be sustained across both phases.*

Result 3. *A single forward market opening can soften competition in duopoly markets.*

Allowing subjects to play the punishing strategy (i.e. firms in a market produce an equal share of the maximum demand) leads to behavioral phenomena such as negative reciprocity. Some subjects play a collusive strategy in early rounds. Their competitors, however, play the defective strategy repeatedly early in the supergame⁶. Subjects then reciprocate by choosing the punishing output in later rounds. We observe these patterns in several markets in both duopoly treatments. This behavior indicates that the punishing strategy is a viable grim strategy. The main results are unaffected by this behavioral effect.

5 Conclusion and Discussion

In this article, we have studied the strategic effect of forward sales on market efficiency and firms' output choices in infinitely repeated experimental duopoly markets. Although there is considerable heterogeneity in market outcomes within each treatment, we obtained the following robust results. First, a forward market does not act like additional competitors in an infinitely repeated setting. Second, several collusive equilibria can be maintained in the presence of forward markets. Although we did not discover any differences in market efficiency between duopoly markets with and without forward sales, we found evidence that forward sales commitments can strengthen collusion as the defective strategy becomes less profitable in the spot-market. In our experiment, the collusive effect outweighed the increased quantity competition effect.

The experimental design in this article differs from previous experimental studies that test the strategic motive of forward contracts. We create a market environment in the laboratory that increases the likelihood of observing collusive outcomes. To facilitate cooperation, we use a fixed matching protocol with an unknown-end termination rule. Our design permits subjects to play several cooperative subgame perfect strategies and we impose strict forward-spot price parity to eliminate possible risk hedging motives. We achieve forward-spot price parity by restricting subjects'

⁶A supergame refers to several consecutive rounds of the same stage-game in a group of matched sellers

quantity choices to a discrete choice set. These design features support four collusive equilibria in the forward market duopoly treatment that were not supported in previous experimental work.

The results of this article can assist antitrust authorities in mitigating market power in oligopolies that are characterized by few firms that interact repeatedly. A good example is the wholesale electricity industry: few sellers, homogeneous products that cannot be stored economically at a large scale⁷, and sound forward markets. This article confirms that merely requiring electricity generators to sell forward, with the intent to limit their market power, can have the opposite effect as forward sales can strengthen collusive outcomes. Without strict regulation, two ways to mitigate market power in oligopolies are incentivizing entry and introducing forward markets. The results of this article provide evidence that incentivizing entry can be a superior market mechanism to forward markets.

⁷Two different spot markets are therefore independent markets and standard storage-based arbitrage arguments do not apply.

References

- Allaz, B. (1992). Oligopoly, uncertainty and strategic forward transactions. *International Journal of Industrial Organization* 10(2), 297–308.
- Allaz, B. and J.-L. Vila (1993). Cournot Competition, Forward Markets and Efficiency. *Journal of Economic Theory* 59(1), 1–16.
- Brandts, J., P. Pezanis-Christou, and A. Schram (2008). Competition with Forward Contracts: A Laboratory Analysis Motivated by Electricity Market Design. *The Economic Journal* 118, 192–214.
- Bushnell, J. (2007). Oligopoly Equilibria in Electricity Contract Markets. *Journal of Regulatory Economics* 32, 225–245.
- Engel, C. (2007, June). How Much Collusion? A Meta-Analysis of Oligopoly Experiments. *Journal of Competition Law and Economics* 3(4), 491–549.
- Ferreira, J. (2003, January). Strategic interaction between futures and spot markets. *Journal of Economic Theory* 108(1), 141–151.
- Ferreira, J. L., P. Kujal, and S. Rassenti (2010). Multiple Openings of Forward Markets: Experimental Evidence.
- Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-Made Economic Experiments - Experimenter’s Manual. *Experimental Economics* 10, 171–178.
- Friedman, J. W. (1971). A Non-cooperative Equilibrium for Supergames. *The Review of Economic Studies* 38(1), 1–12.
- Holt, C. (1995). Industrial Organization: A Survey of Laboratory Research. In J. H. Kagel and A. E. Roth (Eds.), *Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.
- Huck, S., H.-T. Normann, and J. Oechssler (1999). Learning in Cournot Oligopoly - An Experiment. *The Economic Journal* 109, 80–95.
- Huck, S., H.-T. Normann, and J. Oechssler (2004, April). Two are few and four are many: number effects in experimental oligopolies. *Journal of Economic Behavior & Organization* 53(4), 435–446.
- Le Coq, C. and H. Orzen (2006). Do Forward Markets Enhance Competition? Experimental Evidence. *Journal of Economic Behavior & Organization* 61(3), 415–431.

- Liski, M. and J.-P. Montero (2006). Forward Trading and Collusion in Oligopoly. *Journal of Economic Theory* 131(184), 212–230.
- Mahenc, P. and F. Salanié (2004). Softening competition through forward trading. *Journal of Economic Theory* 116(2), 282–293.
- Normann, H.-T. and B. Wallace (2012). The Impact of the Termination Rule on Cooperation in a Prisoner’s Dilemma Experiment. *International Journal of Game Theory* 41(3), 707–718.
- Selten, R. and R. Stoecker (1986, March). End Behavior in Sequences of Finite Prisoner’s Dilemma Supergames - A Learning Theory Approach. *Journal of Economic Behavior & Organization* 7(1), 47–70.
- Van Koten, S. and A. Ortmann (2011). Structural versus Behavioral Measures in the Deregulation of Electricity Markets: An Experimental Investigation Guided by Theory and Policy Concerns.

A Termination Rule

We count each market as a single observation to account for possible interdependence of observations within a single market. The average chosen strategy in the known-end C2 treatment is lower (1.74) than the average strategy in the unknown-end treatment (2.08). This difference is not significant (Wilcoxon rank sum test, $z = -1.04$, $p = 0.30$). There is no observable difference in average chosen strategies between the two different termination rules in the C3 treatment (Wilcoxon rank sum test, $z = -0.17$, $p = 0.87$).

Next, we test for changes in subjects' decisions at the end of the game by comparing chosen strategies in a short period at the end of the game to chosen strategies in 10 prior rounds. Specifically, in the known- (unknown-) end termination treatments, we compare the average chosen strategies in a market in rounds 14-23 (16-25) to the average strategies in rounds 24-25 (26-27), respectively. Counting each market as a single observation, we find no statistically significant difference in the distribution of chosen strategies between the unknown-end C2 treatment, the known- and unknown-end C3 treatments, and the FS2 treatment (matched-pairs Wilcoxon signed-rank test with p-values ranging from 0.21 - 0.97). Although not significant, note that average chosen strategies are lower in the last two rounds in both C3 termination rule treatments. In the known-end C2 treatment, however, the average chosen strategy is significantly greater in the last two rounds compared to the 10 rounds prior (matched pairs Wilcoxon, $z = -2.86$, $p = 0.004$). Comparing average chosen strategies in rounds 15-24 (17-26) to average strategies in the final round yields similar results.

Average chosen strategies in the known-end C2 treatment are lower than average strategies in the unknown-end C2 treatment. Therefore, we test whether average chosen strategies in round 24 of the known-end C2 treatment differ from average strategies in rounds 18-27 in the unknown-end C2 treatment: we find no significant difference in average chosen strategies (Wilcoxon rank sum test, $z = -0.56$, $p = 0.58$). Based on these results, we drop the final round from the known-end C2 data and pool unknown- and known-end termination rule sessions.

B Tables

Table 1: Theoretical Market Outcome Predictions

	f	s	q	p	π_j
Single Phase Subgame	-	$\frac{J}{J+1}\alpha$	$\frac{J}{J+1}\alpha$	$\frac{1}{J+1}\alpha$	$\frac{1}{(J+1)^2}\alpha^2$
Two-Phase Subgame	$\frac{J(J-1)}{J^2+1}\alpha$	$\frac{J}{J^2+1}\alpha$	$\frac{J^2}{J^2+1}\alpha$	$\frac{1}{J^2+1}\alpha$	$\frac{J}{(J^2+1)^2}\alpha^2$
Cooperative Subgame	$\frac{\lambda}{2}\alpha$	$\frac{(1-\lambda)}{2}\alpha$	$\frac{1}{2}\alpha$	$\frac{1}{2}\alpha$	$\frac{1}{4J}\alpha^2$

Note: In the single phase stage-game: $\lambda = 0$. In the two-phase stage-game: $\lambda \in [0, 1]$.

Table 2: Sales Choices by Phase, by Treatment

	f_j	s_j
C2	-	{0, 30, 40, 45, 60}
FS2	{0, 24}	{0, $(30 - f_j)$, $(120 - f) / 3$, $(90 - f_j) / 2$, $(120 - f) / 2$ }
C3	-	{0, 20, 30, 40}

Table 3: Collusive, Cournot, and Defecting Outcome Predictions by Treatment

		f_j	s_j	q_j	f	s	q	p	π_j	Efficiency
Collude	C2	-	30	30	-	60	60	60	1,800	75%
	FS2	{0, 24}	{30, 6}	30	{0, 24, 48}	{60, 36, 12}	60	60	1,800	75%
	C3	-	20	20	-	60	60	60	1,200	75%
Cournot	C2	-	40	40	-	80	80	40	1,600	89%
	FS2	24	24	48	48	48	96	24	1,152	96%
	C3	-	30	30	-	90	90	30	900	94%
Defect	C2	-	45	45	-	75	75	45	2,025	86%
	FS2	0	45	45	{0, 24}	{75, 51}	75	45	2,025	86%
	FS2	24	33	57	{24, 48}	{63, 39}	87	33	1,881	92%
	C3	-	40	40	-	80	80	40	1,600	89%

Note: The defecting outcomes are calculated based on the assumption that the other firm(s) play the collusive strategy.

Table 4: Summary Statistics, Average Market Outcomes by Treatment

	f_j	s_j	q_j	f	s	q	p	π_j	Efficiency
C2	-	37.61 (5.59)	37.61 (5.59)	-	75.33 (14.94)	75.21 (10.75)	44.79 (10.75)	1,572.79 (204.72)	84.52% (6.52%)
FS2	10.35 (7.99)	28.08 (8.00)	38.43 (4.99)	20.70 (13.42)	56.17 (13.67)	76.87 (8.87)	43.13 (8.87)	1,522.62 (184.67)	85.20% (5.34%)
C3	-	29.75 (3.75)	29.75 (3.75)	-	89.22 (14.51)	89.24 (8.42)	30.76 (8.42)	845.43 (175.75)	91.98% (3.92%)

Note: Standard deviations in parentheses. Each market counts as a single observation to control for possible correlation within a market.

Table 5: z-Statistics of Wilcoxon rank-sum and matched-pairs tests (H_0 : row variable = column variable, H_a : row variable \neq column variable)

	48	60	80	90	96	q_{C2}	s_{FS2}	f_{FS2}	q_{FS2}	q_{C3}
q_{C2}	-	4.04*** 0.00	-1.66* 0.10	-	-	-	-	-	-0.54 0.59	-3.87*** 0.00
s_{FS2}	2.54** 0.01	-1.37 0.17	-	-	-	-	-	4.09*** 0.00	-	-
f_{FS2}	-4.29*** 0.00	-	-	-	-	-	-4.09*** 0.00	-	-	-
q_{FS2}	-	4.14*** 0.00	-1.16 0.25	-	-4.29*** 0.00	0.54 0.59	-	-	-	-3.99*** 0.00
q_{C3}	-	3.52*** 0.00	-	-0.49 0.62	-	3.87*** 0.00	-	-	3.99*** 0.00	-

Note: p-values given beneath. Significance at the 1%, 5%, and 10% level is denoted by ***, **, and *, respectively. Each market enters the tests as a single observation to control for possible correlation within a market.

Table 6: Effect of Treatment on Total Output and Efficiency

	Output			Efficiency		
	(1)	(2)	(3)	(1)	(2)	(3)
Constant	76.87*** (1.79)	72.11*** (2.38)	73.47*** (3.13)	85.20%*** (1.08%)	82.49%*** (1.54%)	83.77%*** (2.01%)
C2	-1.54 (2.80)	-1.50 (2.80)	-4.74 (4.34)	-0.61% (1.69%)	-0.58% (1.69%)	-3.62% (2.98%)
C3	12.35*** (2.75)	12.37*** (2.75)	8.73** (4.10)	6.76%*** (1.45%)	6.77%*** (1.45%)	5.18%** (2.52%)
Round	-	1.30** (0.53)	-	-	0.82%*** (0.31%)	-
Round ²	-	-0.09** (0.04)	-	-	-0.06%** (0.02%)	-
Round ³	-	2.0E-03* (1.1E-03)	-	-	0.0014%** (0.0006%)	-
C2·Round	-	-	2.09*** (0.77)	-	-	1.39%*** (0.52%)
C2·Round ²	-	-	-0.17** (0.06)	-	-	-0.11%*** (0.04%)
C2·Round ³	-	-	0.004** (0.002)	-	-	0.003%** (0.001%)
FS2·Round	-	-	1.54 (0.95)	-	-	0.77% (0.53%)
FS2·Round ²	-	-	-0.14* (0.08)	-	-	-0.07% (0.04%)
FS2·Round ³	-	-	0.003* (0.002)	-	-	0.002%* (0.001%)
C3·Round	-	-	1.02*** (0.38)	-	-	0.44%* (0.21%)
C3·Round ²	-	-	-0.03** (0.01)	-	-	-0.012%* (0.006%)
F	13.65	7.26	4.36	15.11	7.45	4.30
R ²	0.12	0.13	0.13	0.11	0.12	0.12

Note: $N = 1,682$ (64 markets with 24 (25) to 27 observations per market). FS2 is the control group. Standard errors in parantheses. Significance of coefficient estimates at the 1%, 5%, and 10% level is denoted by ***, **, and *, respectively.

Table 7: Effect of Type of Two-Firm Treatment on Strategy

	Zero		Cournot		Defect		Punish	
	Coefficient	Marginal Effect	Coefficient	Marginal Effect	Coefficient	Marginal Effect	Coefficient	Marginal Effect
Constant	-1.93*** (0.33)		-0.31 (0.28)		-0.79*** (0.29)		-2.13*** (0.37)	
C2	-1.09** (0.53)	-1.48%*** (0.63%)	0.09 (0.41)	0.93% (7.39%)	0.17 (0.44)	1.64% (4.41%)	0.34 (0.52)	1.72% (2.75%)
Round	-0.081*** (0.022)	-0.09%*** (0.03%)	-2.1E-03 (1.2E-02)	0.12% (0.25%)	-2.6E-02* (1.5E-02)	-0.32%** (0.15%)	-3.2E-05 (2.1E-02)	0.04% (0.12%)

Log-Likelihood = -3,143.62; Wald $\chi^2 = 35.82$; $N = 2,520$ (48 markets)

Note: FS2 is the control group. Base strategy is collude. The multinomial logit model estimates a set of coefficients for each strategy other than the base strategy. Coefficient estimates for different strategies are shown across columns. Standard errors in parantheses. Significance of coefficient estimates at the 1%, 5%, and 10% level is denoted by ***, **, and *, respectively.

Table 8: Effect of Forward Market Outcome on Spot Market Strategy Choice

	Zero		Cournot		Defect		Punish	
	Coefficient	Marginal Effect	Coefficient	Marginal Effect	Coefficient	Marginal Effect	Coefficient	Marginal Effect
Constant	-3.00*** (0.55)		-0.33 (0.46)		-0.56 (0.44)		-3.20*** (0.60)	
Self Sold Forward	0.75 (0.76)	1.62% (1.37%)	0.02 (0.57)	3.21% (10.13%)	-1.40** (0.57)	-15.70%*** (4.56%)	1.37** (0.67)	7.46%** (3.57%)
Competitor Sold Forward	0.62 (0.70)	0.90% (1.12%)	-0.15 (0.54)	-9.35% (9.50%)	0.53 (0.52)	4.87% (3.83%)	1.97*** (0.56)	10.29%*** (2.63%)
Self-Competitor	0.52 (0.93)	0.68% (1.57%)	0.56 (0.81)	14.04% (15.27%)	0.11 (0.82)	-0.79% (5.77%)	-1.75** (0.86)	-5.98%*** (1.70%)
Round	-0.068** (0.028)	-0.11%*** (0.04%)	-0.007 (0.016)	0.03% (0.33%)	-0.031 (0.020)	-0.28% (0.19%)	-0.013 (0.030)	-0.02% (0.12%)

Log-Likelihood = -1,555.16; Wald $\chi^2 = 193.13$; $N = 1,296$ (24 markets)

Note: Control group is no forward sales. Base strategy is collude. The multinomial logit model estimates a set of coefficients for each strategy other than the base strategy. Coefficient estimates for different strategies are shown across columns. Standard errors in parantheses. Significance of coefficient estimates at the 1% and 5% level is denoted by *** and **, respectively.

C Figures

Figure C.1: Average Total Output per Round, All Treatments




Figure C.2: Percentage of Strategies by Two-Firm Treatment

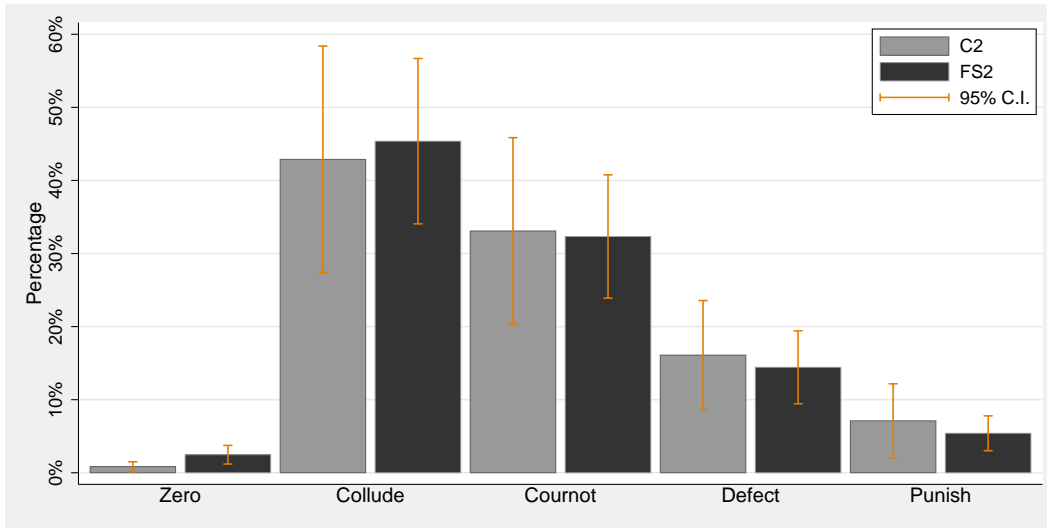


Figure C.3: Percentage of Strategies in C3 Treatment

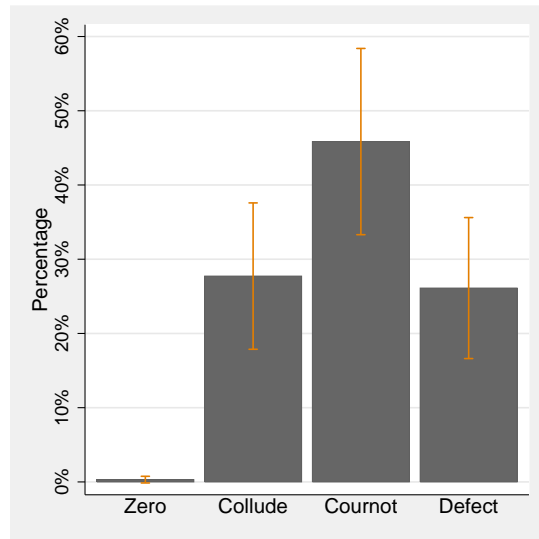
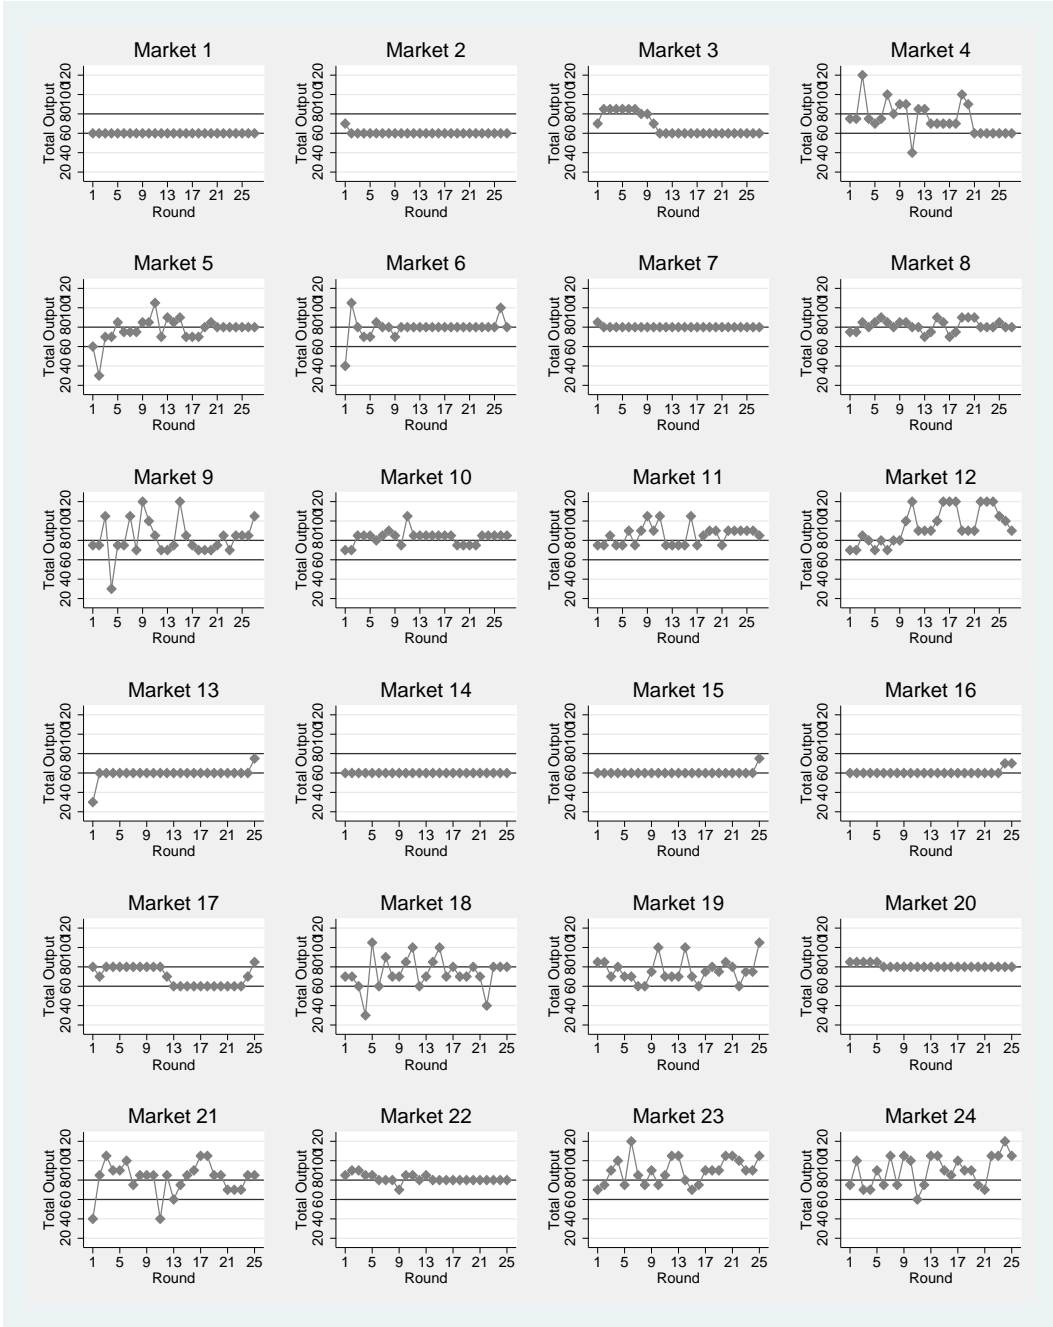
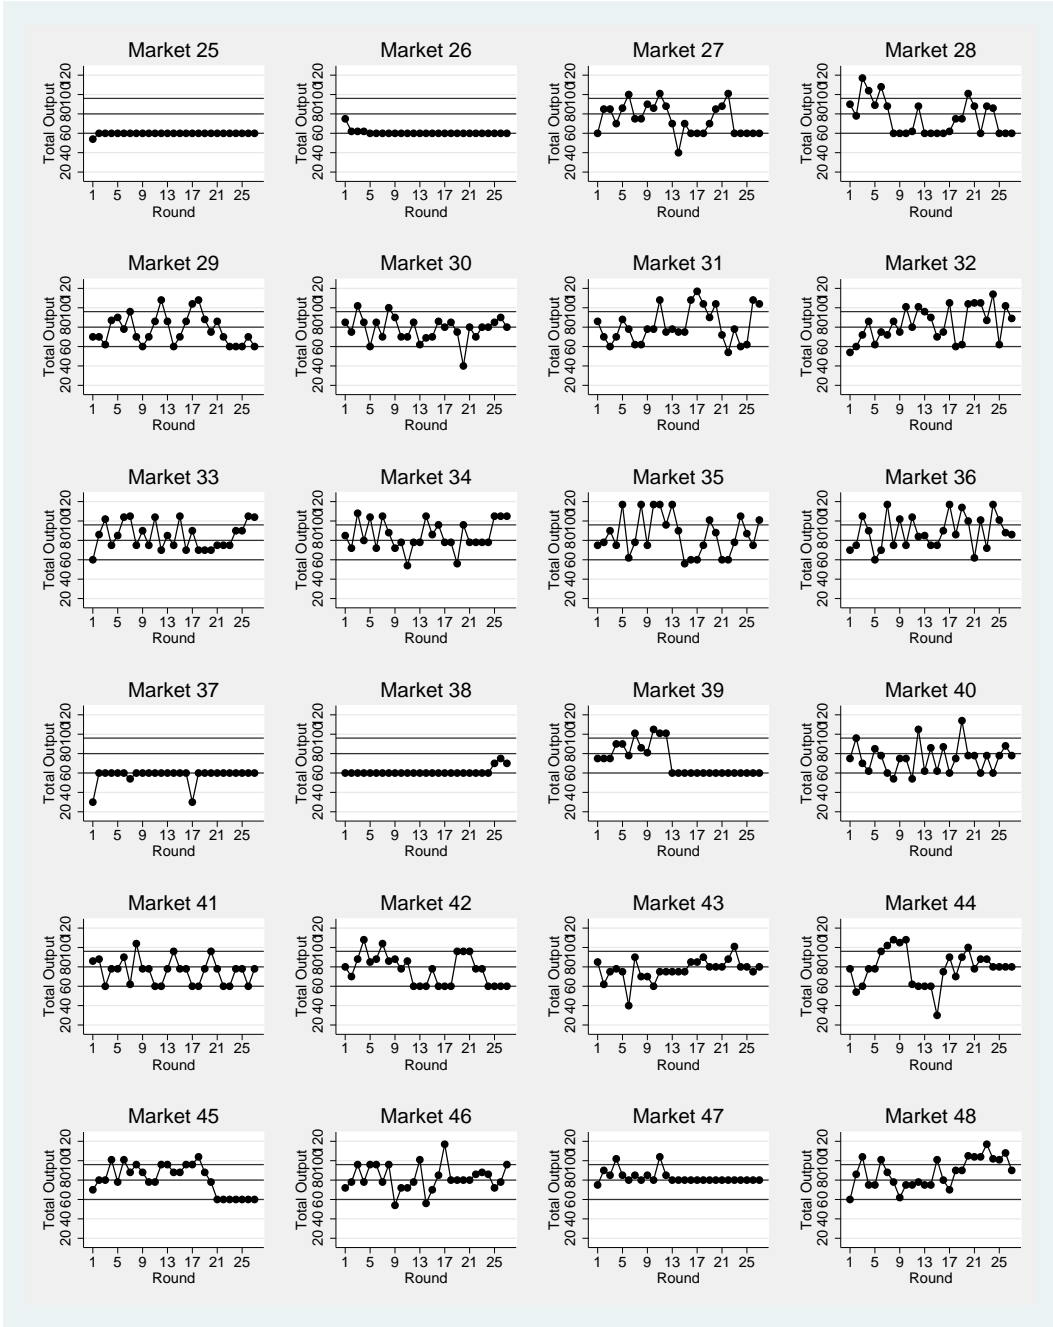


Figure C.4: Total Output per Round, C2 Treatment, All Markets




Note: Markets 1-12 (13-24) in unknown- (known-) end termination rule treatment.

Figure C.5: Total Output per Round, FS2 Treatment, All Markets



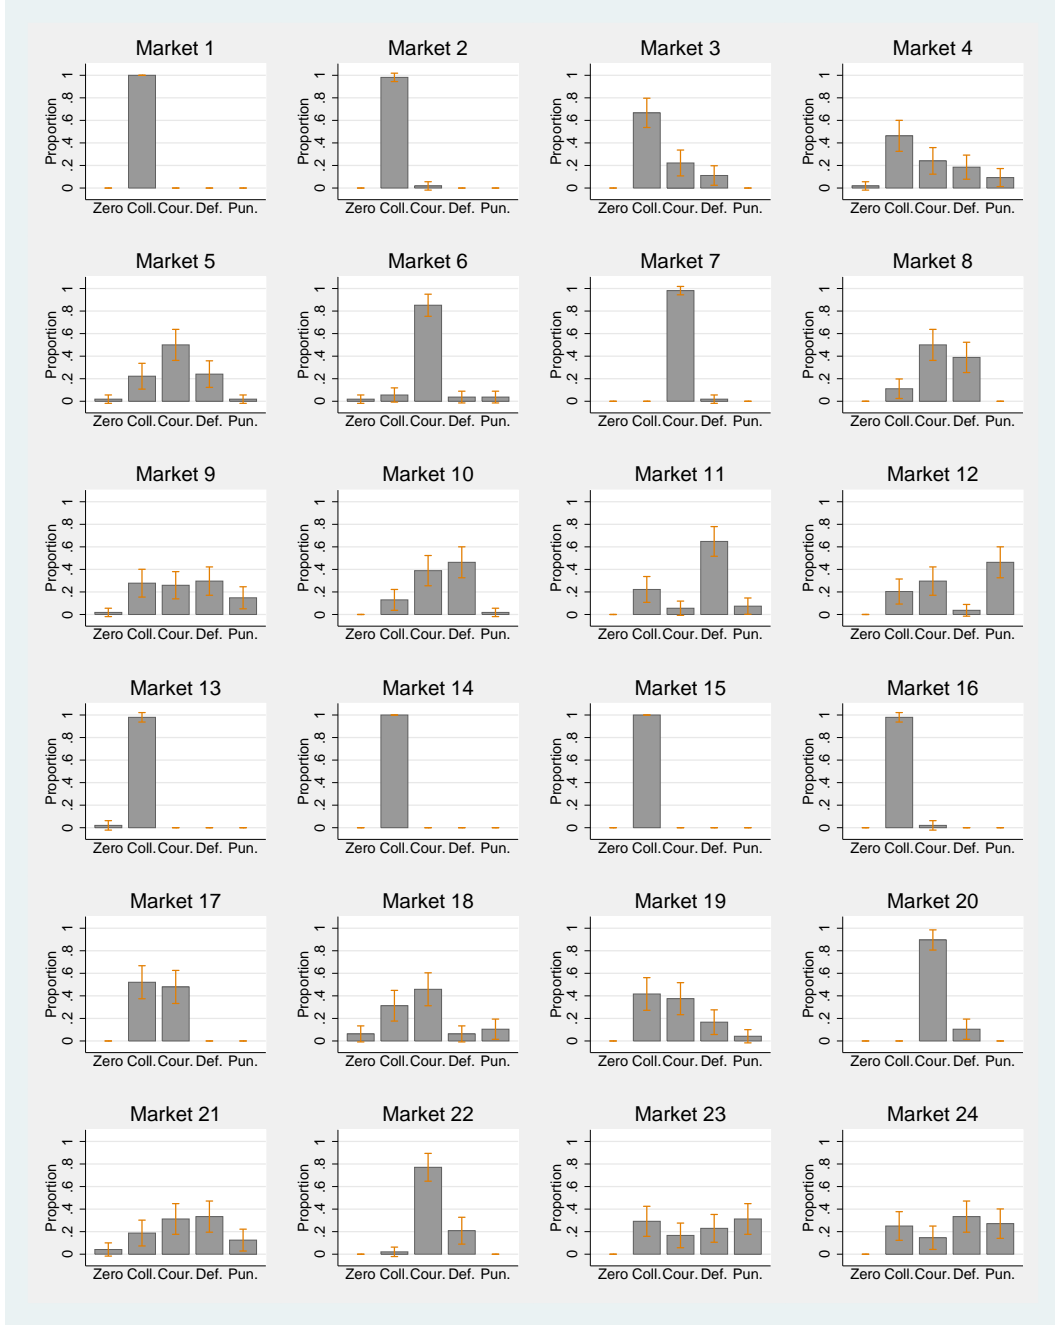
Note: Markets 25-36 (37-48) in Summer (Spring) session.

Figure C.6: Total Output per Round, C3 Treatment, All Markets



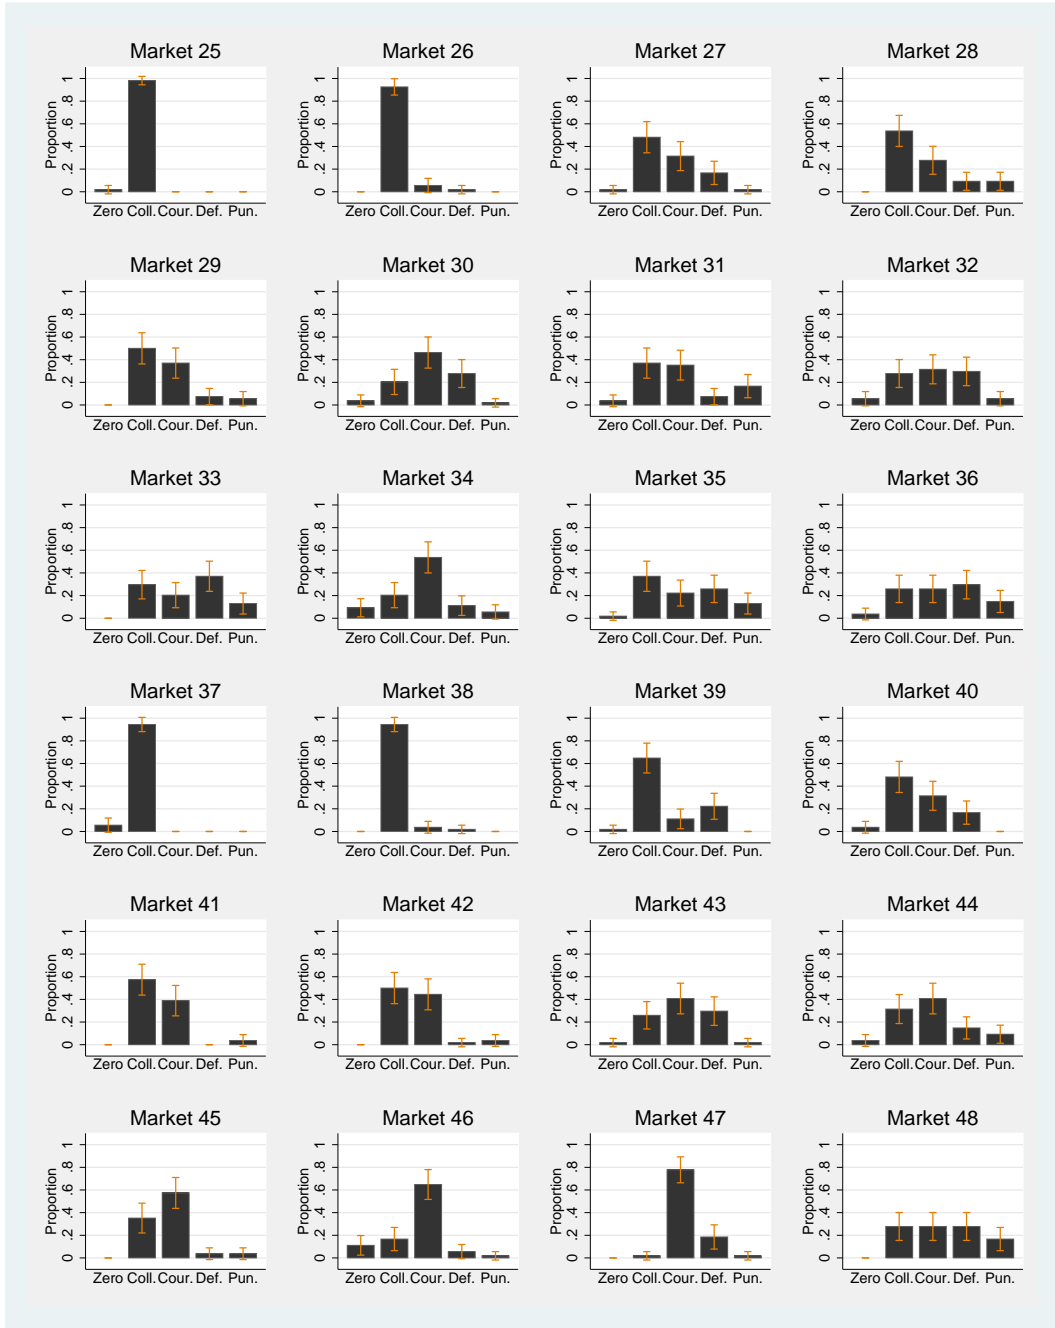
Note: Markets 49-59 (60-64) in unknown- (known-) end termination rule treatment.

Figure C.7: Proportion of Strategies with 95% Conf. Int., C2 Treatment, All Markets



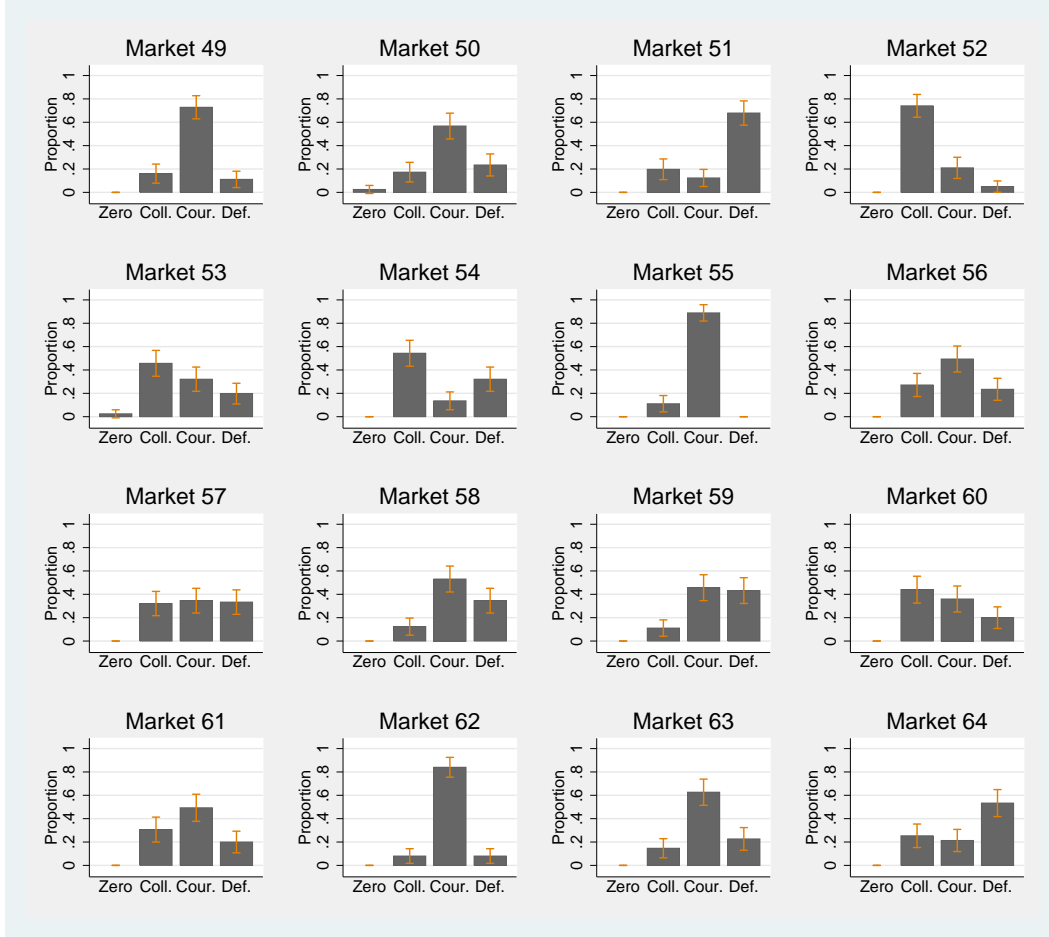
Note: Markets 1-12 (13-24) in unknown- (known-) end termination rule treatment.

Figure C.8: Proportion of Strategies with 95% Conf. Int., FS2 Treatment, All Markets



Note: Markets 25-36 (37-48) in Summer (Spring) session.

Figure C.9: Proportion of Strategies with 95% Conf. Int., C3 Treatment, All Markets



Note: Markets 49-59 (60-64) in unknown- (known-) end termination rule treatment.

D Instructions

D.1 Single Phase Treatment

Brackets, [], denote differences between two-firm and three-firm instructions.

You are about to participate in an experiment in economic decision making. If you follow the instructions carefully, you can earn a considerable amount of money. At the end of today's session, you will be paid your earnings in private and in cash.

Overview

The experiment will last several decision rounds. You will not know the number of rounds until the end of the experiment. At the beginning of the session, you will be randomly and anonymously matched with one [two] other person[s]. The one [two] other person[s] with whom you will be matched will be the same in every round, but you will not learn the identity of the other person[s]. The decisions that you and the other [two] person[s] make will determine the dollar earnings for each of you.

In this session, you are a quantity-setting seller of a hypothetical good. You will earn profits by selling units of the good. At the beginning of each round, you will be asked how many units of the good you want to sell in that round. You make a decision by selecting a number from a list on your computer. The possible choices are 0, 30, 40, 45, or 60 [0, 20, 30, or 40] units. At the same time that you are submitting how many units you want to sell, the other [two] seller[s] in your 2[3]-seller market will also submit how many units he/she [they] wants to sell. None of you will be able to see the decisions of the other [two] seller[s] in your market until both [all three] of you have submitted your decisions. Note that once submitted, all decisions are final and cannot be changed. At the end of each round, you will see how many units you sold, how many units the other seller[s] in your 2[3]-seller market sold, how many total units were sold, the price for that round, and your earnings for that round. Earnings are denoted in tokens and each unit has a cost of 0 (zero) tokens to you.

Price Calculation

Buyers are automated by the computer program. The market price at the end of a round will be determined by the units sold by both [all three] sellers in your 2[3]-seller market in a round. At the end of a round, the computer will calculate the market price (in tokens) as follows:

$$\text{Price} = 120 - \text{Total Units Sold Stage}$$

In general, the higher the number of total units sold the lower the price and vice versa. If the total amount of units sold across both stages is greater than or equal

to 120, the price will be zero. Hint: The number of total sales in your 2[3]-seller market can never be greater than 120 in any round.

Earnings

You will earn profits by selling units. The profit for any unit sold is the selling price in that round. Your total earnings (in tokens) in a round will be calculated as follows:

$$\text{Round Earnings} = \text{Price} \cdot \text{Your Total Units Sold}$$

Your total earnings in this part of the experiment will be your total earnings from all rounds. At the end of the first experiment, tokens will be converted into U.S. dollars at a rate of 1,800 tokens per U.S. dollar.

The following table shows your possible earnings in each round based on the sales choices that you and the other seller in your 2[3]-seller market make:

		Other Seller's Total Units Sold				
		0	30	40	45	60
Your Units Sold	0	0	0	0	0	0
	30	2,700	1,800	1,500	1,350	900
	40	3,200	2,000	1,600	1,400	800
	45	3,375	2,025	1,575	1,350	675
	60	3,600	1,800	1,200	900	0

		Other Two Sellers' Total Units Sold							
		0	20	30	40	50	60	70	80
Your Units Sold	0	0	0	0	0	0	0	0	0
	20	2,000	1,600	1,400	1,200	1,000	800	600	400
	30	2,700	2,100	1,800	1,500	1,200	900	600	300
	40	3,200	2,400	2,000	1,600	1,200	800	400	0

Before making any final decisions, you will be asked to answer 4 (four) practice questions to verify that you understand how your earnings are determined.

Computer Program

At the top of your screen, you will see a payout table similar to the table above (gray frame). In the middle of your screen, you will see the actual decision panel (orange frame). You will make a decision by selecting how many units you want to sell from the list. Once you click the "Submit" button, your sales decision cannot be changed

and will be final. At the end of each round, the computer will display your sales, the other [two] seller's sales, the total sales in your 3-seller market, the price, and your profit for that round. The computer will keep track of your sales, the other seller's sales, the price, and your earnings in each round. This information will be displayed in a table at the bottom of your computer screen (gray frame). The computer will also update your total earnings which will be displayed at the top of your screen (in tokens).

Summary

- The experiment will last several decision rounds. You will not know the number of rounds until the end of the experiment.
- You will be randomly matched with one [two] other seller[s]. The one [two] other seller[s] with whom you will be matched will be the same in every round!
- You will earn profits by selling units. The profit for any unit sold is the selling price in that round.
- The price that you will receive for each unit you sell in a round is calculated as follows: $\text{Price} = 120 - \text{Total Units Sold}$
- Your earnings will be the sum of your earnings from all rounds.

If you have a question at any point during the experiment, please raise your hand! One of the monitors will come to your station and answer it in private.

D.2 Two-Phase Treatment

You are about to participate in an experiment in economic decision making. If you follow the instructions carefully, you can earn a considerable amount of money. At the end of today's session, you will be paid your earnings in private and in cash.

Overview

The experiment will last several decision rounds. You will not know the number of rounds until the end of the experiment. At the beginning of the session, you will be randomly and anonymously matched with one other person. The one other person with whom you will be matched will be the same in every round, but you will not learn the identity of the other person. The decisions that you and the other person make will determine the dollar earnings for each of you.

In this session, you are a quantity-setting seller of a hypothetical good. You will earn profits by selling units of the good. Each round consists of two stages (A and

B). At the beginning of each stage, you will be asked how many units of the good you want to sell in that stage. You make a decision by selecting a number from a list on your computer. (In stage A, the possible choices are 0 and 24 units. In stage B, you will have five choices which depend on the decisions in stage A.) At the same time that you are submitting how many units you want to sell, the other seller in your 2-seller market will also submit how many units he/she wants to sell. None of you will be able to see the decisions of the other seller in your market until both of you have submitted your decisions. Note that once submitted, all decisions are final and cannot be changed.

At the end of stage A, you will see how many units you sold, how many units the other seller in your 2-seller market sold and how many units were sold in total in Stage A. At the end of stage B, you will see how many units you sold, how many units the other seller in your 2-seller market sold, the price for that round, and your earnings for that round. Earnings are denoted in tokens and each unit has a cost of 0 (zero) tokens to you.

Price Calculation

Buyers are automated by the computer program. The market price at the end of a round will be determined by the units sold by both sellers in your 2-seller market in a round in stage A and stage B combined. At the end of a round, the computer will calculate the market price (in tokens) as follows:

$$\text{Price} = 120 - \text{Total Units Sold Stage A} - \text{Total Units Sold Stage B}$$

In general, the higher the number of total units sold the lower the price and vice versa. If the total amount of units sold across both stages is greater than or equal to 120, the price will be zero. Hint: The number of total sales in your 2-seller market can never be greater than 120 in any round.

Earnings

You will earn profits by selling units. The profit for any unit sold is the selling price in that round. Your total earnings (in tokens) in a round will be calculated as follows:

$$\text{Round Earnings} = \text{Price} \cdot \text{Your Total Units Sold}$$

Note: You will receive the same price for any unit sold in stage A and/or stage B. Your total earnings in the experiment will be your total earnings from all rounds. At the end of the first experiment, tokens will be converted into U.S. dollars at a rate of 1,800 tokens per U.S. dollar.

Attached is a table that shows your possible earnings in each round based on the sales

choices that you and the other seller in your 2-seller market make in both stages. Before making any final decisions, you will be asked to answer 4 (four) practice questions to verify that you understand how your earnings are determined.

Computer Program

At the top of your screen, you will see a payoff table similar to the table above (gray frame). In the middle of your screen, you will see the actual decision panel (orange frame). You will make a decision by selecting how many units you want to sell from the list. Once you click the “Submit” button, your sales decision cannot be changed and will be final.

At the end of each round, the computer will display your sales, the other seller’s sales, the total sales in your 2-seller market, the price, and your earnings for that round. The computer will keep track of your sales, the price, and your earnings in each round. This information will be displayed in a table at the bottom of your computer screen (gray frame). The computer will also update your total earnings which will be displayed at the top of your screen (in tokens).

Summary

- The experiment will last several decision rounds. You will not know the number of rounds until the end of the experiment.
- You will be randomly matched with one other seller. The one other seller with whom you will be matched will be the same in every round!
- You will earn profits by selling units in either stage A, or stage B, or both stages. The profit for any unit sold is the selling price in that round.
- The price that you will receive for each unit you sell in a round is calculated as follows: $\text{Price} = 120 - \text{Total Units Sold Stage A} - \text{Total Units Sold Stage B}$
- Your earnings will be the sum of your earnings from all rounds.

If you have a question at any point during the experiment, please raise your hand! One of the monitors will come to your station and answer it in private.

My Sales A	Other Seller's Sales A	My Sales B	Other Seller's Sales B	Total Sales	Price	My Profit	Other Seller's Profit
0	0	0	0	0	120	0	0
0	0	0	30	30	90	0	2,700
0	0	0	40	40	80	0	3,200
0	0	0	45	45	75	0	3,375
0	0	0	60	60	60	0	3,600
0	0	30	0	30	90	2,700	0
0	0	30	30	60	60	1,800	1,800
0	0	30	40	70	50	1,500	2000
0	0	30	45	75	45	1,350	2025
0	0	30	60	90	30	900	1800
0	0	40	0	40	80	3,200	0
0	0	40	30	70	50	2,000	1,500
0	0	40	40	80	40	1,600	1,600
0	0	40	45	85	35	1,400	1,575
0	0	40	60	100	20	800	1,200
0	0	45	0	45	75	3,375	0
0	0	45	30	75	45	2,025	1,350
0	0	45	40	85	35	1,575	1,400
0	0	45	45	90	30	1,350	1,350
0	0	45	60	105	15	675	900
0	0	60	0	60	60	3,600	0
0	0	60	30	90	30	1,800	900
0	0	60	40	100	20	1,200	800
0	0	60	45	105	15	900	675
0	0	60	60	120	0	0	0

Table continued on next page.

My Sales A	Other Seller's Sales A	My Sales B	Other Seller's Sales B	Total Sales	Price	My Profit	Other Seller's Profit
0	24	0	0	24	96	0	2,304
0	24	0	6	30	90	0	2,700
0	24	0	32	56	64	0	3,584
0	24	0	33	57	63	0	3,591
0	24	0	48	72	48	0	3,456
0	24	30	0	54	66	1,980	1,584
0	24	30	6	60	60	1,800	1,800
0	24	30	32	86	34	1,020	1,904
0	24	30	33	87	33	990	1,881
0	24	30	48	102	18	540	1,296
0	24	32	0	56	64	2,048	1,536
0	24	32	6	62	58	1,856	1,740
0	24	32	32	88	32	1,024	1,792
0	24	32	33	89	31	992	1,767
0	24	32	48	104	16	512	1,152
0	24	45	0	69	51	2,295	1,224
0	24	45	6	75	45	2,025	1,350
0	24	45	32	101	19	855	1,064
0	24	45	33	102	18	810	1,026
0	24	45	48	117	3	135	216
0	24	48	0	72	48	2,304	1,152
0	24	48	6	78	42	2,016	1,260
0	24	48	32	104	16	768	896
0	24	48	33	105	15	720	855
0	24	48	48	120	0	0	0

Table continued on next page.

My Sales A	Other Seller's Sales A	My Sales B	Other Seller's Sales B	Total Sales	Price	My Profit	Other Seller's Profit
24	0	0	0	24	96	2,304	0
24	0	0	30	54	66	1,584	1,980
24	0	0	32	56	64	1,536	2,048
24	0	0	45	69	51	1,224	2,295
24	0	0	48	72	48	1,152	2,304
24	0	6	0	30	90	2,700	0
24	0	6	30	60	60	1,800	1,800
24	0	6	32	62	58	1,740	1,856
24	0	6	45	75	45	1,350	2,025
24	0	6	48	78	42	1,260	2,016
24	0	32	0	56	64	3,584	0
24	0	32	30	86	34	1,904	1,020
24	0	32	32	88	32	1,792	1,024
24	0	32	45	101	19	1,064	855
24	0	32	48	104	16	896	768
24	0	33	0	57	63	3,591	0
24	0	33	30	87	33	1,881	990
24	0	33	32	89	31	1,767	992
24	0	33	45	102	18	1,026	810
24	0	33	48	105	15	855	720
24	0	48	0	72	48	3,456	0
24	0	48	30	102	18	1,296	540
24	0	48	32	104	16	1,152	512
24	0	48	45	117	3	216	135
24	0	48	48	120	0	0	0

Table continued on next page.

My Sales A	Other Seller's Sales A	My Sales B	Other Seller's Sales B	Total Sales	Price	My Profit	Other Seller's Profit
24	24	0	0	48	72	1,728	1,728
24	24	6	0	54	66	1,980	1,584
24	24	24	0	72	48	2,304	1,152
24	24	33	0	81	39	2,223	936
24	24	36	0	84	36	2,160	864
24	24	0	6	54	66	1,584	1,980
24	24	6	6	60	60	1,800	1,800
24	24	24	6	78	42	2,016	1,260
24	24	33	6	87	33	1,881	990
24	24	36	6	90	30	1,800	900
24	24	0	24	72	48	1,152	2,304
24	24	6	24	78	42	1,260	2,016
24	24	24	24	96	24	1,152	1,152
24	24	33	24	105	15	855	720
24	24	36	24	108	12	720	576
24	24	0	33	81	39	936	2,223
24	24	6	33	87	33	990	1,881
24	24	24	33	105	15	720	855
24	24	33	33	114	6	342	342
24	24	36	33	117	3	180	171
24	24	0	36	84	36	864	2,160
24	24	6	36	90	30	900	1,800
24	24	24	36	108	12	576	720
24	24	33	36	117	3	171	180
24	24	36	36	120	0	0	0

E Screen-shots

Figure E.1: Decision Stage A, FS2 Treatment

Round 2 : Stage A My Balance: 1800 tokens

Payoff Table (Note: The other seller sees exactly the same payoff table.)

My Sales A	Other Seller's Sales A	My Sales B	Other Seller's Sales B	My Total Sales	Other Seller's Total Sales	Price	My Profit	Other Seller's Profit
0	0	30	30	30	30	60	1800	1800
0	0	30	40	30	40	50	1500	2000
0	0	30	45	30	45	45	1350	2025
0	0	30	60	30	60	30	900	1800
0	0	40	30	40	30	50	2000	1500
0	0	40	40	40	40	40	1600	1600
0	0	40	45	40	45	35	1400	1575
0	0	40	60	40	60	20	800	1200
0	0	45	30	45	30	45	2025	1350

The payoff table above shows the possible outcomes based on the sales choices in Stage A and the sales choices in Stage B.
Use the table in deciding how many units you want to sell in stage A of Round 2.

Choose how many units you want to sell in Stage A of Round 2 by selecting a number from the list to the right.
Click "Submit" to confirm your choice.

Other Seller's Choices A	My Sales A
0	0
24	24

Submit

Round	My Sales A	Other Seller's Sales A	My Sales B	Other Seller's Sales B	My Total Sales	Other Seller's Total Sales	Price	My Profit
1	24	0	6	30	30	30	60	1800

Figure E.2: Decision Stage B, FS2 Treatment

Round 2 : Stage B My Balance: 1800 tokens

Payoff Table

My Sales A	Other Seller's Sales A	My Sales B	Other Seller's Sales B	My Total Sales	Other Seller's Total Sales	Price	My Profit	Other Seller's Profit
24	24	0	0	24	24	72	1728	1728
24	24	0	6	24	30	66	1584	1980
24	24	0	24	24	48	48	1152	2304
24	24	0	33	24	57	39	936	2223
24	24	0	36	24	60	36	864	2160
24	24	6	0	30	24	66	1980	1584
24	24	6	6	30	30	60	1800	1800
24	24	6	24	30	48	42	1260	2016
24	24	6	33	30	67	33	990	1881

The payoff table above shows the possible outcomes based on the sales decisions in Stage B.
Use the table in deciding how many units you want to sell in stage B of Round 2.

My Sales A: 24
Other Seller's Sales A: 24
Total Sales A: 48

Choose how many units you want to sell in Stage B of Round 2 by selecting a number from the list to the right.
Click "Submit" to confirm your choice.

Other Seller's Choices B	My Sales B
0	0
6	6
24	24
33	33
36	36

Submit

Round	My Sales A	Other Seller's Sales A	My Sales B	Other Seller's Sales B	My Total Sales	Other Seller's Total Sales	Price	My Profit
1	0	24	30	6	30	30	60	1800