Cheap Talk and Investment Cycles

G. Brunekreeft and T. McDaniel*

September 2008

Do not quote without permission

Abstract

In markets with some competition and lumpy investment firms face the coordination problem of deciding who invests. These markets can be characterized by periods of over or under capacity, as described by business cycle models, if firms are myopic. We consider a situation in which myopic firms would find themselves in a business cycle. We use experiments to ask how these firms move out of the cycle and whether coordination devices such as cheap talk help to break the cycle. In groups of two, players make output choices and participate in one of the following treatments: (i) no information, where players only see their own output, profit and the market price; (ii) full information, where players additionally see the market output; (iii) symmetric cheap talk, which replicates the no information treatment but allows players to send simultaneous output messages before making choices, and (iv) asymmetric cheap talk where (only) player 1 knows market output and sends an output message to player 2 before choices are made. Results suggest convergence is fastest with symmetric cheap and with no information; prices are highest when there is asymmetric information, and in all case, prices eventually converge to the Cournot equilibrium.

1. Introduction

In markets with some competition and lumpy investment firms face the coordination problem of deciding who invests or how much to invest. These markets can be characterized by periods of over or under capacity, as described by business cycle models, if firms are myopic. Even with a small number of firms prices may be cyclical if agents have little information about supply and demand conditions. We are interested in the general case in which information is imperfect and price cycles are possible when firms are myopic. Even when information is perfect, however, firms still face a coordination problem in deciding how much investment to make.

Many models describe how agents update their information in the context of uncertain market conditions for varying amounts of competition. With belief based learning models such as best response or fictitious play, agents base their choices on some array of opponents' past output decisions. These models are usually applied to oligopoly markets, and if all agents update in this way, outcomes should be close to the Cournot-Nash equilibrium. Alternatively, agents may attempt to imitate the most profitable competitor. Vega-Redondo (1997) describes how this imitation generates competitive outcomes even with a small number of firms. In larger markets business cycles may be short lived if agents have adaptive expectations (Nerlove, 1958). The model that best predicts actual outcomes will depend on the number of agents as well as the information feedback they receive over relevant time periods. To imitate, agents must know others' individual actions as well as the profitability of those actions. To best-respond, agents need only know the aggregate output of their competitors. When the appropriate information is available, imitation should lead to

* Jacobs University Bremen; Bremer Energie Institut, Campus Ring 1; D-28759 Bremen; Germany. Email g.brunekreeft@jacobs-university.de and Department of Economics, Appalachian State University, Boone, NC 28607. Email: mcdanieltm@appsate.edu. The authors would like to thank Tim Perri, Mike McKee, Mark Strazicich, John Dawson, Stephen Cotton and participants at the Pasadena ESA meetings for helpful comments and advice.

competitive outcomes while best response behavior should lead to the Cournot-Nash equilibrium (Huck et. al, 1999; Offerman et. al, 2002).

We are interested in a related but different issue. Namely, in the absence of aggregate and individual data about other firms, can agents use cheap talk to coordinate their actions thereby breaking cycles so that prices converge to the relevant equilibrium? To answer this question, we describe a theoretical model where investment cycles are possible when information is incomplete. We then use experiments to determine if subjects are able to use cheap talk messages to break out of investment cycles. We compare output decisions in treatments with cheap talk but no market information to treatments with full market information but no cheap talk. Cheap talk may be used in a number of ways, coordination being just one. In our experiments subjects' messages are limited to their intended output decisions. Thus, attempted pre-emption is another possible use of messages.

There is a small literature on the credibility of cheap talk in games with both complete (e.g., Farrell, 1988 and Aumann, 1990) and incomplete information (Baliga and Morris, 2002). As we briefly discuss in section 3.2 these conditions do not hold generally for the game at hand.

The remainder of the paper is organized as follows: section 2 discusses the theoretical model for general parameters. In section 3 we describe the experiment and the parameter values chosen. Section 4 discusses our results and section 5 concludes.

2. Theory

In the general model *n* firms make a choice over the number of plants (or machines) to invest in.

Notation and definitions:

- x_i is the number of (identical) plants for firm *i*
- q_i is the output per plant for firm *i*
- Demand in period t is linear: $q^t = \dfrac{a - p^t}{b}$, and inverse demand is: $p^t = a - bq^t$.
- Industry supply in period *t* is: $p^t = \sum_{i=1}^{n} x_i q_i(p^e)$.

A firm's profit function is: $\pi_i = (pq_i - cq_i^2)x_i - \beta x_i^2$ where c is the unit cost of output and β is the unit cost of capital.

The solution to the firm's problem depends on the market conditions. Here we solve the general case for competitive firms. Differentiating the profit function and solving for q_i and x_i:

$$q_i = \frac{p}{2c} \text{ and } x_i = \frac{p^2}{8c\beta}.$$

If firms are symmetric industry supply reduces to: $nq_i x_i = n\dfrac{p^3}{16\beta c^2}$, and demand equals supply if:

$$\frac{a-p}{b} = n\frac{p^3}{16\beta c^2}.$$

This can be solved numerically for chosen parameter values. A competitive cobweb will be stable or not depending on the relative elasticities of the supply and demand curves. A cobweb will eventually converge if the demand curve is flatter than the supply curve. For a given number of firms, we can choose values for *b* and β to achieve either of these effects. A fully optimized model would find values for both q_i and x_i; however, for our purpose it is sufficient to fix q_i and allow agents to choose the level of investment, knowing each 'plant' will produce the same quantity.

Letting $q_i = \bar{q}$, profit is given by $\pi = (p\bar{q} - c\bar{q}^2)x_i - \beta x_i^2$, and the optimal investment becomes:

$x_i = \dfrac{p\bar{q} - c\bar{q}^2}{2\beta}$. For the case of $\bar{q}$ = 1, this reduces to:

$$x_i = \frac{p-c}{2\beta} \text{ giving industry supply = } \mathrm{Q}^s = \mathrm{n}\frac{(\mathrm{p} - \mathrm{c})}{2\beta}.$$

Stability depends on the relative steepness of supply and demand curves. For $\bar{q}$ = 1 the slope of the supply curve is $\dfrac{2\beta}{n}$. The slope of the demand curve is always –*b*.

If firms engage as Cournot competitors we can also observe price cycles in this model if firms use best reply dynamics (ref). Best response 'learning' is similar to myopic behavior in the cobweb model, but in this case there are few players, and the important feedback variable is aggregate quantity, not price. Also, best response cycles converge fairly quickly.

For the case of $\bar{q}$ = 1 and Cournot competition, firms maximize profits given by: $\pi = p(x_i + x_j)x_i - cx_i - \beta x_i^2$. When firms are symmetric the first order conditions can be solved for x_i and x_j giving:

$$x_i = x_j = \frac{a-c}{3b+2\beta}. \text{ Aggregate output is } \mathrm{Q}^s = \frac{2(\mathrm{a} - \mathrm{c})}{3\mathrm{b} + 2\beta}, \text{ and } p = a - bQ^s.$$

3. Experiment

3.1 Parameters

We compare agents' behavior under four scenarios. In all cases two competing firms face a linear demand curve and have increasing marginal cost. In the benchmark treatment (T0) subjects have no information about aggregate output. These subjects see their own profit and price each period. In T1 subjects have full information about price, aggregate supply and own profits each period. Additionally, they are explicitly given an equation for the price calculation in their instructions. In essence, these subjects have both the supply and demand curves, however, they still make their choices in ignorance of what their partner will choose. In T2 subjects can simultaneously send one round of non-binding output messages before making their actual choices. Apart from messages, these subjects have the same information as subjects in the no information treatment. We allow two-way cheap talk since subjects are symmetric in T2. In T3, player 1 is also told aggregate demand each period and can unilaterally send a cheap talk message to player 2.

We chose very simple parameters to lessen subjects' confusion. Our parameter choices are: $\bar{q}$ = 1, *n* = 2, *b* = 0.9, *a* = 20, *c* = 1, β = 1 and p_0 = 2, where p_0 is the initial price. The market clears if:

$$\frac{20-p}{0.9}=\frac{(p-1)}{1}.$$

With these parameters, the slopes of demand and supply are -0.9 and 1, respectively. As a competitive system, the model is stable and converges to p = 11 and $Q^s = nq_i = 10$. When firms engage as Cournot competitors the Cournot-Nash equilibrium is $Q^s = 8$ and p = 12.8.

There are 21 periods in the experiment. Subject made choices in periods 1-20; in period 21 they just received their final feedback. Each period, subjects could choose a quantity between 0 and 18. Relative to many related experiments ours is quite short, but with only 2 players in each group and a small strategy space we expected convergence to occur fairly quickly.[1] In such a set-up subjects who found the equilibrium early could begin to experiment in later rounds thereby adding noise to the data. Twenty periods was a good compromise for our purpose.

We chose *n* = 2 for a number of reasons. First, we are primarily interested in whether cheap talk is an effective mechanism for lessening investment cycles. Adding more subjects to a group increases competition, but we do not have clear theoretical predictions for the cheap talk protocol when there are several players in the group. When group size is two, we expect to observe price fluctuations, but prices should converge to the Cournot-Nash equilibrium; it is the speed of convergence across the four treatments we are most interested in, however. Table 1 shows the predicted price, quantities, and profits for the competitive, Cournot, and cartel outcomes. Also, we give the Stackelberg solution as a way of testing whether or not player 1's superior information in T3 and her ability to send a message regarding her intended play confers a first mover advantage.

If firms are able to collude they should split the monopoly output of 5 thereby each producing 2.5. Because quantity choices are discrete however, the collusive outcome will be for each player to produce 3.[2] Similarly, since subjects can only specify discrete quantities, we do not notice a difference theoretically between the Cournot and Stackelburg outcomes (we use decimal places here only to show there is a minor difference). However, we may in fact observe a leadership advantage for player 1 in T3.

	Price	q_i	q_j	π_i	π_j	sum π
Competitive	11.0	5.0	5.0	25.0	25.0	50.0
Cournot	12.8	4.0	4.0	31.2	31.2	62.4
Cartel	14.6	3.0	3.0	31.8	31.8	63.6
Stackelburg	12.6	4.3	3.98	31.2	30.1	61.3

Table 1: Predicted price, output and profit

3.2 Hypotheses

We use experimental data to test the following hypotheses:

[1] Offerman et. al (2002) use 100 periods; Huck et al. (2000) use 40; Cox and Walker (1998) use 30-37. However, early experiments by Carlson (1967) showed price convergence using only 6 and 9 periods.

[2] Although profits are higher for (q_i,q_j) = (5,0) than (6,0), individual profits are higher at (3,3) than (2.5, 2.5).

H1: price fluctuates are the same with full information (T1) and no information (T0).
H1': price converges as quickly with full information (T1) as with no information (T0).

H2: price fluctuates are the same with symmetric cheap talk (T2) and no information (T0).
H2': price converges as quickly with symmetric cheap talk (T2) as with no information (T0).

H3: price fluctuations are the same with symmetric cheap talk (T2) as with full information (T1).
H3': price converges as quickly with symmetric cheap talk (T2) as with full information (T1).

H4: Prices are higher in T2 than in T3; from Table 1.
H5: The better-informed player produces more and has higher profit with asymmetric information (T3).
H6: Cheap talk does not lead to collusion.

H1 – H3 are stated as null hypotheses while H4 – H6 are stated as the alternative hypothesis. Also, H1 – H3 are stated in both a weak and a strong form. We can test the weak form (fluctuations) looking at the variance of the price trends. For the strong form (convergence) we use the coefficient of variation (CV). We do not have a strict relative price expectation for T3, so we do not state a direct hypothesis for those prices.

H6 arises from the fact that cheap talk is not theoretically credible at the cartel solution. McDaniel (2008) discusses the literature related to the credibility of cheap talk in games with complete and incomplete information. Briefly, Farrell (1988) and Aumann (1990) provide the conditions for non-binding statements to have meaning in the context of full information. Player *i*'s message should be believed if he would want to follow through with the corresponding action if his message were believed, and if he would want to choose that action *only* when he expected player *j* to best-respond. These conditions are not tailored to dynamic games but they have some relevance for our model. For example, assume subjects have arrived at the cartel outcome (3,3). In that case they become involved in a classic prisoner's dilemma where there is an individual incentive to increase output to 4. Thus, if player *i* says she intends to continue producing 3 the next period, player *j* should not believe her (theoretically). Player *j* should reason that player *i* would prefer to choose 4 if she believed *j* would choose 3. On the other hand, if players arrive at the Cournot solution, it is credible to suggest continuing that strategy.

When information is incomplete, Baglia and Morris (2002) provide two additional provisos. The informed player must choose the Stackelburg outcome and the uninformed player must best respond to that choice. Again, these conditions are not adapted to dynamic games. We expect only a message of 4 to be believed in T2 or T3.

4. Results

The experiments were programmed in z-tree (Fischbacher, 1999, version 3 of the software) and conducted at the Appalachian State Experimental Economics Lab (AppEEL). 138 subjects participated earning $17.50 on average. Sessions lasted approximately 75 minutes and each subject only participated in one session. The breakdown of subjects to treatments is shown in Table 2.

(T0) NoInfo	34 subjects (17 groups)
(T1) Full Info	34 subjects (17 groups)
(T2) Sym CT	34 subjects (17 groups)
(T3)Asym CT	36 subjects (18 groups)

Table 2: Break down of subjects to treatments

4.1 Price trends and variability

Figures 1 shows the price trends for the 4 treatments including bands for the first standard deviation. Table 3 shows the mean and variance of prices on average for periods 2-21, the last five periods and for the last period. We use the variance measures in table 3 to discuss H1, H2 and H3. We can reject H1 and conclude that prices fluctuate less with no information. This is true on average, for the last 5 periods and for the final period.[3] Comparing no information to symmetric cheap talk we conclude that prices vary more with no information in the last 5 and the final period. Fluctuations are greater in T0 for periods 2-21, but we cannot reject H2 for those periods (p = .11).

Comparing full information and symmetric cheap talk, we reject H3 and conclude that price fluctuations are greater with full information. Price fluctuations are higher on average with asymmetric information compared to no information and symmetric cheap talk, but statistically no different from full information. For the last five periods fluctuations remain higher in T3 relative to T2, but are less than T1 and not significantly different from T0. In the last period, price fluctuations are significantly lower in T3 than in the other treatments.

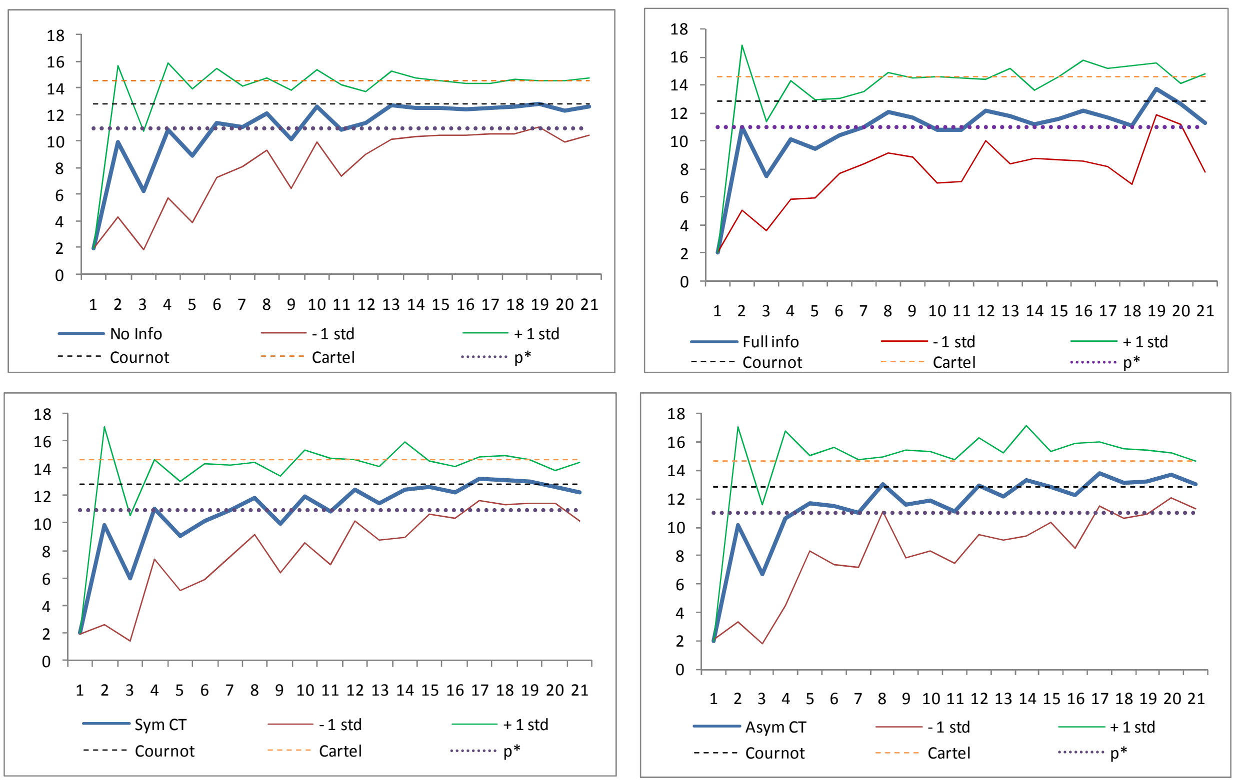

Figure 1: Price Trends

[3] We can have variation in price for the final period because our level of observation is the group.

Treatment	Average	Last 5 periods	Last period
T0: No info	11.42 (10.88)	12.55 (4.29)	12.59 (4.91)
T1: Full info	11.17 (11.63)	12.07 (9.77)	11.26 (12.53)
T2: Sym cheap talk	11.33 (10.97)	12.85 (2.75)	12.27 (4.46)
T3: Asym info	11.94 (13.72)	13.31 (4.25)	12.95 (2.88)

Table 3: Average price and (variance) for periods 2-21, 16-21 and 21

Table 4 shows p-values for bilateral price comparisons across treatments. The sign before the p-value indicates whether the first listed treatment has a higher or lower price than the second listed treatment. For example, the average price is higher with no information than with full information, but not significantly higher. The boxes that are lightly shaded are those showing a significant difference at the 5% level. Notice in particular that there are no significant price differences in the final period. From this data we can reject H4 for periods 2-21 and conclude that prices are higher in T3 than in T2. We cannot reject H4 for the last 5 periods or for the last period.

	Periods 2-21	Periods 17-21	Period 21
No info / Full info	(+) .4923	(+) .8290	(+) .3466
No info / Sym CT	(-) .9693	(-) .3168	(+) .7006
No info / Asym CT	(-) .0022	(-) .0192	(+) .6760
Full info / Sym CT	(-) .3769	(-) .2247	(-) .6611
Full info / Asym CT	(-) .0001	(-) .0058	(-) .1360
Sym CT / Asym CT	(-) .0018	(-) .0891	(-) .2987
All (K-Wallis)	.0003	.0249	.4963

Table 4: sign and p-value for price differences using a ranksum test.

4.2 Price trends and convergence

In this section we use the CV of the price trends to test hypotheses H1', H2' and H3'.[4] The CV is calculated as the standard deviation of prices each period divided by the mean price that period:

$$CV_i = \frac{\dfrac{\sum_{j=1}^{n}\sqrt{(p_j^i - \bar{p}^i)^2}}{n-1}}{\bar{p}^i}$$

[4] See Friedman (1992) for a discussion of CV and convergence tests. For applications, see Nieswiadomy and Strazicich (2004), and Dawson and Sen (forthcoming).

where i = 2 to 21 is the number of relevant periods, j = 1 to n is the number of groups (17 for T0, T1 and T2; 18 for T3), p_j^i is the price in period i for group j, and $\bar{p}^i$ is the mean price in period i.

Figure 2 graphs the CVs for the four treatments. We can say that one price trend is converging more quickly than another if the CV is significantly more negative. T1 is different from the other three treatments in that it does not appear to converge and shows only a weakly negative trend. In fact, a simple linear regression shows all four price trends are highly significantly negative, though the regression coefficient is the least negative for T1.

Using a ranksum test we reject H1' and conclude that prices converge more quickly with no information than with full information. Comparing T0 and T2, T2 shows more convergence, but not significantly more. Thus, we cannot reject H2' that the speed of convergence is the same when agents use symmetric cheap talk as when they have no information.

Comparing T1 and T2, we reject H3' and conclude that prices converge more quickly with symmetric talk than with full information. We also find faster convergence with symmetric than asymmetric cheap talk. Also, convergence is faster (but not significantly faster) in T0 than T3 and in T3 than in T1.

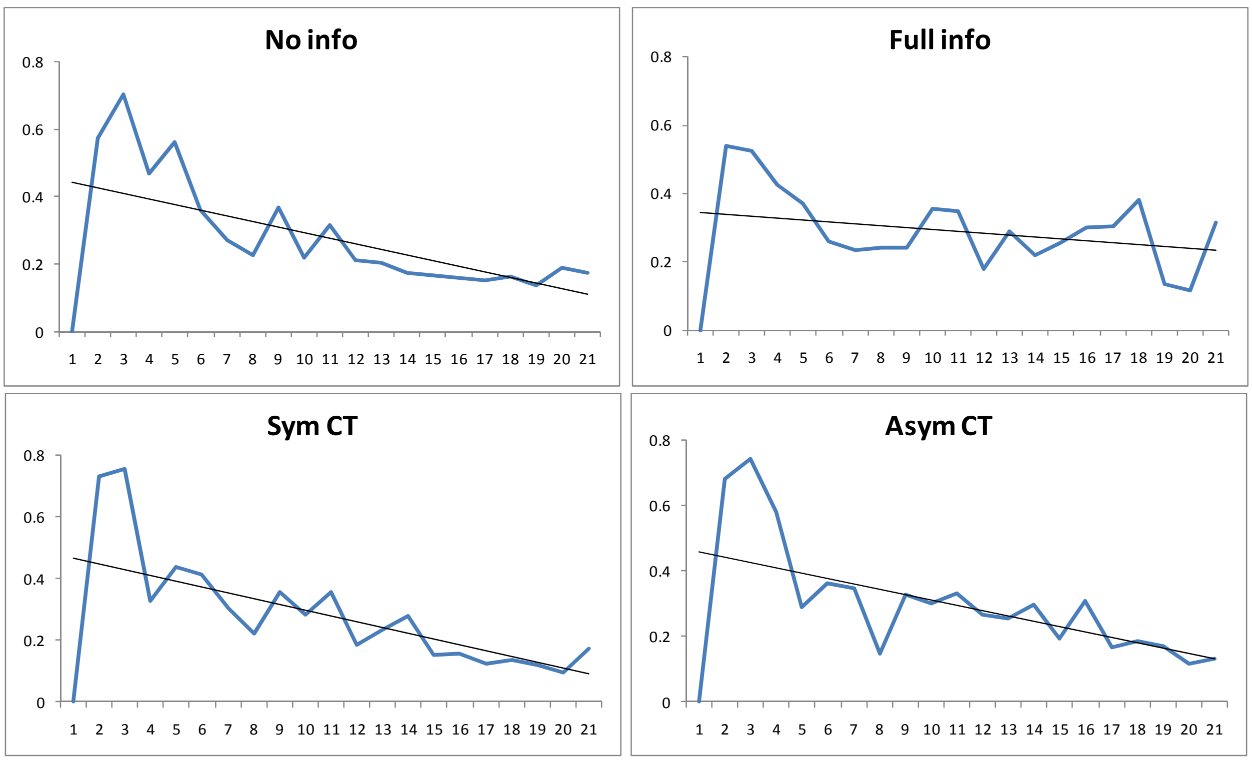

Figure 2: coefficient of variation over periods

Figure 2 and the corresponding tests do not say to which price the trends are converging. What we would like to find is convergence to either the competitive or Cournot equilibrium solutions. With only 2 players we expect the Cournot solution to arise if agents cannot sustain collusion. Using sign rank tests we can ask if prices are equal to given equilibrium values. Combining such a test with the CV tests above, we are able to say if one treatment converges faster than another to the equilibrium value. We conduct the sign rank test on the last 5 periods of each treatment and on the final period.[5]

For all treatments we reject that prices converge to the competitive or the collusive equilibria. Looking at the last 5 periods, prices are not significantly different from the Cournot outcome for T0, T1 or T2. Prices in T3 are significantly above the Cournot equilibrium (but still below the cartel equilibrium). Looking only at the final period, we cannot reject the Cournot equilibrium for any of the four treatments.

From this combination of tests, we conclude that prices are converging to the Cournot outcome in all treatments. The implied ranking of the treatments in terms of convergence speed is: T2 ≥ T0 ≥ T3 ≥ T1.

4.3 Output

Table 5 shows mean output choices across treatments; for T2 and T3 we show choices separately for players 1 and 2. H5 predicts player 1 will produce more than player 2 in T3. This hypothesis assumes player 1's information advantage is akin to a first mover advantage; thus we expect a Stackelburg outcome to arise.

The bilateral treatment comparisons that are significant are all those including T3; i.e., T0, T1 and T2 all have significantly greater output than T3. Thus, we can assume the following output ranking: T1 ≥ T2 ≥ T0 > T3. Comparing players 1 and 2, we find player 1 produces significantly more in T3 but not in T2, as expected. Each player type produces significantly more in T2 than in T3.

	T0	T1	T2	T3
Mean	4.79	4.95	4.87	4.52
Player 1			4.96	4.78
Player 2			4.61	4.44

Table 5: Output by treatment and player

The distribution of output choices is also similar across treatments. Figures 3 and 4 show histograms for group output levels in the four treatments in all periods and the last 5 periods, respectively. Because the output space is discrete, the cartel, Cournot and competitive solutions are all very close. Thus, we are not surprised to find the distributions clustered around these choices. Looking at all periods, the modal output level is the Cournot solution in T0 and T1; in T2, outputs of 7,8, 9 and 10 are closely ranked; in T3 seven is the modal output. The competitive outcome occurs most in T1 and T2; the cartel solution occurs most in T3, and the Cournot solution occurs most in T1. A mode at seven is indicative of a relationship whereby players alternate cheating in a cartel. We explore this possibility more below.

Using signrank tests we compare the instances of competitive, Cournot and cartel outcomes within treatments. The Cournot solution occurs more than the competitive solution in both T0 and T1. There is no significant difference in T2 and T3. The competitive solution occurs more than the cartel solution in T0, T1

[5] As table 1 shows, there is not much profit difference for subjects playing Cournot versus colluding. We are interested in which outcome subjects play, but we are more interested in the speed of convergence across treatments.

and T2 but less in T3. Likewise, the Cournot solution occurs more than the cartel solution in T0, T1 and T2, but less in T3.[6]

It almost appears there is an ‘outlier’ in T3; the same pattern occurs if we look only at the last period. However, using individual group data, we see this is not just one group (or even 2) acting ‘irrationally’. The group producing 22 in the last period is one that had behaved quite reasonably until the last period. Interesting, they progressively (collectively, not necessarily individually) increased their output over the course of the experiment. Nonetheless, the increase in output in the last period was a huge jump and arose because one player decided to produce 18. That action was not in line with the rest of the group’s play.

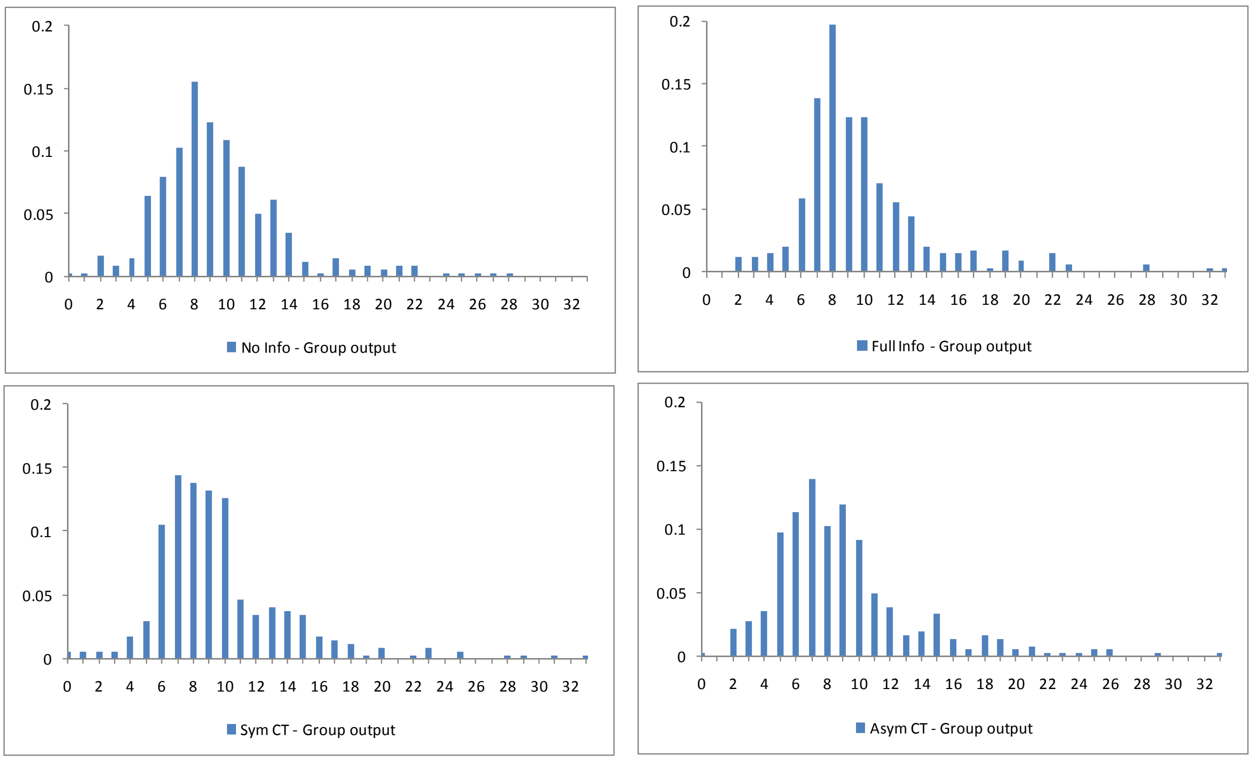

Figure 3: Histogram of group output choices (% of outcomes): all periods.

[6] Looking only at the last 5 periods, the Cournot solution occurs more than the competitive solution in both T1 and T2. There is no significant difference in T0 and T3. The competitive solution occurs more than the cartel solution in T0, equally often in T1 and T2, and less often in T3. The Cournot solution occurs more than the cartel solution in T0, T1 and T2, but less in T3.

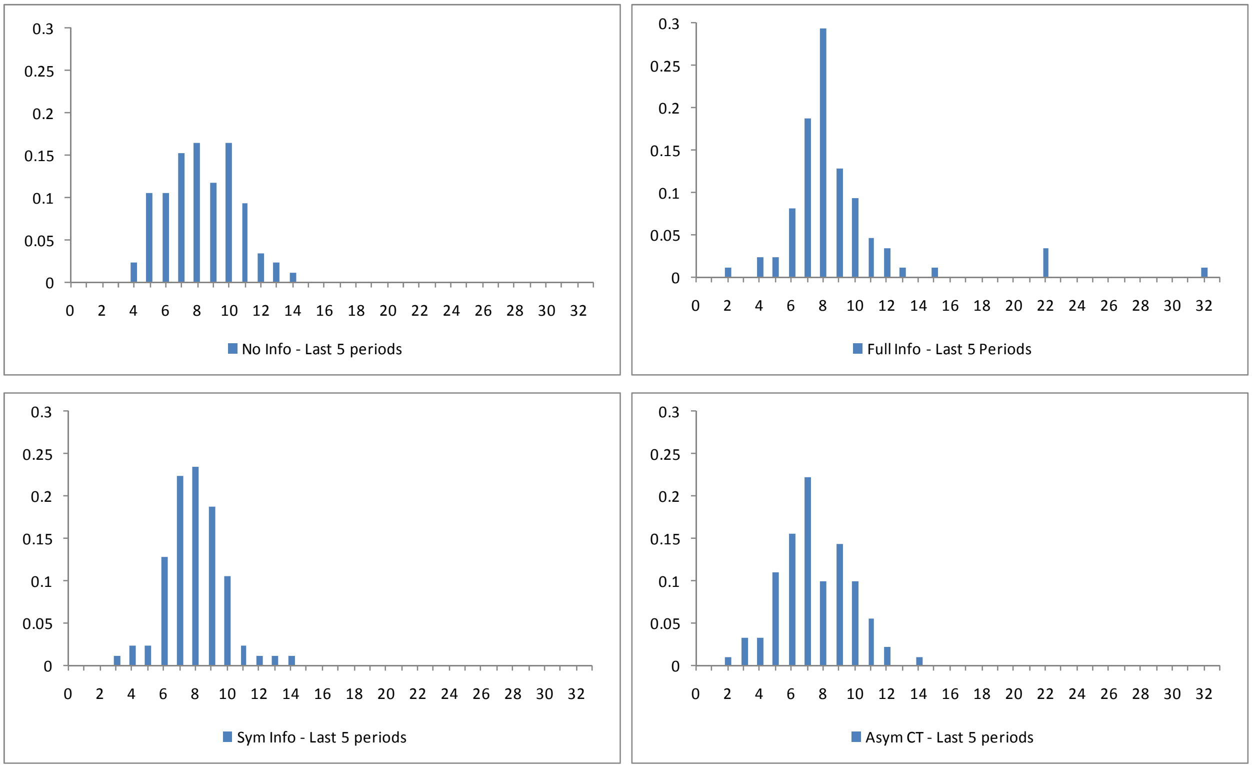

Figure 4: Histogram of group output choices (% of outcomes): last 5 periods.

Because the modal outcome is 7 in T2 and T3 we also looked at the distribution of choices for players 1 and 2 in those treatments. Based on the Stackelburg prediction we expect player 1 to have a higher modal output than player 2 in T3. If the Stackelburg strategy is being followed, one possibility is for the outputs to be very close with player 1 producing 4 and player 2 producing 3. There are two other possibilities, however. First, player 1 may be producing more than 4 in an attempt to pre-empt player 2. Second, player 1 may be producing less than 4 in an effort to move player 2 to the collusive outcome. Interestingly, in both T2 and T3 we find player 1 producing more than player 2. The modal output for player 1 is five while for player 2 it is four. This pattern persists in the last five periods of T3, but not in T2. In T2, the modal output for both players is 4 in the last 5 periods. Since player 1 has no reason to move player 2 to the competitive solution, producing often at 5 only makes sense if player 1 is trying to preempt player 2. We elaborate on this result in section 4.6 where we look at the relationship between the messages players send and the actions they take.

4.4 Profits

Table 6 shows the average points per period, average total points, and average profits for the four treatments. If we compare treatments using only the final profits, there are no significant bilateral differences between treatments. Moreover, the difference in profits between players in T2 and T3 are not significantly different. However, if we compare the distribution of profits using all periods, profits are significantly lower in T0 than T1 (p = .056) and lower in T0 than T3 (p = .01). Also, player 1 makes significantly higher profits than player 2 in T3 but not in T2. From the distribution of profit data and the outcome comparisons in 4.3, we fail to reject H5.

Treatment	Mean points/period	Total points	Final Profit
(T0) No info	20.03	420.61	16.82
(T1) Full info	20.42	428.72	17.15
(T2) Sym cheap talk	20.41	428.68	17.15
Player 1	20.57	431.94	17.28
Player 2	20.26	425.41	17.02
(T3) Asym cheap talk	21.35	448.29	17.93
Player 1	22.04	462.75	18.51
Player 2	20.66	433.83	17.35

Table 6: Profits

4.5 Learning

These experiments were not designed to compare our results to the literature on learning in oligopoly markets (e.g., Vega-Redondo, 1997; Huck et. al, 1999 and 2000; Offerman et. al, 2002). As such our subjects were not given the same information to make these comparisons possible. However, in this section we look more closely at potential learning patterns in T1 and T3 – the only two treatments where one or both players had aggregate output information. Subjects were not given information on the other player's profits in any of the treatments, so imitating the successful other is not straightforward (although it is possible since subjects know how profits are calculated). With only 2 subjects in a group T1 is similar to both treatments Q and Qq in Offerman et al. (2002). In their treatment Q, subjects received aggregate group output, where group size is 3; in treatment Qq subjects additionally receive disaggregated output information. If subjects are imitating successful others, Vega-Redondo (1997) shows that the competitive outcome is likely to result. Based on the *average* results in section 4, T1 appears to be the most competitive of the four treatments. However, using more disaggregated data this is no longer the obvious conclusion. The Cournot outcome is the strong modal choice over all periods as well as the final five periods.

Huck et al. (1999) and Offerman et al. (2002) show that the information subjects receive regarding demand and cost conditions or other subjects' output and profit affects the learning rules players adopt. Offerman et al. suggest giving subjects more information about the other players leads to lower cost decision rules such as mimicking successful others. Reducing this information about others forces subjects to use higher order reasoning such as best response or 'lead and follow' rules. These more sophisticated strategies are more likely to result in Cournot-Nash or collusive outcomes.

Our results show that even aggregate outcome information results in more competitive outcomes on average than no information. However, this is a relative result, and outcomes overall still converge to the Nash solution. If we had additionally provided information on the competitor's profits we might have observed more competition emerging. That would be in line with the findings of Offerman et al. (2002).

Figure 6 shows the trend lines for the difference between subjects' actions and imitation (difIM) and between actions and best response (difBR). Imitation explains the data better in T1 compared to player 1 in T3 and compared to player 2 in the early rounds of T3. In all cases, best response strategies describe the data better than imitation. Strictly speaking, player 2 cannot imitate in T3 since she does not know aggregate output. Therefore, for player 2 it only makes sense to look at the difBR trend.

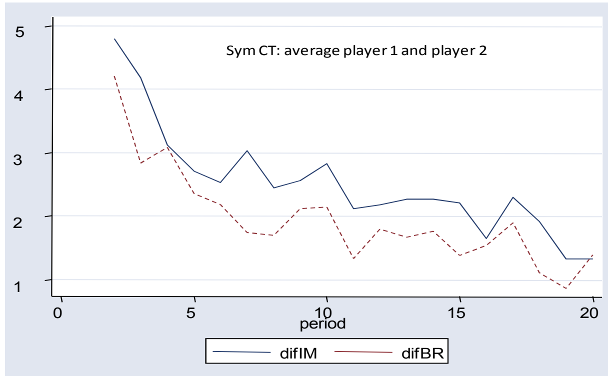

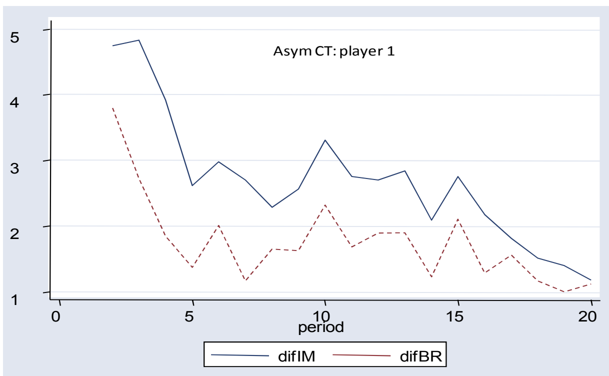

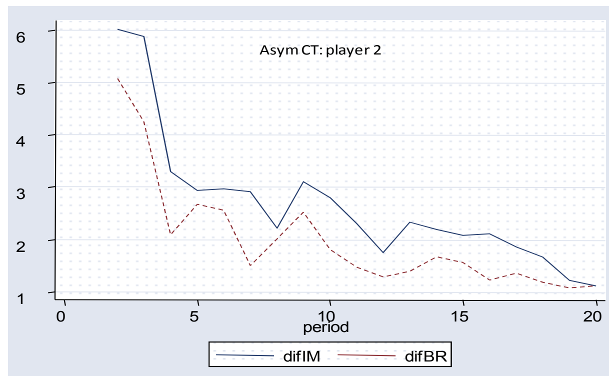

Figure 6: difference between players' actions and imitation or best-response strategies

4.6 Relationship between messages and actions (This section is still in progress)

4.6.1 Collusion, pre-emption or coordination?

For T2 we can compare player 1's message to her own action; 2's message to own action; and j's message to i's action. For T3 we compare 1's message to own action and 2's action. In most cases, players' indicate (through their messages) they intent to produce more than they in fact do. We let m_i represent player i's message, and x_i represent player i's action. Figure 3 shows the relationship between messages and actions in T2 and T3, and table 7(a and b) shows p-values for the relevant, average comparisons. In all cases, the test-statistic is positive and significant, indicating the message sent is greater than the action taken.

If players were trying to cheat on a collusive agreement, the test-statistic would be negative: if player 1 is producing 3, player 2 can make more by producing 4. If subjects state they will produce more than they in fact do, it suggests an attempt at pre-emption. If player 1 claims she will produce 6 when the equilibrium total supply is 8, she is telling player 2 she expects the higher surplus. The fact that player 1 then produces less than 6 implies she understands player 2's incentive to use the same strategy (at least in T2). In T2 neither player possesses the "power" to sustain the higher proportion of the output. The difference in T3 is that both players are not able to send a message, so player 1's message could serve as a focal point. From Table 7 this does not appear to be the case, however, since both players produce less than the message player 1 sends. In both T2 and T3, player 1 uses messages more truthfully in the last 5 periods. This is not true for player 2.

Unfortunately, we cannot disentangle the motives of collusion versus pre-emption or coordination using aggregate data. However, in T3 we have already seen that, on average, player 1 uses its superior knowledge and its ability to unilaterally communicate to obtain a greater share of the surplus. Given the discreteness of the choice space we adopt, the Stackelburg prediction is hardly different from the Cournot outcome. Thus, if we observe output choices (x_1, x_2) = (5, 3) we could interpret this as minimal pre-emption. Using histograms to look at modal responses, we do not find evidence of this pattern. If we condition the data on x_1 = 5, the modal x_2 response is also 5 (29%); player 2 chooses x_2 = 3 or 4 13% of the time each. For player 1 we find (x_1 = 5 | m_1 = 5) = 71% and (m_1 = 5 | x_1 = 5) = 58%. Since the modal message for player 1 overall is m1 = 5 (19% of all messages), it appears that player 1's attempt to pre-empt, but player 2's do not allow it.[7] With this rudimentary analysis, we do not see strong evidence of either (successful) pre-emption or collusion in T3. Further analysis at the individual group level may yield stronger results.

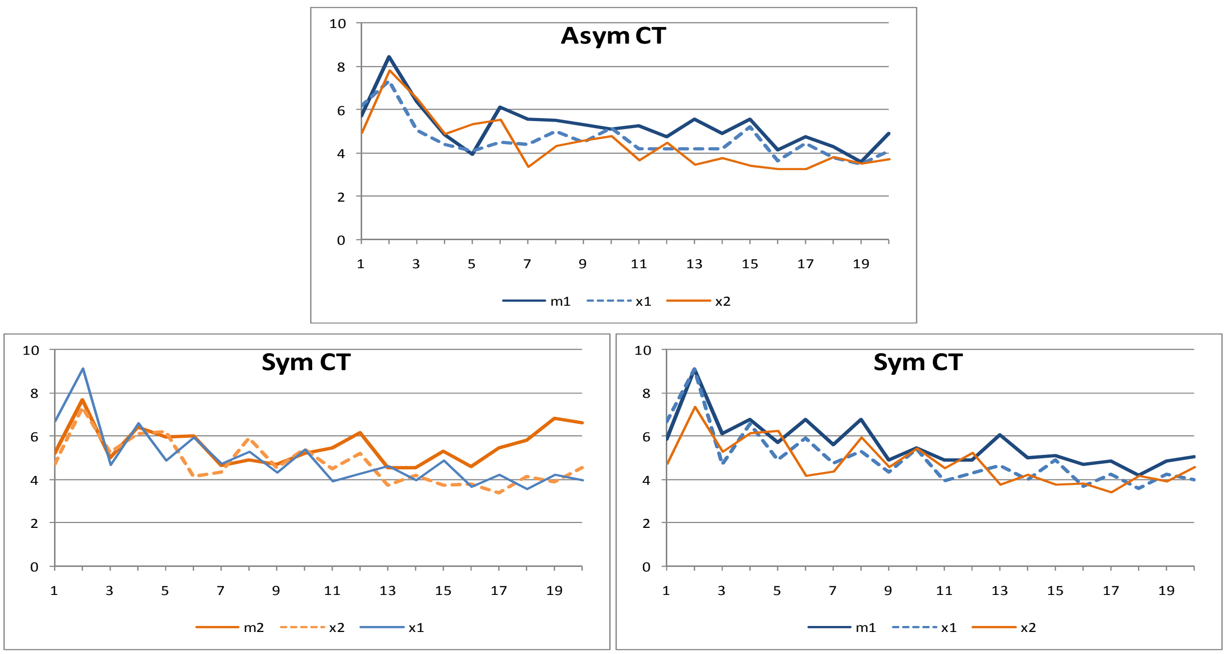

Figure 5: average messages and actions across periods.

T2	x_1	x_2
m_1	.0001	.0002
m_2	.0771	.0040

T3	x_1	x_2
m_1	.0018	.0011
m_2	na	na

Table 7a: p-values; m_i ≡ i's message, x_i ≡ i's action.

[7] In discrete space, player 2's best response to x1 = 5 is x2 = 4. Player 1 is better off at (5,3), and both players prefer (4,4) to (5,4).

T2	x_1	x_2
m_1	.2411	.0447
m_2	.0040	.0078

T3	x_1	x_2
m_1	.2140	.0714
m_2	na	na

Table 7b: p-values; $m_i \equiv$ i's message, $x_i \equiv$ i's action. Last 5 periods

4.6.2 Credibility

Addressing credibility in the context of this game is not straightforward since the strategy space is large. Most often credibility is addressed in the context of "small" normal form games. We measure credibility very crudely by looking at players' modal responses to messages of 3 and 4. Tables A1 and A2 in the appendix show the message-action space for messages 3-6 in T3 and for (matching) messages 3-5 in T2.

In T3, we should find that (i) $m_1 = 3$ is followed by $x_2 = 4$, (ii) $m_1 = 4$ is followed by $x_2 = 2$, and (iii) $m_1 = 5$ is followed by $x_2 = 4$. When $m_1 = 3$ or 4, the modal response by player 2 is four, as expected (32% and 34% respectively); when $m_1 = 5$, however, player 2 also plays five 29% of the time.

Since player 1 tends to be truthful when sending a message of 3, it could either be that player 2 is attempting to cheat on the collusive outcome (or that she expects player 1 to cheat). It seems more likely, however, that players understand that within some range of total output, they are better off producing relatively more. They are just unclear on what the upper bound of that output range is.

In T2, players appear better able to use message to coordinate. Given that both players send the same message, they are able to collude following a message of three; they play Nash following a message of four, and they are competitive following messages of five. This is a liberal conclusion, however, since we are conditioning on matching messages.

As a preliminary conclusion we fail to reject H6 in T3. That is, the data suggests cheap talk does not lead to collusion.

5. Conclusions

We use experiments to look at the convergence of prices in oligopoly markets with different information feedback between players. The model we address is one that would lead to investment cycles if firms were myopic. We ask if cheap talk helps firms break out of cycles as quickly as having full information about supply and demand conditions. We consider two cases of cheap talk. In one case firms have symmetric information and send non-binding output messages simultaneously. In another case, one firm is better informed about market supply and can unilaterally send an output message to the other firm. Our main contribution is our use of the coefficient of variation to show that that price cycles are broken more quickly with no information and symmetric cheap talk than with full information.

We also find providing more information leads to more competitive outcomes *on average*. While this result is surprising it is line with some of the literature on learning in oligopoly games. As discussed in the text, others likewise find that providing more information to subjects provides opportunities for low cost strategies such as mimicking which tend to converge to competitive outcomes. We note, however, that this increase in competition is strictly relative and that prices still converge to the Cournot-Nash equilibrium. Also, the Cournot Nash outcome is the modal outcome with full information. Our experiment was not designed to test various learning models but we were able to ask if subjects in the full information and asymmetric information treatments were using best response or imitation strategies. In both cases best response strategies fit the data better than imitation; imitation does better in T1 than T3.

Even though group size is just 2 in our experimental subjects are not able to maintain collusion in any of the treatments. Our experiments are short in comparison to similar experiments, however, and it is possible that collusion may have been achieved with more rounds. On the other hand, we may have simply observed more volatility. Prices were highest in the asymmetric treatment and remained above the Cournot Nash equilibrium for most of the periods. We find that the informed player receives a greater share of profits with asymmetric information.

It is not entirely clear how subjects use messages. On average messages in both cheap talk treatments are above the action taken. While this suggests subjects may be attempting to pre-empt their competitor we do not find strong evidence in support of such behavior when we look at the data on a more disaggregated level. We do find support for the theoretical prediction in the asymmetric cheap talk treatment that messages of 3 (the collusive action) are not credible.

There a several directions for further research. The most obvious would be to compare the convergence properties of T0, T1 and T2 under unstable supply and demand conditions. The model is also adapted to allow for a multi-dimensional choice space, so subjects could be allowed to choose both the number of plants and the output per plant. Larger group sizes would allow one to look at convergence to the competitive equilibrium in the true spirit of the cobweb model; the theory for large groups and cheap talk is not established, however.

References

Aumann, R. (1990). "Nash Equilibria are not Self-Enforcing," in Economic Decision Making: Games, Econometrics and Optimization: Essays in Honor of Jacques Dreze, J.J. Gabzewicz, J-F. Richard, and L. Wolsey, pp. 201-206, Elsevier.

Baglia, S. and S. Morris (2002). "Co-ordination, Spillovers, and Cheap Talk," *Journal of Economic Theory*, 105, pp. 450-468.

Carlson, (1967). "The Stability of an Experimental Market with a Supply Response Lag," *Southern Economic Journal*, 33, 305-321

Cox, J. and M. Walker (1998), "Learning to Play Cournot Duopoly Strategies," Journal of Economic Behavior and Organization, 36, no. 2.

Dawson, J. and A. Sen (forthcoming) "New Evidence on the Convergence of International Income from a Group of 29 Countries," *Empirical Economics*.

Farrell, J. (1988). Communication, coordination and Nash equilibrium, Economics Letters, 27:3, 209-14.

Fischbacher, U., 1999. z-tree: Zurich toolbox for readymade economic experiments- experimenter's manual. Working Paper, University of Zurich.

Friedman, M. (1992) "Do Old Fallacies Ever Die?" *Journal of Economic Literature*, Vol. 30. No. 4, pp. 2129-2132.

Huck, S., H. Normann and J. Oechssler (1999) "Learning in Cournot Oligopoly – an experiment," *The Economic Journal*, vol. 109 no. 454.

Huck, S., H. Normann and J. Oechssler (2000) "Does information about competitiors' actions increase or decrease competition in experimental oligopoly markets?" *International Journal of Industrial Organization*, 18 pp. 39-57.

McDaniel, T.M. (2008) "Strategic use of cheap-talk with incomplete information: experimental results." Working Paper, Appalachian State University.

Nieswiadomy, M. and M. Strazicich (2004) "Are Political Freedoms Converging?" *Economic Inquiry*, April 2004, Vol. 42, No. 2, 323-340.

Nerlove (1958) “Adaptive Expectations and Cobweb Phenomena”, *Quarterly Journal of Economics*, Vol. 72, No. 2. pp. 227-240

Offerman, T., J. Potters, and J. Sonnemans (2002), “Imitation and belief learning in an oligopoly experiment,” *The Review of Economic Studies*, vol. 69, no. 4, pp. 973-997.

Appendix

Table A1: messages 3 – 6 and corresponding actions for T3.

m1 = 3				*x2*			
x1		**1**	**2**	**3**	**4**	**5**	**Total**
	1	0	0	2	0	0	2
	2	2	0	0	6	4	12
	3	6	20	18	**24**	4	72
	4	0	0	4	2	0	6
	5	0	2	4	2	4	12
	6	0	0	0	0	2	2
Total		8	22	28	34	14	106

m1 = 4			*x2*										
x1		**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**9**	**10**	**18**	Total
	0	0	0	0	0	2	0	0	0	0	0	0	2
	1	0	0	0	0	2	2	0	0	0	0	0	4
	2	2	0	0	2	0	0	0	0	0	0	0	4
	3	0	0	0	0	2	2	2	0	2	2	2	12
	4	0	2	**20**	12	**20**	12	10	2	0	0	0	78
	5	0	2	0	2	16	2	0	0	0	0	0	22
Total		2	4	20	16	42	18	12	2	2	2	2	122

m1 = 5				*x2*								
x1		**1**	**2**	**3**	**4**	**5**	**6**	**7**	**9**	**10**	**13**	**Total**
	1	0	0	0	0	2	2	0	0	0	0	4
	2	0	0	0	0	4	4	2	0	0	0	10
	3	2	0	0	4	0	0	0	0	2	0	8
	4	0	4	2	0	4	0	2	0	0	0	12
	5	6	12	16	14	**26**	6	6	2	8	2	98
	6	0	0	0	0	4	2	0	0	0	0	6
Total		8	16	18	18	40	14	10	2	10	2	138

m1 = 6				*x2*					
x1		**1**	**2**	**3**	**4**	**5**	**6**	**11**	Total
	2	0	2	4	0	0	0	0	6
	3	0	0	2	0	0	4	0	6
	4	2	0	0	4	0	0	0	6
	5	0	0	0	0	4	2	0	6
	6	6	2	4	4	**8**	4	0	28
	7	0	0	0	0	0	0	2	2
Total		8	4	10	8	12	10	2	54

m1 = m2 = 3

		x2		
x1	**3**	**4**	**5**	Total
3	24	6	6	36
4	12	2	0	14
5	2	0	2	4
Total	38	8	8	54

m1 = m2 = 4

		x2		
x1	**4**	**5**	**6**	Total
2	2	0	0	2
3	0	2	0	2
4	14	0	2	16
5	8	0	0	8
Total	24	2	2	28

m1 = m2 = 5

			x2	
x1	**4**	**5**	**6**	Total
4	0	4	0	4
5	2	14	2	18
Total	2	18	2	22

Table A2: messages and actions for m1=m2