Strategic Delegation in an Experimental Mixed Duopoly

Ninghua Du*

John S. Heywood**

Guangliang Ye***

**Associate Professor, School of Economics, Shanghai University of Finance and Economics, 777 Guoding Road, Shanghai 200433 P. R. China* ninghua.du@mail.shufe.edu.cn

***Distinguished Professor, Department of Economics, University of Wisconsin- Milwaukee, P.O. Box 413, Milwaukee, WI 53201 USA heywood@uwm.edu*

****Professor, Hanqing Advanced Institute of Economics and Finance, School of Economics, Renmin University of China, 59 Zhongguancun Street, Beijing 100872 P. R. China gye@ruc.edu.cn Tel: 86-10- 62514797 Fax: 86- 10- 62511343 (Corresponding Author)*

Abstract: We provide the first experimental test of the consequences of delegation in a mixed duopoly. Such delegation allows a profit maximizing private owner and a welfare maximizing public owner to weight sales in managerial contracts. Theory predicts that such contracts improve welfare. Our evidence indicates that both public and private subject owners do provide a weight on sales consistent with the subgame perfect equilibrium. Critically, however, this emerges in the experiment only when playing with a robot manager and when playing with a human subject manager after the experience of playing with a robot manager.

Keywords: Mixed Duopoly, Incentive Contract, Experiment Markets

JEL Codes: H11, L13, L32

Acknowledgements: The authors thank Vivian Lei for detailed comments on an earlier draft.

1. Introduction

This paper presents the first experimental test of delegation in a mixed duopoly. A mixed oligopoly consists of a public firm presumed to maximize welfare competing against private firms that maximize profit.[1] Since Merrill and Schneider (1966) an enormous theoretical literature isolates when such a public firm disciplines a private oligopoly and so improves welfare, the potential to "regulate by participation." In basic Cournot models, the presence of a public firm increases welfare when the number of private firms is small (DeFraja and Delbono, 1989). As a base line, we test whether laboratory participants make choices consistent with the Cournot mixed duopoly model. Given our evidence that they do, we test the consequences of delegation.[2]

The literature on strategic delegation builds from the separation of ownership from control to examine the incentive contracts that owners provide managers. In private oligopolies owners typically provide their managers a contract that rewards both profit and revenue (Vickers, 1985; Fershtman and Judd, 1987; Sklivas, 1987). Rewarding revenue causes the manager to produce more than the simple Cournot quantity with the expectation that rivals will reduce their quantities. Yet, what would be profitable for a single manager does not carry over to the market as when each owner provides such a contract, total output increases and profit for each firm decreases relative to a Cournot market without delegation.

Delegation in a mixed duopoly has been shown to generate similar consequences (Barros 1995; White 2001; White 2002; Fernandez-Ruiz 2009; Heywood and Ye 2009). Both the public and private owners have an incentive to provide their managers with a contract that rewards additional output because of the strategic substitutability.[3] When the managers adopt quantities based on these contracts, total output and welfare both increase. While the private owner suffers reduced profit, the increased welfare enjoyed by the public owner suggests that governments should encourage such delegation. However, this policy conclusion depends on owners and managers behaving in the fashion outlined by the delegation model.

While we provide the first experimental test of delegation in a mixed duopoly, Huck et al (2004) conduct an experiment of delegation in a fully private Cournot duopoly. In their experiment owners choose between a contract that maximizes profit and a more aggressive contract that weights sales as predicted by the subgame perfect equilibrium (SPE). In a result the authors identify as "surprising," the owners rarely choose the subgame perfect prediction. It is important to note that while the SPE can result in ex post harm for owners in the private oligopoly (their profits are lower when both managers are told to weight sales), the opposite is true for the public owner in the mixed oligopoly (welfare is higher when both managers are told to weight sales). Thus, the mixed duopoly differs from the private duopoly in that there is no incentive for the two firms to collude eliminating one source of potential strategic interaction. This difference partially motivates our examination. The absence of the usual prisoner's dilemma tension, "should I collude or cheat," implies that we can more directly test the logic of the backward induction inherent in strategic delegation.

We recognize testing stage games and the resulting subgame perfect equilibrium can be challenging. Binmore et al. (1985) first conducted alternating-offer sequential bargaining experiments in a shrinking pie model to test the SPE prediction. They found that the opening offers lie between equal splits and the selfish SPE prediction. The experiments by Neelin et al. (1988) and Ochs and Roth (1989) get similar findings with different parameter structures. Johnson et al. (2002) return to such alternating-offer experiments finding a role for social preference (regard for the fellow player) in explaining deviations from SPE. They show that human subjects move closer to the SPE when they knowingly bargain with selfish and rational robot players. Moreover, the subjects with experience playing against the robot players remained closer to the SPE when they later played against other human subjects. Despite evidence of social preferences, Johnson et al. (2002) also suggest that computation difficulty may make subjects less likely to complete backward induction. We bring these insights to our study of delegation. We first conduct a trial with inexperienced owners setting contracts for human subject managers. We next conduct a trial in which inexperienced owners set contracts for robot managers. Finally, we conduct a trial in which those owners with experience setting contracts for the robot managers do so for human subject managers.

We confirm the basic model of the mixed oligopoly without delegation. We do so in a simple model that ignores issues of multiple and boundary equilibria that have proven difficult in private duopoly games (Cox and Walker, 1998). Next, we introduce delegation and inexperienced owners play against actual human subject managers. The private owners significantly overweigh sales relative to the SPE. When inexperienced owners play against robot managers, both the public and private owners adopt weights consistent with the SPE. This carries over when the owners with robot manager experience play with actual human subject managers.

We take these results as broadly confirming the predictions of the delegation model in mixed markets but note that the same failure evident in Huck et al (2004) seems evident for our private sector owners in our initial delegation treatment. At issue is the cause of this initial failure. We note that in the robot treatment both the uncertainty associated with manager's choices and the possibility of owner's having social preferences toward managers are eliminated. With these complications removed, the SPE is achieved. In the subsequent experienced owner treatment the owners' retain information about the backward induction from the robot treatment but have the possibility of social preferences reintroduced. As this does not alter their choice of the SPE, we suggest that the original failure to adopt the SPE is more likely related to the uncertainty introduced by the managers' choices and the difficulty introduced in solving the backward induction rather than reflecting social preference.

As key sectors of both transition and developed economies remain mixed oligopolies (Bortolotti et al. 2004), it is highly relevant that we test whether or not laboratory subjects behave as implied by the implicit two-stage game of delegation in a mixed duopoly. Among the many examples, China Telecom remains government owned and faces private rivals. In India, banking reform started in 1992 encouraged private entry that has eroded but not eliminated the previously dominant public banks (Bhaumik and Dimova, 2004). The US Post Office competes with UPS and FedEx in the US parcel delivery market. In Germany, large private banks compete with small local publicly owned banks while air travelers in Europe and Asia currently choose between a variety of public and private airlines.

In what follows, the next section describes the underlying theory. The third section explains the experimental design. Section four presents the basic experimental results. Section five provides further statistical tests and section six draws conclusions.

2. The Theoretical Model to be Taken to the Laboratory

This section briefly details the model of delegation in a mixed duopoly. One public firm indexed by A and one private firm indexed by B compete in a domestic market. The private owner maximizes profit and the public owner maximizes welfare. Managers maximize objective functions as specified in their incentive contracts. The firms produce a homogeneous good with identical quadratic cost functions typical of the literature: $C(q_i) = q_i^2, i = A, B$. An linear inverse demand curve, $P = 16 - q_A - q_B$, yields firm profits:

$$\pi_i = q_i(16 - q_A - q_b) - q_i^2 \quad , \quad \text{i}=A,B \tag{1}$$

Social welfare, the sum of consumer surplus and total profits, is the public owner's (the government's) objective function:

$$W = CS + \pi_A + \pi_B \quad \text{where} \quad CS = (q_A + q_B)^2 / 2 \tag{2}$$

Following Barros (1995), owners offer their managers an incentive contract that is a linear combination of profit and sales revenue:[4]

$$IC_A = x\pi_A + (1-x)Pq_A \text{ or } IC_B = y\pi_B + (1-y)Pq_B \tag{3}$$

In the first stage owners play a game over the incentive proportions x and y. In the second stage managers play a Cournot quantity game given their respective incentive contracts. Backward induction yields the SPE. Given x and y, the quantity game in stage two results in:

$$q_A = \frac{16(2y+1)}{4xy+4x+4y+3} \quad \text{and} \quad q_A = \frac{16(2x+1)}{4xy+4x+4y+3} \tag{4}$$

Substituting these outcomes into profits and welfare, the owners simultaneously choose x and y to maximize their respective objective functions. The SPE yields $x^* = 0.57$ and $y^* = 0.84$ implying that $q_A^* = 4.06, q_B^* = 3.24$ and $W^* = 63.17$.

It is this model that is moved to the laboratory and we now turn to describing the design of the experiment.

3. Experimental Design

We conducted our experiments at the Experimental Economics Laboratory at the Shanghai University of Finance and Economics (SHUFE). Participants were recruited from a campus-wide list of undergraduate students who responded to advertisements in courses or on the web. Most participants are junior students with knowledge of calculus. The participants were paid a 5 Yuan show-up fee as well as their earnings during the experiment. To make the payoffs symmetric across participants, the earnings were counted by experimental dollars and the exchange rate between experimental dollars and yuan (RMB) differed with the treatment and the role as explain in more detail later. Moreover, the owners' payoff functions were normalized as follows: $[16 – (q_A + q_B)] (q_A + q_B) – 0.5(q_A – q_B)^2 – 60$ for public owners, and $[16 – (q_A + q_B)] q_B – q_B^2 –14.5$ for private owners. The choices of q_A and q_B are limited between 0 and 8 (0 and 8 included), and are rounded to two decimal places.

The subjects' average earnings were 35 Yuan for 1 hour and 40 minutes. The typical pay of a part time job on campus for an undergrad (say, tutoring) was approximately 15 Yuan per hour (1 USD=6.85 yuan). All laboratory sessions were computerized using “z-tree” software (Fischbacher 2007). Sample translated versions of the instructions are provided in the

Appendix. Note that the names for owners and managers remain A or B in order to avoid any framing effects caused by the associations that subjects may have with the terms "state owned" or "privately owned" firms.

There are four treatments as briefly outlined in the introduction. There is one session in the No Delegation (ND) treatment. This session is designed to make sure that the basic mixed oligopoly predictions tend to hold before introducing strategic delegation and the implied stage game. In the session, 16 participants are divided into 8 markets with fixed parings and fixed roles. In each market, the public firm competes with the private firm in outputs. Each market repeats the competition for 20 rounds. The advantages of fixed pairings include the ability to test for learning and increasing the number of observations given the limited number of participants. There are no managers in these treatments so owners simply play a repeated single stage mixed oligopoly quantity game given the demand and cost structures outlined in the previous section. In the ND treatment 2 experimental dollars is equal to 1 yuan (RMB) for both the public and the private owner. The equilibrium predicted outputs and welfare are $q_A^* = 4.36$, $q_B^* = 2.91$ and $W^* = 62.41$.

New participants are drawn and two sessions (indexed by IO1 and IO2) of the Inexperienced Owner with Human Manager treatment are next conducted. This serves as our initial test of the SPE with delegation. There are 16 participants in each session, who play the game for 20 rounds. Participants are randomly divided into 4 markets. In each market, 4 participants each take the role of public owner, public manager, private owner and private manager. The matching and roles are fixed during all rounds. For the owners 2 experimental dollars is equal to 1 yuan and 16 experimental dollars is equal to 1 yuan for the managers.

Participants know the exchange rate could be different by role but they only observe their own exchange rate. Each round lasts for 5 minutes, 3 minutes for the owners and 2 minutes for the managers.

Participants are provided a “payoff calculator.” By entering a tentative *x* and tentative *y*, the owners can observe the equilibrium q_A and q_B in the corresponding subgame and both owner’s payoffs given the equilibrium q_A and q_B. Given the public owner’s choice *x* and the private owner’s choice *y*, by entering the tentative outputs in the payoff calculator, the managers can see their corresponding payoffs. The players are provided the payoff functions derived from the demand and cost functions.[5] The owners’ choices of x and y are between 0 and 1 and are rounded to two decimal places. Importantly, Requate and Wichman (2011) confirm that experimental results in duopoly market games are not influenced by the choice of using a payoff calculator or providing a complete payoff table. Nonetheless, we note that the calculator may help subjects explore optimizing choices by providing subgame results.

Next there are four sessions (indexed by RM1, RM2, RM3 and RM4) in the Robot Manager treatment. In this treatment owners are told that the robot managers produce optimally given the incentive parameters. The RM treatment serves to encourage learning about the backward induction. It does so by providing practice in a session that removes both uncertainty and social preference. The type of social preference we remove is between the owner and his or her manager. While we can't be sure what this preference will look like, we've placed them on the "same team" suggesting that social concerns might get in the way of solving the backward induction. Examples of social preference include distribution based concerns such as inequity aversion or altruism and intentional concerns such as reciprocity or guilt aversion. Camerer and Fehr (2006) provide descriptions and evidence on the role of such concerns.

This RM treatment also uses new participants assigned randomly to a session. There are 8 participants in each session, who play the game for 20 rounds. In each round, participants are randomly divided into 4 markets with fixed pairings and fixed roles. In each market, one takes the role of public owner while the other takes the role of the private owner. Both the public manager and the private manager are simulated by computer (i.e., robot managers). In each round, given the public owner's choice x and the private owner's choice y, the robot managers automatically provide the Nash equilibrium outputs in the corresponding subgame. This is described in the instructions and the experimental software again provided each participant a "payoff calculator." If one enters the tentative public owner's choice x and the tentative private owner's choice y in the payoff calculator, one can read the corresponding q_A and q_B by the robot managers and both owner's payoffs given the robot managers' outputs. Each round lasts 3 minutes. In the RM treatment 2 experimental dollars equals 1 yuan (RMB) for both the public and the private owner.

Finally there are two sessions (indexed by EO1 and EO2) in the Experienced Owner with Human Manager treatment. The arrangements are identical to those of the Inexperienced Owner with Human Manager treatment except that the owners are participants who had completed the Robot Manager treatment (RM3 and RM4). EO1 was two days after RM3 and EO2 was two days after RM4. The managers in the Experienced Owner treatment are new participants who do not know that the owners had completed the RM treatment. We examine how the experiences of the robot round might persist as these experienced owners move to later engage human subject managers. In particular, this treatment reintroduces the possibility of social preferences but it does so after the robot round has allowed the owners to more easily engage in the backward induction implied by the model.

We recognize that the finitely repeated game that we implement is not identical to the one time game model in Section 2. Yet, as the participants know the number of rounds in the experiment, the SPE in each round will not be different from the equilibrium in the single round game and folk theorem type results should be avoided. Perhaps even more important, we stress again that in the mixed oligopoly there is no incentive for the firms to cooperate, as there would be in the private oligopoly. Quantity restrictions increase the profit of the private firm even as they decrease the welfare objective of the public firm. Thus, we feel confident that differences across rounds represent learning not dynamic incentives.[6]

4. Experimental Results

The choices by participants in each experimental market are interdependent and these shared experiences could be carried over periods. To control for this issue we treat each group of fixed pairings in each session (i.e., each experimental market) as one independent observation from which we take the mean across rounds. We move through the results following the treatments outlined in the previous section. We first examine the basic mixed duopoly game, the no delegation treatment. We contrast this with the stage game with inexperience owners, IO. We then create the experienced owners treatment by first examining the treatment with robot managers and then replacing the robot managers with human subject managers while retaining the same owners.

Table 1 presents the means and standard deviations associated with each treatment.

<Table 1 about here>

Our formal hypothesis tests begin in the next subsection.

4.1 The Mixed Duopoly Treatment without Delegation

In this subsection, we test whether the experimental results of the mixed duopoly without delegation fit the theoretical predictions. We first check that the production of the public firm exceeds that of the private firm as predicted by theory. This hypothesis test is shown in the first row of Table 2. The hypothesis that the two firms have the same output is easily rejected against the one-tail alternative that the public firm produces more. The test is the non-parametric Wilcoxon signed-ranks test used to compare two related samples to assess whether their population mean rank differs. This paired difference test can be interpreted as having a null that the measure of central tendency of one distribution is equal to that of another distribution.[7]

Second, we test the hypothesis that the actual values chosen by the public firm and private firm subjects equal their hypothesized equilibrium values. We use the same sign test by comparing the distribution of experimental values across the subjects to a degenerate distribution equal to the hypothesized value.[8] The second column shows that the mean of the no delegation model for the public firm is only 0.14 away from the equilibrium and that for the private firms is only 0.07 away from the equilibrium. The statistical tests cannot reject equality.

<Table 2 about here>

Finally, we test the hypotheses that the total welfare in the experimental markets equals that predicted by the no delegation model. As row 4 shows, the difference between the mean experimental welfare and that in equilibrium is only 0.31 on a base of over 62 and the hypothesis of equality cannot be rejected. In short, the no delegation model seems to move successfully to the experimental laboratory and we will use this as a benchmark when considering the innovation of allowing delegation.

4.2 Delegation with the Inexperienced Human Owner

We now consider the results of our delegation stage game. Recall our initial treatment considers inexperienced owners who play against each other and with human managers. We make a series of comparisons both with the equilibrium predictions and with the previous treatment without delegation. As will be seen, the results are generally not a good match to the theoretical predictions.

Without experience, the output of the public firm continues to exceed that of the private firm but the difference misses statistical significance. Moreover, when examining the owner's choice of the incentive parameter, both owners are overly aggressive in rewarding sales compared to the model predictions. The public owners average an incentive parameter on profit that is 0.05 below that predicted while the private owners average an incentive parameter that is an enormously large 0.20 below that predicted. In each case the signed-ranks test rejects equality of the measure of central tendency with the hypothesized value.

<Table 3 about here>

Despite the lower than predicted incentive parameters, the public managers adopt output levels that are insignificantly different from that predicted by the model. This contrasts sharply with the private owners who adopt output levels significantly above those predicted by the model. As a consequence, the total output in the market and the total welfare exceeds that predicted by the market and the difference is statistically significant. Thus, relative to the model's predictions, the owners are more aggressive which should result in over production relative to predictions but, instead, only private managers over produce relative to predictions. Indeed, as shown in Row 7, even relative to the overly aggressive incentive parameters, the private manager over produces. Specifically, given the owner chosen values of x and y, the private manager chooses an output on average 0.17 above what would be predicted. The experimental markets do not do a good job mimicking the theoretical predictions. This happens despite the presence of the earnings calculator designed to help with computational difficulty inherent in the stage game. Thus, while the model of mixed duopoly performed well without delegation, the structure of the stage game and its continued complexity or perhaps social preferences yields a much less successful fit with the model with delegation.

In the final three rows of Table 3 we test the experimental results from the delegation game directly against those from the experimental model without delegation. As a result of the managers in the delegation game overshooting the predicted outputs, we anticipate that the outputs will also be greater than those in the game without delegation. The three tests reject the hypothesis that the experimental markets with and without delegation yield the same values. Note that these tests are across treatments and so we use the alternative Wilcoxon-Mann-Whitney rank sum test. This is preferable as it does not assume paired observations that are drawn from a distribution such that the median difference is zero.

Thus, the experimental results are consistent with very few of the predictions of the delegation model. The basic presumption that the public firm will produce more does not emerge. Moreover, the results show owners adopting incentive parameters that are too small and aggregate output that is too large relative to the model predictions.

4.3 Delegation with Robot Managers

We now begin the process of creating experienced owners. We start by drawing new subjects for owners and by having them play the delegation game with what they are told are robot managers. The robot managers optimize the objective function given by the owners' incentive parameters. Thus, the results of the earlier delegation game in which managers over produced given the predicted incentive parameters could not happen. More fundamentally, playing with robot managers eliminates a source of uncertainty and potential complexity. It also eliminates any issue of social preference of the owner for the manager. At issue is how these changes might alter the experimental market and whether any changes that do result might persist as these experienced owners move to later engage human subject managers.

The key comparisons of the experimental market and the predicted equilibrium are shown in the top three rows of Table 4. The results appear mixed. On the one hand, the greater number of observations still allows rejection of the hypothesis that the experimental markets have reproduced the predicted equilibrium values as each incentive parameter and welfare are significantly different from its predicted value.

<Table 4 about here>

One the other hand, it appears clear that the chosen values have move much closer to the predicted values. The public owner is predicted to choose 0.57 but chose 0.52 in the inexperienced owner retreatment and 0.55 in this robot round. More dramatically, the private owner is predicted to choose 0.84 but choose only 0.64 in the inexperience owner treatment and then 0.86 in the robot round. As the final two columns make clear, this last change represents a significant difference between the inexperienced owner treatment and the robot round. The owners (especially the private owner) thus do appear to be moving closer to the predicted values when playing with robots.

4.4 The Experienced Owner Treatment

In this treatment we retain the owners from the robot treatments RM3 and RM4 but they now play against human subject managers. These subject managers are new participants who have no experience. We examine whether or not any of the progress toward the predicted values remains. The tests against the hypothesized values are in the top rows of Table 5.

We again start with the most basic hypothesis that the output of the public firm equals that of the private firm. As in the robot treatment but not in the inexperienced owner treatment, this hypothesis can be easily rejected supporting the theoretical predictions. Turning to the incentive parameters, both remain close to their predicted value. The experimental result for the public owner averages 0.53 and has shown only modest variation across treatments. The ranked sum test suggests a statistical difference between the experimental results and the equilibrium prediction but only at the 10 percent level. The experimental result for the private owners averages 0.83, essentially identical to the result from the robot round and the ranked sum test indicates no difference between the experimental result and the prediction.

<Table 5 about here>

The quantities chosen by the public firm human subject managers are very close to the predicted equilibrium and the hypothesis test cannot reject equality. The private managers produce more than the equilibrium prediction and the hypothesis test of equality can be rejected albeit only at the ten percent level. As shown in Row 7, this results largely from producing more than implied by the owner chosen incentive parameters. Finally, the welfare resulting from the experimental markets appears very close to that predicted and the hypothesis of equality cannot be rejected.

We now turn to a series of tests comparing this treatment with experienced owners to the first delegation treatment with inexperienced owners. These are shown in the last five rows of Table 5. The incentive parameter of the public owner has generally shown only modest change across all treatments with the mean difference between the inexperienced and experienced owner treatment being only 0.01. The hypothesis test cannot reject equality. On the other hand, the experienced private firm owners adopt a parameter that matches that predicted and the hypothesis test indicates it is significantly greater than that adopted by the inexperienced private firm owners.

The experienced public firm managers choose outputs that are insignificantly different from those of the inexperienced public firm managers. However, the experienced private firm managers choose mean outputs, 3.41, closer to those predicted by the model, 3.24, and outputs that are significantly smaller than those chosen by the inexperienced private firm managers, 3.82. Finally, welfare in the experienced owner treatment is not only insignificantly different from the predicted but also significantly different from that emerging from the inexperienced owner treatment.

The results show that the public owners' choices and the public managers' choices essentially do not vary across treatments. The critical difference that emerges is for private owners. This raises a potential concern in experimental design that the payoff function for the public owner may be too flat to encourage optimization (see Drago and Heywood 1989 and Harrison 1989 among others). The usual concern is that deviations from equilibrium predictions do not indicate the failure of the theory but the flat payoff function and too small an incentive to optimize. In our case, the public owners appear to routinely optimize across treatments but that optimum is rather close to one-half, the value in the middle of the range, and on that might be chosen by subjects with too little incentive to optimize. As a check, we have graphed the payoff function for the two types of owners in Figure 1. We examine the influence of deviations from the respective optima on the payoffs of each type of owner assuming the other properly optimizes. After simply shifting vertically the owners' payoff functions to have the same optimum, the pattern is very clear that the public owners have more, not less, incentive to optimize. The public owners have a steeper payoff function making it less likely they are simply adopting one-half to avoid calculation effort. Indeed, the steeper payoff function may help explain their more routine optimization.

<Figure 1 about here>

When comparing across the treatments we then focus on players other than the public owner. When private owners play with either robots and or with human subjects after gaining the robot experience, they improve their performance toward the SPE and adopt the incentive parameter predicted by the model. For private managers, playing in the round with experienced owners helps them improve toward the SPE. We emphasize the similarity of the robot round and the human subject round that followed. The human subject round reintroduces the potential for social preferences regarding managers yet it did not change the behavior of the owners thus we think this concern is largely irrelevant in explaining the initial failure to achieve equilibrium. Instead, the evidence across the treatments suggests that the complexity of the backward induction when facing the uncertainty of managerial choices may be the source of off equilibrium choices. When new owners face the two stage problem without the uncertainty of managerial choices, they appear to successfully backwardly induct. This ability is carried into the experienced owner treatment when the uncertainty and the potential for social preferences are reintroduced.

5. Examining Deviations

In this section we briefly change focus from aggregated measures of central tendency to deviations from the model predictions. We compare measures of mean squared deviations in quantities and show the pattern of mean squared deviations across rounds.

For any variable *k*, the deviation is the difference between its realization and its predicted value, *k*-*k** and the mean squared deviation (MSD) measure is then:

$$\mathrm{MSD}(k, k^*) = \frac{1}{T}\sum_{i=1}^{T}(k - k^*)^2 \qquad (5)$$

We examine this measure for the four choice variables, where $x^*=0.57$, $y^*=0.84$, $q_A^*=4.06$, and $q_B^*=3.24$ and where T is the total number of rounds within an experimental treatment. As

Table 6 shows, the mean squared deviation is about four times larger in the inexperienced owner treatment than in the experienced owner treatment. The rank sum test shows a significant difference in the MSD between the inexperienced and experienced owner treatments is significantly different for all four choice variables. Moreover, in tests available from the authors, the dispersion (standard deviation) is routinely smaller for the experienced owner treatment. Thus, not only does the experienced owner treatment come closer to the predicted value but deviations around that value are smaller.

<Table 6 about here>

Critically, we note that the mean squared deviation for the robot treatment is even lower than for the experienced owner treatment. We suspect this reflects the reduced uncertainty inherent in the second stage choices. Yet, the fundamental point is that experience by the owners is associated with reduced mean squared deviation that reflects both less bias and dispersion.

As an alternative to taking the mean squared deviation across periods, we also take it across subjects so that it can be examined over periods. Figure 2 presents this for each of the roles and treatments. It confirms that the deviations are far smaller when the owner is experienced. It also suggests that even the inexperienced owner treatment may show smaller deviations in later rounds.

<Figure 2 about here>

Finally, we decomposed the individual quantity deviations that make up the MSD measure in Figure 2. For any round we identify a variable, $q_i(x, y)$ i=A,B as the equilibrium response of the manager given the incentive choices of the owner. Given this we can decompose the deviation into two components: $q_i - q_i^* = [q_i - q_i(x,y)] + [q_i(x,y) - q_i^*]$. The first component represents the managers' second stage production choice error and the second component represents the owners' coefficient choice error. This decomposition is available from the authors but it confirms that both owners and managers contribute to the deviation of the experimental result from the equilibrium level. Their contributions vary from round to round. Sometimes they enhance one or another and sometimes they offset each other. Generally speaking, both stages are important in deciding total deviation and both stages have smaller deviations in the experienced owner treatment.

6. Conclusion

This paper is the first to establish experimental mixed oligopoly markets and explore the consequences of strategic delegation. Theory predicts that in the duopoly the public firm produces more than the private firm and increases welfare. Theory also predicts that delegation further increases outputs and welfare as owners choose give weight to revenue as well as profit (Barros, 1995). Our evidence generally confirms these predictions while recognizing that the influences of the specific treatment and of experience can be critical.

In the one stage game without delegation, the public firm subjects produces more and, indeed, the results conform closely to the theoretical predictions. In the delegation game without experienced owners, there was little evidence confirming the theoretical predictions. The owners were too aggressive and the quantities did not match the predicted equilibrium. We then conducted a treatment in which new owners played with robot managers. The incentive parameters moved to more nearly match those predicted. The robot round reduced uncertainty, removed potential social concerns and provided experience. When those owners with the robot manager experience moved into a treatment with subject managers, this improvement was retained. The delegation game in the robot manager treatment reduces the uncertainty and complexity of the second stage choices and removes a potentially important source of social preference.

We also emphasize that the mixed oligopoly setting differs in a critical respect from the private oligopoly. In the private duopoly there exists the well-known prisoner's dilemma in which each firm would like to restrict quantity if an enforcement mechanism exists but each has an incentive to violate such a restriction otherwise. In the mixed duopoly this dynamic is absent as the incentive for the public firm is to increase quantity to enhance welfare even as the private firm wishes to restrict output to increase profit. While testing experimental mixed oligopoly markets is in its infancy, future work might create otherwise identical private and mixed duopoly markets to see if this difference in incentives might explain differences in the success of delegation (see Huck et al. 2004).

In additional tests, we confirm that the experienced owner treatment remains close to the hypothesized predictions on average despite reintroducing the possibility of social preferences. Thus, we suggest these preferences are not responsible for the failure to reach equilibrium. Nonetheless, we note that we have not tested for any particular theory of social preferences, say inequity aversion, and that future work might design delegation models that allow such tests. We do, however, show that there is less variation in the choices in the experienced owner treatment compared to the inexperienced owner treatment. Similarly, a decomposition of the derivations shows that both owners and managers have smaller deviations from predictions. These additional tests added to the robustness of the basic findings and round out the picture that the predictions of delegation in a mixed duopoly can be confirmed in the laboratory but only when owners are provided sufficient experience without the uncertainty of managerial choices or potential social relations.

Table 1: Means and Standard deviations

	X	Y	q_A	q_B	# of obs.	quantity correlation
Theoretical Predictions (ND)	--	--	4.36	2.91	--	--
ND	--	--	4.500 (0.142)	2.840 (0.258)	8	0.467(**)
Theoretical Predictions (IO/RM/EO)	0.57	0.84	4.06	3.24	--	--
IO	0.522 (0.052)	0.643 (0.121)	4.025 (0.163)	3.821 (0.241)	8	0.325
RM	0.548 (0.037)	0.862 (0.032)	--	--	16	--
EO	0.530 (0.051)	0.834 (0.123)	4.108 (0.190)	3.412 (0.273)	8	0.49(**)

Note: Standard deviations in parentheses; **: $p < 0.05$.

Table 2: Testing the No Delegation Treatment

Null Hypothesis	Experimental Result	Statistical Test	Test Statistics	# of obs
$q_A - q_B = 0$	1.65 (0.338)	Wilcoxon signed-ranks test	2.521 [0.0117]	8
$q_A = 4.36$	4.50(0.142)	Wilcoxon signed-ranks test	1.400 [0.1614]	8
$q_B = 2.91$	2.84(0.258)	Wilcoxon signed-ranks test	-1.120 [0.2626]	8
$W = 62.41$	62.10(0.456)	Wilcoxon signed-ranks test	-1.540 [0.1235]	8

Note: Standard deviations in parentheses and p-values in square brackets.

Table 3: Testing the Inexperienced Owner Treatment

Null Hypothesis	Experimental Result	Statistical Test	Test Statistics	# of obs
$q_A - q_B = 0$	0.20(0.340)	Wilcoxon signed-ranks test	1.400 [0.1614]	8
x =0.57	0.52(0.052)	Wilcoxon signed-ranks test	-2.380 [0.0173]	8
y =0.84	0.64(0.121)	Wilcoxon signed-ranks test	-2.521 [0.0117]	8
q_A =4.06	4.03(0.263)	Wilcoxon signed-ranks test	-0.420 [0.6744]	8
$q_A - q_A(x, y)$ =0	-0.04(0.328)	Wilcoxon signed-ranks test	-0.140 [0.8886]	8
q_B =3.24	3.82(0.241)	Wilcoxon signed-ranks test	2.521 [0.0117]	8
$q_B - q_B(x, y)$ =0	0.17(0.269)	Wilcoxon signed-ranks test	1.680 [0.0929]	8
W = 63.17	63.78(0.155)	Wilcoxon signed-ranks test	2.521 [0.0117]	8
$W^{ND} = W^{IO}$	--	Wilcoxon-Mann-Whitney rank sum test	-3.361 [0.0008]	16
$q_A{}^{ND} = q_A{}^{IO}$	--	Wilcoxon-Mann-Whitney rank sum test	2.941 [0.0033]	16
$q_B{}^{ND} = q_B{}^{IO}$	--	Wilcoxon-Mann-Whitney rank sum test	-3.361 [0.0008]	16

Note: Standard deviations in parentheses and p-values in square brackets.

Table 4: Testing the Robot Manager Treatment

Null Hypothesis	Experimental Result	Statistical Test	Test Statistics	# of obs
x =0.57	0.55(0.037)	Wilcoxon signed rank test	-2.017 [0.0437]	16
y =0.84	0.86(0.032)	Wilcoxon signed rank test	2.223 [0.0262]	16
W = 63.17	63.07(0.143)	Wilcoxon signed rank test	-2.327 [0.0200]	16
$x^{RM} = x^{IO}$	--	Wilcoxon-Mann-Whitney rank sum test	1.041 [0.2979]	24
$y^{RM} = y^{IO}$	--	Wilcoxon-Mann-Whitney rank sum test	3.797 [0.0001]	24

Note: Standard deviations in parentheses and p-values in square brackets.

Table 5: Testing the Experienced Owner Treatment

Null Hypothesis	Experimental Result	Statistical Test	Test Statistics	# of obs
$q_A - q_B$=0	0.70(0.419)	Wilcoxon signed rank test	2.380 [0.0173]	8
x =0.57	0.53(0.051)	Wilcoxon signed rank test	-1.820 [0.0687]	8
y =0.84	0.83(0.123)	Wilcoxon signed rank test	0.840 [0.4008]	8
q_A=4.06	4.11(0.191)	Wilcoxon signed rank test	0.700 [0.4838]	8
$q_A - q_A(x,y)$ =0	-0.07(0.196)	Wilcoxon signed-ranks test	-1.120 [0.2626]	8
q_B=3.24	3.41(0.273)	Wilcoxon signed rank test	1.680 [0.0929]	8
$q_B - q_B(x,y)$ =0	0.17(0.207)	Wilcoxon signed-ranks test	1.960 [0.0499]	8
W = 63.17	63.41(0.345)	Wilcoxon signed rank test	1.540 [0.1235]	8
$x^{EO} = x^{IO}$	--	Wilcoxon-Mann-Whitney rank sum test	0.315 [0.7527]	16
$y^{EO} = y^{IO}$	--	Wilcoxon-Mann-Whitney rank sum test	2.521 [0.0117]	16
$q_A^{EO} = q_A^{IO}$	--	Wilcoxon-Mann-Whitney rank sum test	0.420 [0.6744]	16
$q_B^{EO} = q_B^{IO}$	--	Wilcoxon-Mann-Whitney rank sum test	-2.415 [0.0157]	16
$W^{EO} = W^{IO}$	--	Wilcoxon-Mann-Whitney rank sum test	-1.995 [0.0460]	16

Note: Standard deviations in parentheses and p-values in square brackets.

Table 6: Measuring the Deviations between Predicted and Experimental Results

	IO	RM	EO
MSD(x, x*)	0.0428	0.0113	0.0117*
MSD(y, y*)	0.0966	0.0059	0.0214*
MSD(q_A, q_A^*)	0.9989	--	0.1782*
MSD(q_B, q_B^*)	1.2424	--	0.3527*
# of obs.	8	16	8

Note: The asterisk indicates that the hypothesis that MSE of the experienced owner treatment and the inexperienced owner treatment have the same measure of central tendency can be rejected at the five percent level in the rank sum test.

Figure 1 : Payoffs of Public and Private Owners

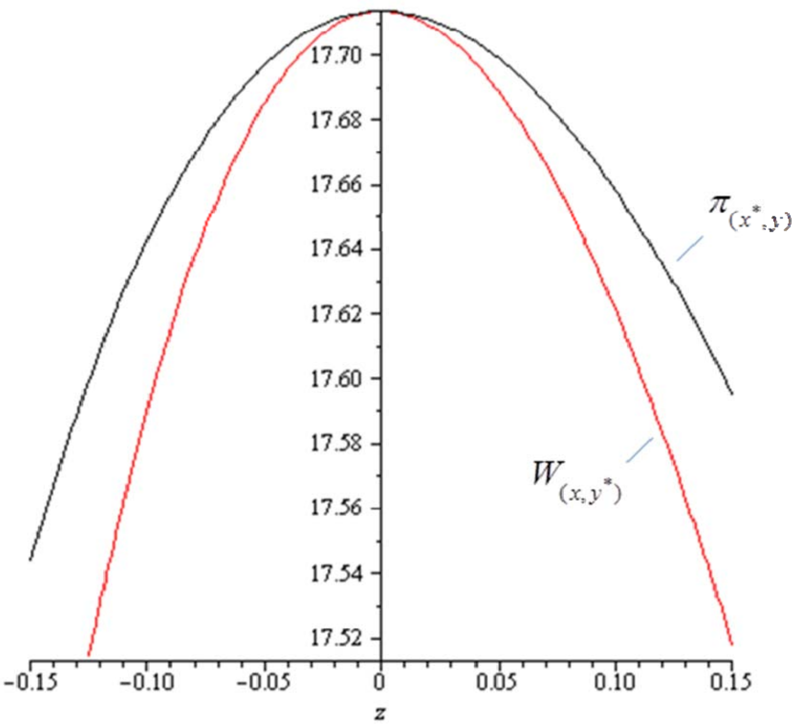

The payoffs of each owner are shown as a function of deviations from her optima; i.e. $z = x - x^*$ or $z = y - y^*$. For ease of comparison we parallel shift the public owners curve to anchor to $\pi_{(x^*,y)}$.

Figure 2: Mean Squared Deviation between Predicted and Experimental Results over Periods

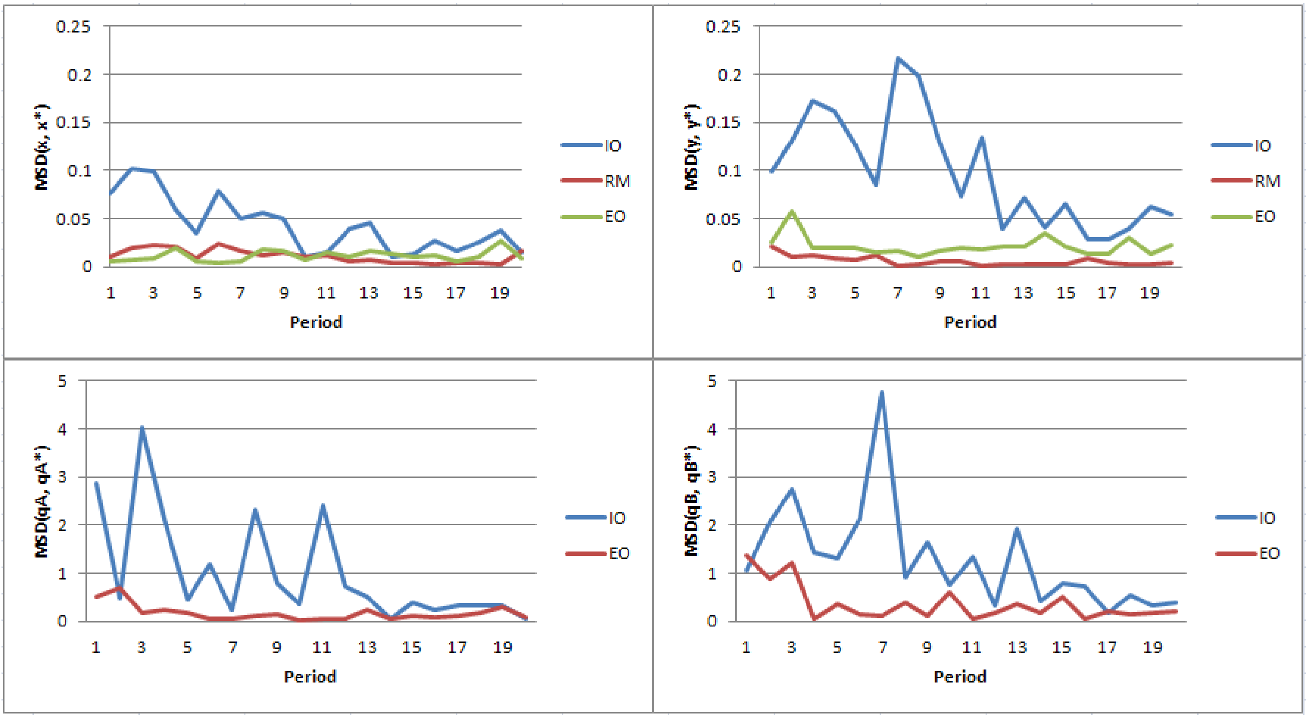

References

Barros, F., 1995. “Incentive Schemes as Strategic Variables: An Application to a Mixed Duopoly,” *International Journal of Industrial Organization*, 13: 373-86.

Bhaumik, S.K. and Dimova, R. 2004, ‘How Important is Ownership in a Market with a Level PlayingField’, *Journal of Comparative Economics*, 32: 165–180.

Binmore, K., A. Shaked and J. Sutton. 1985. "Testing Noncooperative Bargaining Theory: A Preliminary Study," *American Economic Review*, 75: 1178 – 80.

Bortolotti B., M. Fantini and D. Siniscalco. 2004. "Privatization around the World: Evidence from Panel Data," *Journal of Public Economics*, 88:305–332

Camerer, C.F. and E. Fehr. 2006. “When Does ‘Economic Man’ Dominate Social Behavior?” *Science,* 311: 47-52.

Cox, J.C. and M. Walker. 1998. “Learning to Play Cournot Duopoly Strategies,” *Journal of Economic Behavior and Organization,* 36: 141 – 61.

DeFraja, G. and F. Delbono. 1989. "Alternative Strategies of a Public Enterprise in Oligopoly," *Oxford Economic Papers,* 41: 302 – 11.

Drago, R. and J.S. Heywood. 1989. "Tournaments, Piece Rates and the Shape of the Payoff Function," *Journal of Political Economy,* 97: 992 – 98.

Fernandez-Ruiz, J., 2009. “Managerial Delegation in a Mixed Duopoly with a Foreign Competitor,” *Economics Bulletin,* 29: 90-99.

Fershtman, C. and K. L. Judd, 1987. "Equilibrium Incentives in Oligopoly", *American Economic Review*, 77: 927 – 40.

Fischbacher, U., 2007. “z-Tree: Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics*, 10: 171-78.

Harrison, G. 1989. "Theory and Misbehavior of First-Price Auctions," *American Economic Review*, 79: 749-762.

Heywood, J.S. and G. Ye, 2009. “Delegation in a Mixed Oligopoly: the Case of Multiple Private firms,” *Managerial and Decision Economics*, 30: 71 - 82.

Huck, S., H.T. Normann and J. Oechssler. 1999. "Learning in Cournot Oligopoly: An Experiment," *Economic Journal,* 109: c80 – c95.

Huck, S., W. Mueller and H.T. Normann. 2004. "Strategic Delegation in Experimental Markets," *International Journal of Industrial Organization,* 22: 561 – 74.

Johnson, E., C. Camerer, S. Sen and T. Rymon, 2002. "Detecting Failure of Backward Induction: Monitoring Information Search in Sequential Bargaining Experiments," *Journal of Economic Theory*, 104: 16-47.

Matsumura, T. 1998. "Partial Privatization in Mixed Duopoly," *Journal of Public Economics*, 70: 473 – 83.

Merrill, W.C and N. Schneider, 1966. "Gvoernment Firms in Oligopoly Industries: A Short Run Analysis," *Quarterly Journal of Economics,* 80: 400 – 12.

Neelin, J., H. Sonnenschein and M. Spiegel, 1988, "A Further Test of Noncooperative Bargaining Theory: Comment," *American Economic Review*, 78: 824-36.

Ochs, J. and A. Roth, 1989. "An Experimental Study of Sequential Bargaining," *American Economic Review*, 79: 355-84.

Requate, T. and I. Waichman. 2011. "A Profit Table or a Profit Calculator? A Note on the Design of Cournot Oligopoly Experiments," *Experimental Economics,* 14: 36 -46.

Sklivas, S.D., 1987, “The Strategic Choice of Managerial Incentives,” *Rand Journal of Economics,* 18: 452–58.

Vickers, J., 1985, “Delegation and the Theory of the Firm,” *Economic Journal*, 95: 138-47.

White, M.D. 2001. "Managerial Incentives and the Decision to Hire Managers in Markets with Public and Private Firms," *European Journal of Political Economy,* 17: 877 – 96.

White, M.D. 2002. "Political Manipulation of a Public Firm's Objective Function," *Journal of Economic Behavior and Organization,* 49: 487 – 99.

ENDNOTES

[1] Despite the move toward privatization, many countries retain publically owned firms that frequently compete against private firms. See Bortolotti et al. (2004).

[2] We recognize that there exists an enormous experimental literature on Cournot duopoly models in a fully private market. A valuable summary of over 40 experimental studies is provided by Requate and Waichman (2011).

[3] This can be contrasted with unilateral decision of a government owner to partially privatize a public firm generating an objective function that weights both welfare and firm profit. This can cause a reduction in the public firm output that is nonetheless welfare enhancing as it reduces cost asymmetry associated with convex cost structures (Matsumura, 1998).

[4] Note that the public owner encourages the public manager to give greater weight to consumer surplus by decreasing x.

[5] See Huck et al. (1999) on the consequences of not revealing such structural information.

[6] Thus, we follow the view of Johnson et al. (2002) that taking the experience of playing games into account is "a sensible shift away from interpreting game-theoretic equilibria as solutions that brilliant players figure out, toward thinking of equilibria as resting points which are the result of evolution or learning by players of limited rationality (Page 41)."

[7] As the fundamental test involves a ranking of all observations and not their specific values, the measure of central tendency is more appropriately thought of as the median.

[8] As an alternative, we dropped the theoretical flexibility of the nonparametric test and tested the hypothesis that experimental means differ from hypothesized values with t-tests. There are no substantive differences from those we present and these are available from the authors.