

Learning to Make Strategic Moves: Experimental Evidence on Strategic Information Acquisition

Anders Poulsen*

April 16, 2007

Abstract

Do players make strategic commitments, such as ‘burning one’s bridges’? Since such strategic moves can appear somewhat counter-intuitive, we conducted an experiment to see whether people make the predicted strategic move. The experiment uses a simple bargaining situation. A player can make a strategic move of committing to not seeing what the other player will demand. Our data show that the subjects do, but only after substantial time, learn to make the predicted strategic move. Our results lend support to models that stress the relevance of strategic moves, but they also show that players need time to learn to make the strategic move.

COMMENTS VERY WELCOME.

Keywords: Strategic moves; commitment; bargaining; information; strategic value of information; physical timing effects; virtual observability; endogenous timing; learning; experiment.

JEL Classification: C72; C78; C90; C92; D63; D80.

1 Introduction

A crucial insight from game theory is that a player involved in an interactive situation can gain from making what Schelling (1960) calls a *strategic move*. Well-known examples are moving before someone else to get a first-mover advantage (Schelling, 1960,

*School of Economics, University of East Anglia, Norwich NR4 7TJ, United Kingdom. E-mail: a.poulsen@uea.ac.uk. I thank Niels Bruhn Christensen for invaluable help with the design and the programming. Thanks to Chaim Fershtman, Werner Güth, Graham Loomes, Ben McQuillin, Bob Sugden, the seminar audiences at University of East Anglia, and at the 2006 Atlanta International meeting of the Economic Science Association for helpful and encouraging comments. Toke Fosgaard, Mads Harmsen, and Frederik Øvlisen from the LEE lab in Copenhagen gave professional lab assistance. All errors are mine. Financial support from the Danish Social Science Research Council is gratefully acknowledged.

and Tirole, 1988); signing contracts with third-parties (Aghion and Bolton, 1985, and Bensaïd and Gary-Bobo, 1993); burning money (Ben-Porath and Dekel, 1992, van Damme, 1989, Huck and Müller, 2005); strategic delegation (Fershtman and Gneezy, 2001, Fershtman and Kalai, 1997, and Schelling, 1960); changing the information structure (Hauk and Hurkens, 2001, Hurkens and Vulkan, 2006, and Schelling, 1960).

These important game-theoretic results about strategic moves lead naturally to an empirical question: Do decision makers in practice understand and appreciate the usefulness of making these strategic moves? This, we believe, is a non-trivial issue: the strategic moves mentioned above are characterized by a deliberate restriction of one's freedom of action, a reduction of one's payoffs from certain outcomes, or an avoidance of information. It could be that most decision makers would prefer not to engage in such activities since they may appear counterintuitive and irrational. Perhaps, due to bounded rationality and cognitive biases, decision makers would underestimate or completely ignore the strategic impact of the strategic move on the opponent's behavior. Such biases have been experimentally documented in other, mostly non-strategic, decision problems. See Loewenstein, Moore, and Weber (2006) and the references therein.

To explore these issues in a simple and controlled environment, we run experiments using a simple game, namely a sequential Nash demand game (Nash, 1953). One player, A, makes a demand for a share of an exogenously given amount of money. Then another player, B, makes a demand. If the sum of the demands do not exceed the available amount of money, each player gets the money he demanded; otherwise each player gets zero. We give player B the option of making a strategic move that we call *strategic information acquisition*.

In our first treatment, player B can, before player A makes a demand, commit himself to see A's demand or not before B makes his demand, and A observes B's commitment decision. We refer to this as the *Commitment* game. In this game, if B decides not to see A's demand, it is common knowledge that B is an uninformed second-mover. If, on the other hand, B decides to see A's demand, it is common knowledge that B is an informed second mover. Our theoretical prediction is that B will make the strategic move of not seeing A's demand.

In our second treatment player B again commits either to see or else not to see player A's demand before B makes his own demand, but now A does not observe what B has committed himself to. This is the *Unobserved Commitment* game. Theory predicts that player B should now commit to seeing player A's demand, since A cannot condition his demand on B's commitment decision.

In our third treatment, the *No Commitment* treatment, B has no commitment move available: he decides whether or not to be informed about player A's demand after player A had made his demand. Indeed, since player B's decision whether or not to see player A's demand is not observed by player A in the Unobserved Commitment and the No Commitment games, game theory views the two games as strategically equivalent. The experimental literature has however documented behavioral differences in games, such as the Unobserved and No Commitment games that differ only in the order but not observability of moves. This is referred to as physical timing effects, or virtual observability. See Güth, Huck, and Rapoport (1998), Huck and Müller (2005), and Weber, Camerer, and Knez (2006). By comparing behavior in

the Unobserved Commitment and the No Commitment games, we can see whether physical timing significantly influences the making of strategic moves.

Our main purpose is to see how a given commitment opportunity is used by players, not how commitments are created. We therefore deliberately avoid specifying the exact way in which players achieve commitment. A richer design would specify the exact commitment technology (third-party contracts, strategic delegation, inability to receive information, and so on), how costly and easily reversible the commitment is (see Muthoo, 1996) and the extent to which (contractual) commitments can be re-negotiated. We discuss this further in Section 5. Our simple set-up can be interpreted as a reduced-form model of a more complicated situation where commitment is (sufficiently) cheap and where it is (sufficiently) costly to re-negotiate the commitment.

Our experimental findings can be summarized as follows: Players learn over time to make the strategic moves predicted for our treatments, and the physical timing effects mentioned above tend to disappear over time. The time required for learning depends, however, on the treatment: in the Commitment treatment the time needed for player Bs to learn that it is optimal to avoid information is considerably longer than the time they need in the Unobserved Commitment and No Commitment treatment to learn that information is optimal. We interpret this learning asymmetry as evidence that players bring with them into the lab an ‘information is good’ bias. Such a bias is plausible since information is typically viewed as desirable in most everyday decision problems; see for example Loewenstein, Moore, and Weber (2006). When put in a strategic situation such as the Commitment treatment, however, this bias leads to information being collected. Since this tends to give player B a low payoff, player Bs over time unlearn the bias. More generally, therefore, our experiment provides support for models stressing the relevance of strategic moves, when it is acknowledged that decision makers need ample time to learn to make the strategic moves.

Our results contribute to a quite small literature of experimental work on whether people make optimal strategic moves. We are only aware of the unpublished paper by Andreoni (2005), and the published works of Fershtman and Gneezy (2001), Huck and Müller (2005), Morgan and Vardy (2004), and Poulsen and Tan (2006). We describe in Section 6 how our work relates to these contributions.

The rest of the paper is organized as follows. In Section 2, we develop the theoretical predictions. Section 3 explains the design of the experiment and how it was conducted. Section 4 reports our findings. In Section 5 we discuss our findings and the design, and outline possible future research. We relate our experimental design and results to the existing literature in Section 6. Section 7 concludes. The instructions are in the Appendix.

2 Theoretical Predictions

There are two players, $i = A, B$, and a sum of money, X . Let ϵ denote the smallest monetary unit. The set of feasible demands is $Z = \{0, \epsilon, 2\epsilon, \dots, X - \epsilon, X\}$. Denote player i 's demand by x_i , $i = A, B$, where $x_i \in Z$. If $x_A + x_B \leq X$, each player gets what he demanded. If $x_A + x_B > X$, each player gets zero. We assume all players are rational, seek to maximize their expected money earnings, and that this is common

knowledge.

2.1 The Benchmark Game (BM)

The Benchmark game is a sequential Nash demand game with perfect information. Player A first makes a demand, x_A . Player B sees player A's demand and then B makes his demand, x_B . The theoretical prediction is that player A in equilibrium gets almost all the money. More precisely, there are two pure-strategy subgame-perfect equilibria. In the first subgame-perfect equilibrium, player A demands all the money, $x_A^* = X$, and player B demands $x_B^* = 1 - x_A$ for any demand x_A made by player A. In the second subgame-perfect equilibrium, player A demands $x_A^* = X - \epsilon$, where ϵ is the smallest monetary unit, and player B demands $x_B^* = X - x_A$ for any demand x_A with $x_A < X$, and makes some demand $x_B^* > 0$ for $x_A = X$.

2.2 The Commitment Game (C)

In the Commitment game player B first irrevocably decides whether or not to see the demand that player A will make. We refer to this as B's information decision. Player A observes which information decision B made, and then player A makes his demand. Player B then sees player A's demand or not, as determined by his information decision. Finally player B makes his demand.

Suppose first player B decides to see player A's demand. The subgame that follows is identical to the Benchmark game, so in any subgame-perfect equilibrium of the Commitment game player B gets almost zero if he sees player A's demand. Suppose next player B decides not to see player A's demand. In the subgame that follows it is common knowledge that B does not see A's demand when B makes his demand. Any feasible pair of demands (x_A, x_B) such that $x_A + x_B = X$ is a Nash equilibrium of this subgame. Our prediction will be based on the unique symmetric pure-strategy Nash equilibrium, namely that each player demands half of the money. Thus in any subgame-perfect equilibrium of the overall game, player B gets half of X if he decides not to see player A's demand. It is thus optimal for B to decide not to see A's demand. The subgame-perfect equilibrium of the Commitment game is therefore: Player B decides not to see player A's demand and B demands half of the money when it is his turn to make a demand; if player A observes that player B decided not to see player A's demand, player A demands half of the money; and if player A observes that player B decided to see player A's demand, then player A demands (almost) all the money and player B demands the rest.

This theoretical prediction is we believe quite intuitive: if player B decides not to see player A's demand, a sequential Nash demand game with imperfect information is played, and this is strategically equivalent to a simultaneous-move Nash demand game. In this game we expect each player to demand half of the money. If B decides to see player A's demand, B becomes second mover in a sequential Nash demand game with perfect information, and in this game B should expect to get significantly less than half of the money. We predict that B prefers to play the former game and so decides not to see player A's demand.

2.3 The Unobserved Commitment Game (UC)

In the Unobserved Commitment game, Player B first irrevocably decides whether or not to see player A's demand, as in the Commitment game. Player A does however not learn B's information decision. In other words, player A knows that B has made *some* information decision (to see or not to see player A's demand), but player A does not know *which* information decision B made. Then player A makes a demand. After this player B sees player A's demand or not, as determined by player B's information decision. Finally player B makes his demand.

There are many Nash equilibria in this game. For example, the B strategy of not seeing player A's demand and demanding one-half and the player A strategy of demanding one-half is a Nash equilibrium. We observe however that the B strategy of seeing player A's demand and playing a best reply to player A's demand (that is, demanding $x_B = X - x_A$ when seeing that player A demanded $x_A < X$ and making any feasible demand when seeing $x_A = X$) weakly dominates any other B strategy. Assuming that player B avoids dominated strategies, it follows that player B decides to see player A's demand. This being common knowledge, the players' equilibrium demands are identical to those in the Benchmark game. Intuitively, since B's information decision is not observed by player A, there can be no strategic effect of B's information decision on the demand that player A subsequently makes. It is then best for B to decide to see A's demand.

2.4 The No Commitment Game (NC)

In the No Commitment game, player B has no opportunity to commit to see player A's demand or not before player A makes a demand. The order of moves are: player A first makes a demand. Then player B decides whether or not to see player A's demand. Finally, player B makes his demand. As in the Unobserved Commitment game, the B strategy of seeing player A's demand and playing a best reply weakly dominates all other strategies. Assuming once more that B does not play dominated strategies, the equilibrium demands for the No Commitment game is the same as for the Unobserved Commitment game.

2.5 Remarks

2.5.1 Physical Timing Effects

From a game-theoretic point of view, the Unobserved Commitment and No Commitment games are the same: it does not matter that player B's decision whether or not to see player A's demand or not is made at different times (before or after player A makes his demand), since in both games B's information decision is not observed by player A. In both games, therefore, player B should decide to see player A's demand. So why study both games? There were two reasons. First, we wish to compare the Commitment game with a game where there is the same strategic move available, but where this move is not observed by player A. This is the Unobserved Commitment game. Second, it is from the experimental games literature known that situations

that differ only in the physical timing but not the observability of moves can generate different observed behavior. This phenomenon is referred to as ‘physical timing effects’, or ‘virtual observability’. See, for example, Güth, Huck, and Rapoport (1989), Huck and Müller (2005), and Weber, Camerer, and Knez (2004). In our context, the fact that player B in the Unobserved Commitment game makes his information decision first could, although the decision is not observed by player A, make player B more likely to avoid information than in the No Commitment game, where B makes his information decision after player A. Comparing behavior in the Unobserved Commitment and the No Commitment game allows us to see whether physical timing effects per se influences B’s strategic move. Differences in physical timing alone could also have an effect on player A’s demand. In the No Commitment game, player A makes his demand before, whereas player A in the Unobserved Commitment game makes his demand after B has made his information decision. Although theory says there should be no difference, does player A make more aggressive demands in the former than in the latter game? Our data allow us to see whether this is the case.

2.5.2 Fairness and Reciprocity

The experimental game literature has documented the relevance of reciprocity and fairness concerns; see Camerer (2003). Could it be that subjects perfectly understand and appreciate the logic of strategic moves, but they avoid making them because they believe that their opponent due to such motivations could react negatively to the strategic move? If so, fairness and reciprocity would confound the cognitive aspects that are our key focus. Note, however, that in our treatments such fairness and reciprocity concerns should not deter player B from making the strategic move: in the Commitment game, player B, by making the strategic move of avoiding information, merely establishes an equal payoff distribution, and not making the strategic move would give B very little. While it is true that player B by making the strategic move lowers player A’s payoff compared to what player A would get if B did not make the strategic move, the strategic move crucially does not give player B more than player A, and this should reduce the extent to which a strategic move by player B violated player A’s sense of fairness or reciprocity.

Of course, in the Commitment game player B’s own desire to establish a fair payoff distribution reinforces the desirability of making the strategic move of avoiding information about player A’s demand. However it remains the case that B, if he decides not to see player A’s demand, must have understood that it is only by making the strategic move of not seeing player A’s demand that he can induce player A to help establishing a fair payoff distribution.

2.5.3 Comparing Demands

The theory predicts that if in the Commitment game player B decides to see player A’s demand then player A will demand almost all the money and B will demand almost zero. The same is predicted to happen in the Unobserved Commitment and No Commitment games. It is well-known from the experimental literature that such extreme payoff distributions often do not occur. In the Ultimatum game, for example, the Pro-

poser typically gets between 60 and 70 % of the available money. These moderated payoff distributions have been attributed to fairness concerns or to the first mover’s fear that the second mover will reject very low offers. See, for example, the discussion in Camerer (2003). We should expect a similar moderation in our experiment. By collecting data from the Benchmark game, we can measure the strength of these moderating factors. Taking these considerations into account, the moderated hypothesis is that demand behavior in the Unobserved Commitment and No Commitment treatments will be the same and, furthermore, that these should be indistinguishable from the demand behavior in the Benchmark game.

3 Experimental Design and Procedure

The experiments took place in the spring and fall of 2006 at the Laboratory for Experimental Economics (LEE) at University of Copenhagen, Denmark. The experiment was fully computerized, using the Z-Tree software (Fischbacher, 1999). For recruitment Greiner’s ORSEE system (Greiner, 2004) was used.

3.1 Participants

In total 254 participants, recruited from across the University of Copenhagen, participated in the experiment. Participants received a show-up fee of 50 Danish Kroner (DKK), equal to about US \$ 9 at the time of the experiment. On average a session lasted 45 minutes. Across all treatments, average earnings, including the show-up fee, was DKK 171.6, or about \$ 30.9.

Table 1 provides some information about the treatments, sessions, and number of participants.

Treatment	No. of sessions	No. of participants
BM	4	74
C	4	72
UC	3	44
NC	4	64

Table 1: Overview of treatments and sessions. BM = Benchmark; C = Commitment; UC = Unobserved Commitment; No Commitment treatment.

3.2 Experimental Procedure

After entering the laboratory, each participant was seated in front of a computer. All computers were separated by cubicles and no verbal or visual communication between participants took place during the experiments. Once all participants had read the instructions, a test was distributed (instructions and tests are in the Appendix). When all students had answered the test questions, the experimenters checked all answers. If a participant gave an incorrect answer to a question, he was asked to try again.

Any questions about the instructions or the test were answered privately. Once all participants had answered all test questions correctly, this was announced and the experiment started.

To allow for learning, the experiment consisted of 15 periods. At the start of the experiment, each participant was randomly given the player A or the B player role. A participant stayed in the same role for all 15 periods. In each period one A participant was randomly matched with a B participant.

In each period the sum X was given by 100 points. The set of feasible point demands were $\{0, 1, 2, \dots, 99, 100\}$. The A and the B participant each had to demand a number of points from this sum. If the sum of the demanded points did not exceed 100 points, each participant earned the number of points that he/she demanded. If the sum of the two demands exceeded 100 points, each earned zero points. After each period both participants were informed about each other's choices and their own earnings. After the last period the points a participant had earned in each period were added and converted into Danish Kroner (DKK), using the following exchange rate: 5 points is equal to DKK 1 (so 100 points equals DKK 20, or about \$ 3.6). After the experiment, this number of Danish Kroner was, together with the show-up fee, paid to each participant in a separate room.

The specific treatments were implemented as follows:

The Benchmark game: First participant A entered a demand. Then participant B saw A's demand on his screen. Then B entered his demand.

The Commitment game: First participant B irrevocably decided whether or not to see participant A's demand. Participant A was on his screen informed about B's decision and A entered a demand. If B decided to see A's demand, B saw A's demand on his screen, after which B entered his demand. If B decided not to see A's demand, B was only informed that A had made his demand. B then entered his demand.

Unobserved Commitment game: First participant B irrevocably decided whether or not to see participant A's demand. Participant A, knowing that B had made a decision but not which one, then entered a demand. If B decided to see A's demand, B saw A's demand on his screen, after which B entered his demand. If B decided not to see A's demand, B was only informed that A had made his demand, after which B entered his demand.

No Commitment game: First participant A entered a demand. Then participant B decided whether or not to see A's demand. If B decided to see A's demand, B saw A's demand on his screen, after which B entered his demand. If B decided not to see A's demand, B was on his screen informed only that A had made his demand, after which B entered his demand.

4 Experimental Data

4.1 Player B's Information Decision

We first compare across treatments the percentage of participant Bs who decided to see A's demand. Figure 1 shows these percentages for each of the 15 periods, for the Commitment (C), Unobserved Commitment (UC), and No Commitment (NC) treatments.




Figure 1: Percentage of B participants who decide to see A's demand in the various treatments.

In the Commitment treatment the percentage of participant Bs who decide to see A's demand initially exceeds 60 %. The percentage falls steadily over time, however, and towards the end fewer than 15 % of Bs decide to see A's demand. The average percentage is 29.3 %. In the No Commitment treatment the percentage of participant Bs who see A's demand fluctuates between 80 and 95 %. The average percentage is 88.8 %. Unlike the Commitment treatment, it does not seem that much learning takes place. In the Unobserved Commitment treatment there is a initially a slight drop in the percentage of Bs who become informed, but the percentage picks up over time and approaches the theoretical prediction. The average percentage is 81.8 %.

On average, and especially in the earlier periods, fewer participant Bs decide to see A's demand in the Unobserved Commitment than in the No Commitment treatment. This can be interpreted as evidence of a physical timing effect: since B moves first in the UC game but second in the NC game, he decides to avoid information about A's demand more often in the first than in the second game. As shown by Figure 1, how-

ever, over time behavior in the Unobserved and No Commitment treatments become indistinguishable. From period nine and onwards, it is hard to see any difference.

Table 2 repeats the average across-period percentages of B participants who decided to see A’s demand in the treatments. These across-period averages confirm the message given by figure 1: The proportion of Bs who decided to see A’s demand in the Unobserved and the No Commitment treatments differ significantly (chi-square test, $X^2 = 14.21$, $p < 0.001$). This result is however made less relevant by the already mentioned fact that B’s information decision in the two treatments become more and more similar over time.

C	UC	NC
29.3	81.8	88.8

Table 2: The average across-period percentage of B participants who decided to see A’s demand in the C, UC, and NC treatments.

All in all, there is strong evidence that the B participants over time in each treatment learn to make the predicted strategic move. The learning is however asymmetric: it takes more time to learn to avoid information in the Commitment game than it takes to obtain the information in the Unobserved Commitment and in the No Commitment game. This is quite intuitive: in most everyday situations information is thought of as being desirable (see Loewenstein, Moore, and Weber, 2006). When put in an unusual strategic situation like the Commitment game, decision makers need time to unlearn their homegrown propensity to seek information. There is some evidence of a physical timing effect on B’s information decision in the first periods, but the effect diminishes over time.

4.2 A Closer Look at The Commitment Treatment

Figures 2 and 3 below show participant A’s and B’s average demands and earnings in each period in the Commitment treatment, conditional on whether or not B decided to see A’s demand. Figures 4 and 5 in the Appendix shows the relative frequency distributions of A and B demands and earnings, respectively.

Demands Figure 2 shows that in the Commitment treatment the A and B demands vary strongly with B’s information decision, in the direction predicted by the theory. When B decides not to see A’s demand, both participants tend to demand half of the 100 points. Indeed, A demanded 50 in 343 out of the 382 cases (89.8 %) where B decided not to see A’s demand. In 337 out of the same 382 cases (88.2 %), B demanded exactly half. Pooling all A and B demands across periods, the average A and B demand when B decides to see A’ demand is 49.5 and 50.2, respectively. Figure 4 in the Appendix shows the relative frequency distributions of participant A and B demands. When B decides to see A’s demand, A demands more and B demands less. The average across-period A and B demand are 63.3 and 42.8, respectively. Comparing across treatments, these A demands are statistically significantly different at any conventional significance level (Wilcoxon-Mann-Whitney test, $z=12.8$), as are the B demands ($z=-8.22$).

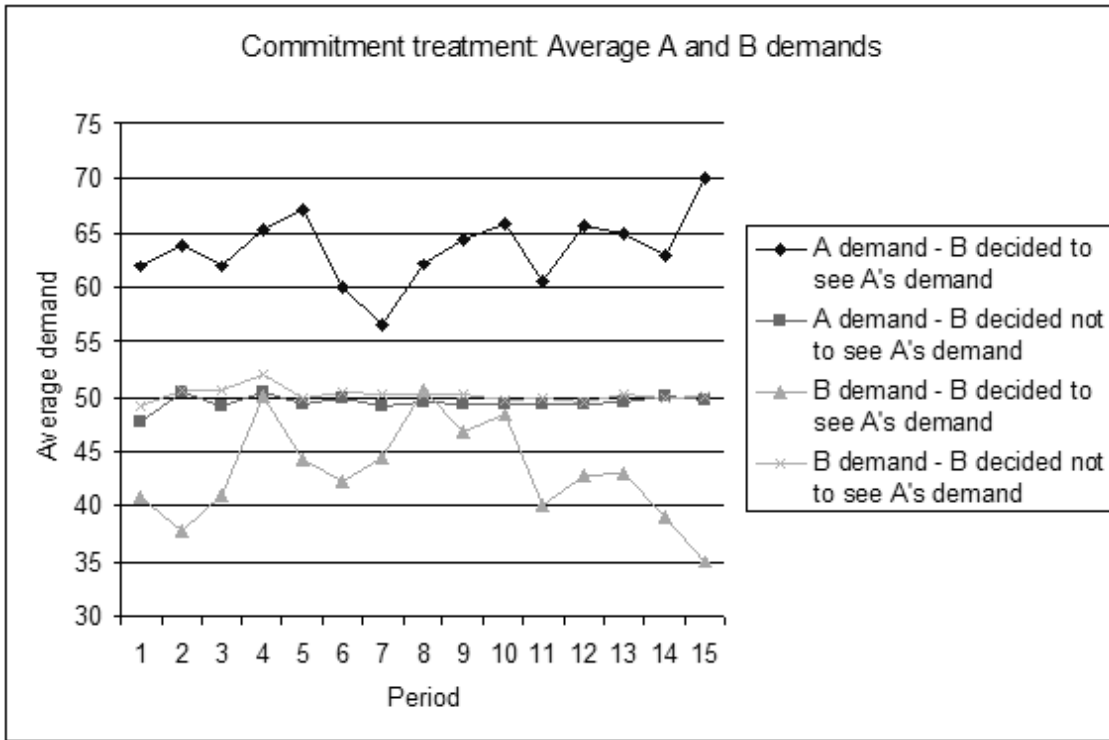


Figure 2: Average participant A and B demands in the Commitment treatment, conditional on B's information decision.

Earnings As is the case for their demands, participants' earnings in the Commitment game varies systematically with B's information decision. Figure 3 shows that when B decides not to see A's demand, the A and B earnings are both strongly concentrated on 50, the earnings from the equal split. In more than 88 % of all encounters, the 100 points are split equally. When B decides to see A's demand, the earnings distributions are much more dispersed. Figure 5 in the Appendix shows that the A and B earning distributions differ significantly across the treatments (Wilcoxon-Mann-Whitney test. A earnings: $z = 4.91$. B earnings: $z = -10.21$).

Participant B is as predicted significantly better off when he decides not to see A's demand than when he decides to do so. The across-period average B earnings are 46.3 and 34.6, respectively. For participant A, the comparison of earnings is less obvious. The across-period averages are 46.3 and 45.9, respectively. Considering only these very similar averages, however, neglects the fact that, as shown in Figure 5 in the Appendix, there is considerable dispersion in A's demands and earnings when A learns that B decided to see his demand. About 25 % of A's earn half of the points, and almost the same percentage earn zero. The remaining 50 % of the A participants earn between 55 and 85 % of the points. The top panel in Figure 6 in the Appendix shows the expected payoff to participant A from making various demands in the Commitment treatment when B decides to see A's demand, given the observed B demand behavior. The demand of 65 maximizes A's expected money earnings. Nevertheless, other A demands such as 50, 60, and 75 do not give a radically lower expected money payoff, and this can explain the observed dispersion in the observed A demands.




Figure 3: Average participant A and B earnings in the Commitment treatment, conditional on B's information decision.

Agreements and Efficiency We measure efficiency as the proportion of the total available payoff (100 points in each encounter) that was paid out to participants. When participant B decides not to see A's demand, efficiency is 92.4 %, while it is 77.1 % when B decides to see A's demand. This difference in efficiency can be explained by considering the difference in agreements (when $x_A + x_B \leq 100$). Out of the 158 A-B encounters where B decides to see A's demand, agreement occurred in 120 cases, a proportion of 0.76. Out of the 382 cases where B decided not to see A's demand, there was agreement in 356 cases, a proportion of 0.93. This difference in proportions is statistically significant at any conventional significance level (Chi-square test, $X^2 = 30.2$).

The reason for this strong difference in agreements can be explained as follows. When participant B decides to see A's demand, A reacts by making higher demands, and these in turn make B more likely to 'punish' A by demanding more than the residual, $100 - x_A$, thereby giving each participant a zero payoff. The lower panel in Figure 6 in the Appendix shows, for the case where B decided to see A's demand in the Commitment treatment, the proportion of B demands, made conditionally on observing some A demand, that lead to disagreement. We see that as A's demands increases from 50 and beyond, such 'spiteful' B behavior increases.

Is There Any Evidence That Participant A Is Influenced By Fairness or Reciprocity When Participant B Makes the Strategic Move? In Section 2.5 it was noted that players might avoid making strategic moves if they feared an adverse reaction by the other player. Such a reaction could be fueled by fairness or reciprocity

concerns. Our design sought to minimize the presence of this factor. The data suggest that this has been accomplished: If in the Commitment treatment participant A becomes upset when learning that B decided not to see A’s demand, we should expect A to ‘punish’ B by making a ‘spiteful’ demand, that is a demand large enough to result in disagreement. More precisely, to the extent that A plausibly believes that B will demand 50 when B decides not to see A’s demand, a spiteful A should demand more than 50. The data show that out of the 382 encounters where B decides not to see A’s demand, A made a demand exceeding 50 in only 5 encounters. There is thus no evidence that any fairness or reciprocity concerns matter for participant A when he learns that B decides to commit to not seeing A’s demand.

4.3 Across-Treatment Comparisons

Table 3 shows the average A and B demands (denoted \bar{x}_A and \bar{x}_B), average point earnings (π_A and π_B), and efficiency ($\pi_A + \pi_B$) in the treatments. In the Unobserved Commitment (UC) and No Commitment (NC) treatments A’s demands do, by definition, not vary with B’s information decision. Also, the reported B earnings in the first and last row is not a typo; the earnings differ only in the second decimal (they equal 34.62 and 34.64, respectively).

	\bar{x}_A	\bar{x}_B	π_A	π_B	$\pi_A + \pi_B$
BM	60.2	43.4	48.9	34.6	83.5
C: B does not see A’s demand	49.6	50.2	46.1	46.3	92.4
C: B sees A’s demand	63.3	42.8	45.9	31.2	77.1
UC: B does not see A’s demand	59.9	49	23.8	21.4	45.2
UC: B sees A’s demand	59.9	43.6	46.6	33.7	80.3
NC: B does not see A’s demand	59.5	48.7	28.3	25.9	54.2
NC: B sees A’s demand	59.5	47.2	46.3	34.6	80.9

Table 3: Average participant A and B demands, earnings, and efficiency, conditional on B’s information decision, computed as averages across all periods. BM=Benchmark. C= Commitment, UC = Unobserved Commitment, NC = No Commitment.

Does Observable Commitment Enhance Efficiency? In the BM treatment, participant B exogenously sees A’s demand, whereas B in treatment C can visibly commit to not seeing A’s demand. According to the theory, outcomes should be fully efficient in both the Benchmark and the Commitment treatment. Let us see how this compares with observed behavior. In the Benchmark treatment efficiency is 83.5. As shown in Table 3, when B in the Commitment treatment commits to not seeing A’s demand, efficiency is 92.4 (in terms of the frequency of agreements, in the Benchmark treatment 464 out of 555 encounters, a proportion of 0.83, results in agreement; when B does not see A’s demand in the Commitment treatment, agreement occurs in 356 out of 382 encounters, a proportion of 0.93; the difference in proportions is strongly significant: Chi-square test, $X^2 = 18.2$). But when B in the Commitment treatment commits to see A’s demand, efficiency is only 77.1 (when in the Commitment treatment participant B decides to see A’s demand, agreement occurs in 120 out of 158 encounters, a proportion of 0.76; on comparing with the Benchmark treatment, the

difference in proportions is strongly significant: Chi-square test, $X^2 = 4.36$). Overall, efficiency in the Commitment treatment is 87.4 (in the Commitment treatment overall, there is agreement in 476 out of 540 encounters, a proportion of 0.88; on comparing with the Benchmark treatment, the difference in proportions is strongly significant: Chi-square test, $X^2 = 4.28$). The opportunity for B to commit himself to not seeing A's demand thus has a positive effect on overall efficiency, and when B makes the predicted commitment to not seeing A's demand, the positive effect is strongest, giving an almost 10 percentage point efficiency increase, from 83.5 to 92.4 %.

When B in the Commitment treatment decides to see A's demand, the game is formally the same as the Benchmark game. Why, then, is efficiency lower (77.1 %) in the former than in the latter game (83.5 %)? Table 3 shows that A on average makes tougher demands in the former than in the latter situation. The top panel in Figure 8 in the Appendix shows that there is a tendency for participant A to make more aggressive demands (demands above 70) when A in the Commitment treatment learns that B *chose* to see A's demand, than when A in the Benchmark treatment knows that B is exogenously informed about A's demand. One interpretation is that participant A, when he in the Commitment treatment observes an 'out-of-equilibrium' decision by B to see A's demand, revises his beliefs about B, namely assigns higher probability to B playing a best reply to even large A demands. Participant A consequently reacts by increasing his demand. A comparison of Figure 6 (top panel) and Figure 7 (top panel) in the Appendix shows, indeed, that the optimal A demand in the Benchmark treatment is 60, while it is 65 in the Commitment treatment when B decides to see A's demand. The more aggressive A demands in the latter treatment are therefore quite rational.

Does Observability Enhance Efficiency? The overall efficiency rate in the Commitment treatment is 87.4. In the Unobserved Commitment treatment, overall efficiency is 74. Not surprisingly, observability of commitment significantly raises efficiency: When participant B's information decision is observed by A, B tends not to see A's demand, and this in turn induces A to make moderate efficiency-enhancing demands. Non-observability results in frequent disagreement, especially when B decides to avoid information about A's demand, since demands then tend to be incompatible.

Is There a Physical Timing Effect on Participant A's Demands? The theory predicts that the A participants' demands in the Unobserved Commitment and the No Commitment treatment will be the same – it will not matter to A whether B has already made his information decision (as in the UC treatment) or whether B makes it after A moves (the NC treatment). The relative frequency distributions of demands are shown in Figure 8 in the Appendix. The average A demands in the Unobserved and No Commitment treatments are indeed very close (59.9 % and 59.5 %, respectively). There is no significant difference between the demand distributions (Wilcoxon-Mann-Whitney test, $z = -0.63$, $p = 0.26$). We thus find no evidence of a physical timing effect on A's demands.

Does Participant A Make the Same Demands in the Benchmark as in the No Commitment Treatment? The theory also predicts that participant A should

make the same demands in the Benchmark treatment as in the No Commitment and in the Unobserved Commitment treatments, because B in the two latter treatments will decide to see A's demand and so B will be in exactly the same situation. An alternative hypothesis is that A will make more cautious demands in the No Commitment treatment, since he cannot be certain that B will decide to see A's demand.

The distributions of A demands in the Benchmark and the No Commitment treatment are shown in Figure 8 in the Appendix. The average A demands are very close, 60.2 and 59.5, respectively. However more As demand half the points in the No Commitment than in the Benchmark treatment, while the opposite is true for demands in the 70-75 range. Comparing the two distributions, a Wilcoxon-Mann-Whitney test gives $z = -2.09$ and $p=0.018$. The null hypothesis of equal demand distributions can be rejected against the one-sided alternative, that the Benchmark demands stochastically dominate those in the No Commitment treatment, at the 5 %, but not at the 1 % level. There is thus some evidence that A's tend to make more cautious demands in the No Commitment treatment in the Benchmark game.

These somewhat more cautious A demands in the no Commitment treatment are rational, given the empirical B behavior. Figure 7 in the Appendix shows the expected payoff that various A demands earn, given the empirical B behavior in the Benchmark and No Commitment treatment. We have omitted those participant A demands that are only rarely made (frequency below 10). It can be seen that whereas the money maximizing A demand is 60 in the Benchmark treatment, is it only 50 in the No Commitment treatment. We can attribute this difference to the fact that about 11 % of the B participants decide to avoid information in the No Commitment treatment (cf. Table 4.1), in which case they tend to demand 50. Their presence is enough to make the 50 demand optimal for participant A.

5 Discussion

Our experiment is simple: One player has the opportunity to make a strategic move that is predicted to affect another player's behavior. The experiment measures whether the first player learns to make the optimal strategic move. Is our design *too* simple? We believe not. First, our results show that even in our simple Commitment game it takes considerable time for the players to learn to make the theoretically predicted commitment. Presumably, in a 'too simple', or 'too easy' set-up, players would either jump straight to the prediction or would get bored and behave erratically. Our data reveal a different behavior. Second, we wanted to study a simple situation where the player with the opportunity to commit would not have to worry too much about whether the other player understands what the commitment move means. Clearly this may fail to hold in more complicated situations, and future research should address this issue. Third, a more complicated design could have introduced many potentially confounding variables, and thus jeopardize internal validity.

Our main focus was on whether decision makers understood to exploit any given commitment possibility, not on whether they understood to create the opportunity in the first place. This is why we deliberately avoided specifying the 'commitment technology' through which B can achieve the commitment. This creative aspect of

making strategic moves possible is, of course, important and fundamental. As described in the Introduction, the theoretical literature has identified several possible commitment technologies: third-party contracts and strategic delegation are particularly well-known. It is also well-known that the possibility to re-negotiate contracts can sometimes neutralize any commitment value of contracts (see for example Cailaud, Jullien, and Picard, 1995). Whether this prediction is descriptively adequate seems to be yet another fruitful area for experimental research. We are only aware of one experiment that explicitly models the commitment technology (but assumes that contracts cannot be re-negotiated), namely the strategic delegation experiment by Fershtman and Gneezy (2001). We describe their design in Section 6 below.

6 Related Research

Our work is related to several other strands of the literature.

6.1 Commitment Experiments

Fershtman and Gneezy (2001) study strategic delegation (see for example Fershtman and Kalai, 1997, and the references therein) using the ultimatum game. In the first version of this game, Proposer delegation, the Proposer has a delegate who will make a proposal on behalf of the Proposer. The Proposer gives the delegate a compensation scheme that states the payment the delegate will receive, conditional on the amount of money the delegate brings back from the bargaining table. The compensation scheme is not re-negotiable. The Proposer is given a separate hiring budget with which to pay the delegate. Hiring a delegate is thus free in the sense that it does not reduce the earnings from bargaining. An important restriction on the set of feasible compensation schemes is that if the delegate end up in disagreement and so bring no money back, the delegate is paid zero. In the second version, Responder delegation, the Responder has a delegate and gives the delegate a similar compensation scheme. Fershtman and Gneezy also run a treatment with optional Proposer delegation. All their experiments used one-shot games, so unlike us they do not study learning.

The data show that Proposers benefit from having a delegate, compared to the no-delegate game. Under optional delegation, 75 % of Proposers decides to hire a delegate and the observed behavior is much the same as under mandatory delegation. Theoretically, however, the Proposer should not benefit (or be hurt) from using a delegate: he can make a take-it-or-leave-it offer himself. On the basis of further experiments, Fershtman and Gneezy conclude that a ‘hostage’ interpretation is the best explanation for why Proposers are better off using delegates. The reason is that, in their design, if the Responder rejects the delegate’s offer, both the delegate and the Proposer get zero. This can induce the Responder to be less likely to reject any given offer than when there is no delegate. The delegate can thus be used by the Proposer as a hostage or a shield behind which low offers can be made. The data show that Proposers indeed exploit this by inducing their delegates to make lower offers.

For the Responder delegation is theoretically optimal since this acts a commitment device to reject even large offers. Indeed, when the compensation scheme is observable

to the Proposer the Responder should choose a compensation scheme that induces the delegate to only accept the *largest* feasible offer. This can be done by only paying the delegate if she returns with exactly this share of the available money. The data show that with delegation Responders get 48.8 % of the money and Proposers get 39.4 %. Without Responder delegation, the Responder gets 39.2 % and the Proposer 49.2 %. Responders tend to design the compensation schemes such that delegates who return with substantial shares of the money are rewarded more, but these shares are far from the predicted maximal share. Indeed, not a single Responder chooses the theoretically predicted compensation scheme (see their Table A6). Responders are thus reluctant to fully exploit their commitment power by choosing the predicted compensation scheme. This is perhaps not so surprising: When using a delegate with an observable compensation scheme, the Responder in effect is Proposer in the ultimatum game and the Proposer becomes the Responder. The finding that the Responder does not induce the delegate to accept only very large offers is thus evidence of the usual moderating fairness and reciprocity effects seen in ultimatum games.

Poulsen and Tan (2006) use the ultimatum game to study how the willingness to exercise a commitment opportunity interacts with whether the opponent will choose to observe the commitment or not. In their main treatment the Responder commits to a smallest acceptable offer (SAO). The Proposer at the same time decides whether or not, to observe (at no cost) the Responder's SAO. Then the Proposer makes an offer and the Responder accepts or rejects, as determined by the chosen SAO. This move protocol is somewhat similar to the No Commitment treatment in the present study, except that here player A (corresponding to the Responder in Poulsen and Tan) moves physically before player B (Proposer). Moreover, and crucially, their experiment does not contain an analogue to our Commitment treatment.

The data show that many Proposers refuse to be informed about the Responder's SAO and simply offer one-half. Also, among those Proposers who decide to be informed many refuse to offer most of the money even when they know that the Responder has chosen the largest SAO. Responders tend to choose larger SAO when Proposers can be informed, compared to the baseline where Proposers cannot decide to condition offers on the Responder's SAO. Nevertheless, the increase is much smaller than predicted. Thus, much as in Fershtman and Gneezy (2001), the willingness to aggressively exercise the commitment option is dampened by the, justified, belief that the opponent will be influenced by fairness and reciprocity concerns.

As already mentioned in the Introduction, in the present experiment the strategic move of avoiding information does not give the person making it much more than the other, but merely neutralizes the opponent's initial advantage. In this sense the commitment move of avoiding information is less 'aggressive' than the commitments in Fershtman and Gneezy (2001) and Poulsen and Tan (2006). Viewed in this light, it is natural that most players in the current experiment over time learn to optimally exercise their commitment option.

Andreoni (2005) examines commitment in the trust game (Berg, Dickhaut, and McCabe, 1995). In his buyer-seller version of the game, the buyer (trustor) decides how much to pay to the seller (trustee) and the seller then decides on the quality of the item. In this game, the buyer will pay little since the seller will always choose low quality. In the game with commitment, the seller can choose to move first and

irrevocably commit to a ‘satisfaction guaranteed’ policy, according to which the buyer has the right to annul the deal and get his money back. This guarantee is predicted to enhance efficiency, since the buyer no longer has to worry about being cheated by the seller.

Andreoni’s experimental data shows that most sellers over time learn to offer the guarantee, and efficiency is significantly enhanced. In this sense, his results complement ours. However, an important difference between Andreoni’s and our design is that the strategic move of offering a satisfaction guaranteed is predicted to (weakly) enhance *both* parties welfare. In our experiment the strategic move is predicted to redistribute surplus from one side to another (although, as discussed earlier, our data show that there is indeed a positive effect on overall efficiency). Intuitively, the fact that in Andreoni’s design the strategic move makes both parties better off makes it more attractive and obvious for subjects to make the strategic move in his than in our design.

Van Huyck, Battalio, and Walters (1995) study commitment in a peasant-dictator game, a game very similar to the trust or investment game (Berg, Dickhaut, and McCabe, 1995). In the Discretion version of their game, the peasant decides how much to invest and the dictator then decides how much of the investment return to tax. In the subgame-perfect equilibrium the dictator appropriates any surplus created by the peasant, and hence the peasant makes no investment. In the second version, Commitment, the dictator irrevocably commits to a tax rate before the peasant makes his investment decision. This leads the dictator to choose a lower tax rate and so the peasant makes a positive investment. Both the dictator and the peasant are (for the peasant, weakly) better off under Commitment than under Discretion. Their experimental data support these predictions: tax rates and investments depend significantly on the treatment, and efficiency in the Commitment treatment is twice as large as under Discretion.

In Van Huyck et al’s experiment the Dictator exogenously moves before or after the peasant; he cannot choose *when* to set the tax rate (before or after the peasant moves). Thus their experiment does not consider the issue we are interested in, namely whether dictators would understand that their tax revenue is larger if they committed to a tax rate before the peasant moves. Also, in their study the commitment move is, as in Andreoni (2005), predicted to be efficiency-enhancing, whereas it in our design is predicted to be efficiency-neutral and to re-distribute money from one player to another.

6.2 Commitment and Physical Timing Effects

In an interesting paper, Huck and Müller (2005) investigate whether the option to ‘burn money’ in a Battle-of-the-Sexes game helps a player to select his preferred outcome (see Ben-Porath and Dekel, 1992, and Van Damme, 1989). The theoretical prediction, using a forward-induction argument, is that a player will decide not to burn money and that the players’ will select the player’s preferred equilibrium. Huck and Müller find that when subjects move sequentially, the theory predicts actual behavior well. But when the game is implemented in the normal form, behavior is very different. This difference leads Huck and Müller to consider the hypothesis (among

others) that the difference is driven by a physical timing effect, namely by the fact that the player with the opportunity to burn money moved first. Investigating certain variations of the basic game, their data reveal a strong physical timing effect that goes beyond the ‘virtual observability’ effect already known in the literature (see for example Weber, Camerer, and Knez, 2006).

We did not find strong physical timing effects, neither on player B’s information decision nor on the demands made. The results from the existing experimental literature help to explain why. First, our underlying bargaining game is sequential and asymmetric. This means that physical timing cannot work as a symmetry breaker (as in Huck and Müller, 2005). Second, since player B makes his demand after player A in all our treatments, player B is in a ‘weak’ position. The fact that in the Unobserved Commitment treatment player B makes an unobserved information decision before player A can not, according to the data, compensate for this weakness. Third, in the Unobserved Commitment game, deciding not to see player A’s demand is a weakly dominated strategy for player B. As pointed out by Güth, Huck, and Rapoport (1998), these factors tend to weaken any physical timing effect. See also the findings in Weber, Camerer, and Knez (2006) and Huck and Müller (2005).

6.3 The Observability of Commitment

It is intuitive and well-known that observability of the strategic move is crucial, see Schelling (1960). Just how crucial observability is made clear in Bagwell’s (1995) influential article. Analysing the Stackelberg game, Bagwell assumes there is a, possibly very small, probability that the follower receives a wrong signal about the leader’s chosen action. He shows that this eliminates the leader’s first mover advantage. More precisely, the only pure-strategy Nash equilibrium of the game with imperfect observability is the Cournot equilibrium. Bagwell’s provocative result sparked a rich subsequent theoretical and experimental literature, investigating the robustness of his analysis. See, for example, Van Damme and Hurkens (1997), Huck and Müller (2000), Güth, Müller, and Spiegel (2006), Hurkens and Vulkan (2006), and the references therein.

Fischer et. al. (2006) conduct an experiment using a sequential bargaining game. When the first mover makes a demand, he knows only that there is an exogenous probability that the second mover will observe the demand. If the second mover observes the first mover’s choice, an ultimatum game is played, otherwise payoffs are determined as in the Nash demand game. Their theoretical model predicts that as the observation probability increases, there should be a transition from Nash demand game behavior (an equal split) to ultimatum game behavior (the first mover gets the lion’s share). Fischer et. al. experimentally vary this probability and find that, as predicted, the higher is the observation probability, the (weakly) higher are the first movers’ demands. Also, the second movers, when uninformed, make lower demands when the probability is high, correctly anticipating a higher demand by the first mover.

In Fischer et. al.’s setup, the probability with which the first mover’s choice will be seen by the second mover is exogenous. Given the importance of the extent and nature of observability, it is interesting to endogenize the follower’s decision whether or not to observe the leaders’ choice, and our work can be seen as such a contribution. We give

the follower a simple observe/not observe choice. Our theoretical and empirical results are straightforward: the follower should decide not to see the leader's action only when this decision is observed by the leader. Our data show that subjects learn this over time. In the light of the set-up in Fischer et. al. (2006), it could be interesting to adapt our model such that player B, instead of making a yes/no decision, decides on a probability of seeing A's demand. In the Commitment game, player A then observes the probability but not the realization. This version seems to correspond to the set-up in Fischer (2006) if the observation probability in the latter model were chosen by the second mover and so endogenous. We conjecture that the second mover would then decide to set the probability equal to zero, as is done in the Commitment game. I thank Werner Güth for indicating this relationship between the models.

In an interesting paper, Vardy (2004), using a theoretical model, examines the role of costly observability. In Vardy's model, the follower can, after the leader has made his choice, pay to observe the choice (this move structure corresponds to our No Commitment Game). If the follower pays, he observes the leader's action without error. Vardy shows that any strictly positive observation cost eliminates the Stackelberg equilibrium. However, for low cost there exists a mixed 'noisy' equilibrium that, as the cost becomes small, converges to the Stackelberg equilibrium. In an experimental companion paper, Morgan and Vardy (2004) investigate the empirical relevance of these different equilibria. They vary the observation cost and find that the frequency with which followers pay to observe the leader's choice is a decreasing function of the observation cost. The leader keeps his first mover advantage for low cost, but observed overall play is inconsistent with the noisy equilibrium.

6.4 Games with Endogenous Moves

Our paper is related to the literature endogenising the order of moves in games. These games are called 'timing games'. See for example Fonseca, Huck, and Normann (2005), and the references therein. Most of the experiments have been using duopoly games, and our design and results can be seen as making a contribution to the endogenous timing literature by studying bargaining. More precisely, the formal structure of our Commitment game corresponds to a bargaining game with endogenous timing and with observable delay (Hamilton and Slutsky, 1990). Player B decides visibly to player A whether to play a sequential Nash demand game with himself as informed second-mover or as an uninformed second-mover. The latter game is equivalent to a simultaneous-move Nash demand game. In the Unobserved Commitment game, player B makes the same decision but player A does not know what game B has chosen to play. Still, choosing to be an informed second mover is as described earlier a weakly dominant strategy for player B in the Unobserved Commitment game, and we observe that most player Bs behave as predicted. The predictions made by the theory of endogenous timing are borne out for our situation, when sufficient learning has taken place.

7 Conclusion

Game theory predicts that players make strategic commitments, such as ‘burning one’s bridges’. Since such strategic moves can appear counterintuitive and could be too complicated for people to use in practice, we conducted an experiment to see whether people actually make the predicted strategic move. Our data show that the experimental subjects do, but after substantial time, learn to make the predicted strategic move. Our experiment was simple in that it did not explicitly model the commitment technology. Future work should consider this and also how players create their commitment opportunities. The author is currently involved in such experimental work and hopes he can soon report some experimental results.

8 Appendix

8.1 Additional Figures




Figure 4: Relative frequency distribution of participant A and B demands in the Commitment treatment, conditional on B's information decision.

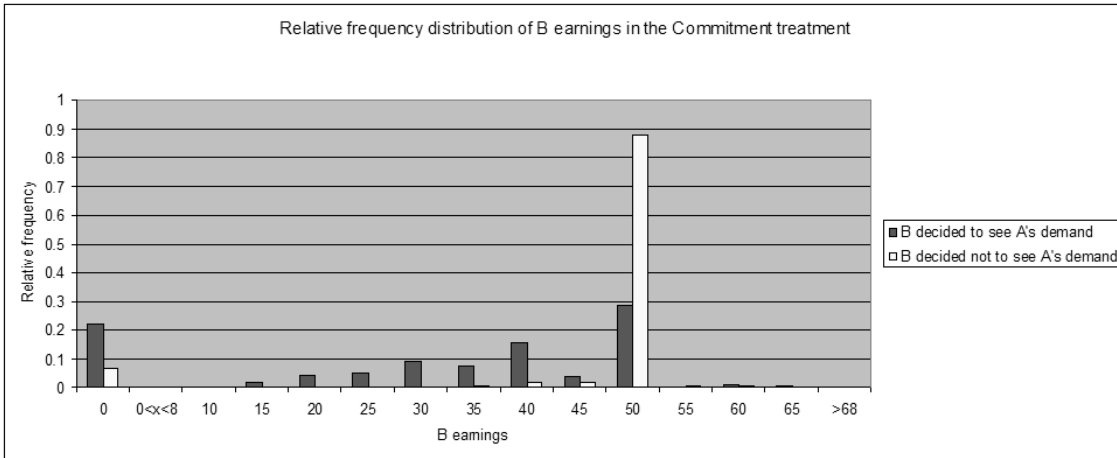
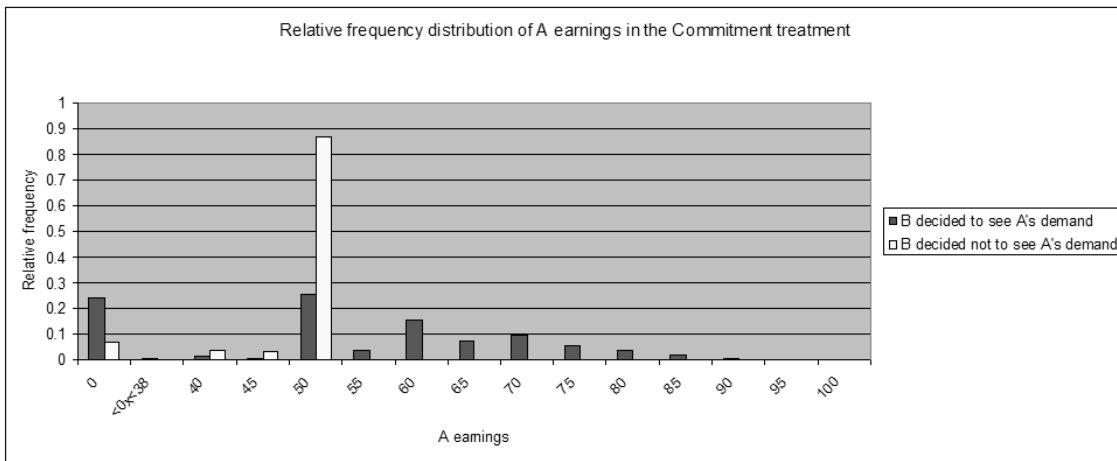


Figure 5: Relative frequency distribution of participant A and B earnings in the Commitment treatment, conditional on B's information decision.

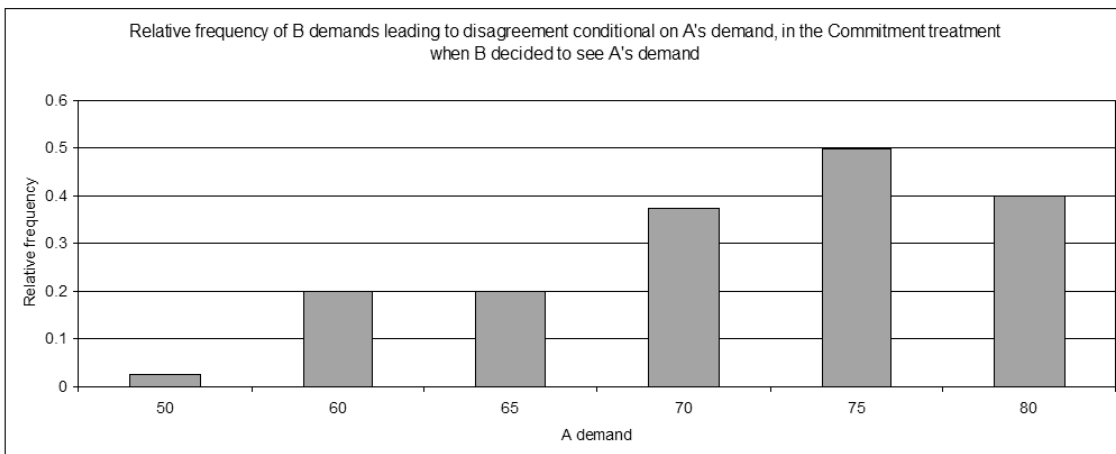
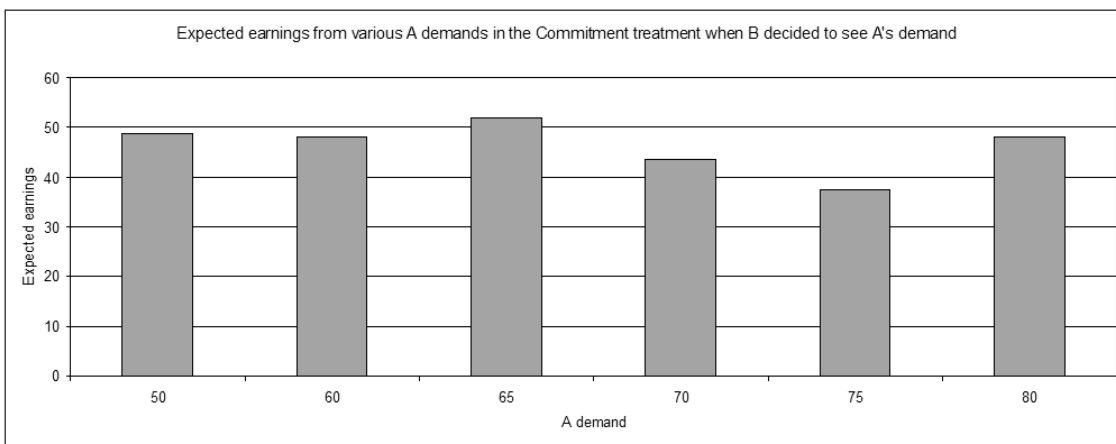


Figure 6: The top panel shows, for the Commitment treatment when B decided to see participant A's demand, the expected payoffs to various A demands given the empirical B demands. The lower panel shows, for the Commitment treatment when B decided to see A's demand, the proportion of B demands that, conditional on a given A demand, led to disagreement. Note: participant A demands with frequency below 10 have been ignored.


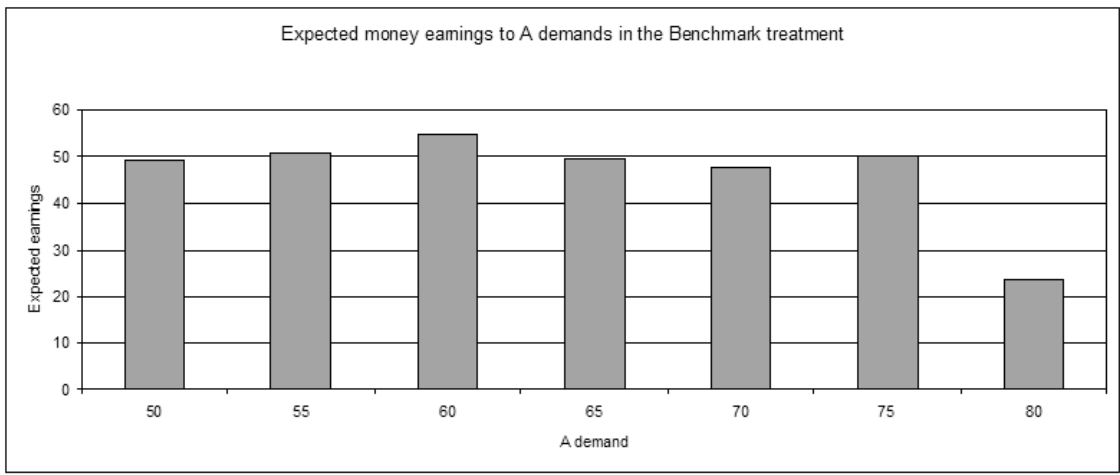


Figure 7: The expected payoffs to various participant A demands, given the observed B demand behavior, in the Benchmark (top panel) and No Commitment treatments (lower panel). Note: participant A demands with frequency below 10 have been ignored.





Figure 8: Relative frequency distribution of A and B demands in various treatments.

8.2 Instructions

The instructions were in English and are reproduced below.

8.2.1 Benchmark Treatment

Welcome to this experiment! The purpose of the experiment is to analyze decision making. Please read these instructions carefully. From now on please do not communicate with the other participants. If you have any questions, give notice by raising your hand. We will then answer your questions privately.

All participants in this experiment have been recruited in the same way as you, and you all have the same instructions.

In the experiment you can earn money. How much money you earn depends on your decisions and on the decisions made by the other participants you will interact with. During the experiment all money amounts are given in ‘points’. At the end of the experiment, the total amount of points you have earned is converted into Danish Kroner, by using the following exchange rate:

5 points is equal to 1 Danish Kroner

All decisions remain anonymous. You will not be informed about the identity of other participants, and no other participant will be informed about your identity. You receive your final payment in Danish Kroner in cash and this is done in private.

The experiment consists of 15 periods. At the start of the experiment, you will be randomly given one of two different roles, A and B. Half of the participants are given the A role and the other half are given the B role. Those participants given the A role are called ‘A participants’, or just ‘A’. Those given the B role are called ‘B participants’, or just ‘B’. Your role is the same in all 15 periods.

In each period one A participant will interact with one B participant. It will be randomly decided at the beginning of each period which A participant will interact with which B participant. You interact with the other participant through the computer terminal you are placed in front of.

Calculation of earnings:

In each period, there is a sum of 100 points available. The A and the B participant each have to demand a number of points from this sum. This demand must be a whole number between 0 and 100 (both included). If the sum of the demanded points made by A and B does not exceed 100 points, each earns the number of points that he/she demanded. If the sum of the two demands exceeds 100 points, each of you earns zero points. This is repeated here: Suppose A demands a points and suppose B demands b points. Then earnings for this period are calculated as follows:

- If $a + b \leq 100$, A gets a points and B gets b points.
- If $a + b > 100$, A gets zero points and B gets zero points.

Timing of decisions:

In each period, A makes his/her demand first. Then B sees what A has demanded, and B next makes his/her own demand. Finally both participants are informed about the two demands, the total demand, and their own earnings. This is repeated here:

1. A makes his/her demand
2. B sees what A has demanded
3. B makes his/her demand
4. The two participants are informed about both participant's choices and their own earnings

After this a new period begins, and you will be randomly matched with one of the participants in the other role. After the last period, the points you have earned in each of the 15 periods will be added and converted into Danish Kroner, using the exchange rate from above. This number of Danish Kroner is then paid to you in cash, together with the show-up fee of 50 Danish Kroner.

The computer screens:

The computer screens used in this experiment are quite simple to navigate in. However, if you during the experiment have any questions regarding how to use the screens, please raise your hand and one of the experimenters will come to you. On the computer screen, in the upper right corner, you will see how much time (in seconds) remains to make your decision. Please note that even if you run out of time, you can still make a decision, but please try to respect the time limit. Demands are made by activating the 'My demand' box using the mouse and then typing in your demand using the keyboard. When you have made your decision, please press the 'OK' button in the lower right corner. Similarly, when you have read the summary of the decisions, press the 'Continue' button to proceed.

If you at any time need to make some calculations, you are welcome to use the pen and the blank sheets of paper that have been provided.

In a moment we will ask you to answer some questions in a short test to make sure you have understood the rules.

8.2.2 No Commitment Treatment

The instructions for this treatment differs from the other treatment instructions only in the paragraph 'Timing of Decisions', which is reproduced here:

Timing of decisions:

In each period, A makes his/her demand first. After this B decides whether or not he/she would like to see what A has demanded. Then B makes his/her own demand. Finally both are informed about whether B observed A's demand, the two demands, the total demand, and their own earnings. This is repeated here:

1. A makes his/her demand
2. B decides if he/she wants to see what A has demanded

3. One of the following happens:
 - a. If B in 2 decided to see A's demand, then B sees what A has demanded and B makes his/her own demand
 - b. If B in 2 decided not to see A's demand, then B makes his/her demand without knowing what A has demanded
4. The two participants are informed about both participant's choices and their own earnings

8.2.3 Commitment Treatment

The instructions for this treatment differs from the other treatment instructions only in the paragraph 'Timing of Decisions', which is reproduced here:

Timing of decisions:

In each period, A makes his/her demand before B. But before this B decides whether or not he/she would like to observe A's demand, and A is informed about B's decision. Then A makes his/her demand. Depending on B's decision, B either sees A's demand or not. B then makes his/her own demand. Finally both are informed about whether B decided to observe A's demand, the two demands, the total demand, and their own earnings. This is repeated here:

1. B decides whether or not he/she wants to see A's demand before B makes his/her own demand
2. A is informed about B's decision in 1
3. A makes his/her demand
4. One of the following happens:
 - a. If B in 1 decided to see A's demand, then B sees what A has demanded and B makes his/her own demand
 - b. If B in 1 decided not to see A's demand, then B makes his/her demand without knowing what demand A made
5. The two participants are informed about both participant's choices and their own earnings

8.2.4 Unobserved Commitment Treatment

The instructions for this treatment differs from the other treatment instructions only in the paragraph 'Timing of Decisions', which is reproduced here:

Timing of decisions:

In each period, B first decides whether he/she wants to see A's demand or not before B makes his/her own demand. B's decision to see A's demand or not is not seen by A. Then A makes his/her demand. Depending on B's decision, B either sees

A's demand or not. B then makes his/her own demand. At the end of the period, both are informed about whether B decided to see A's demand or not, the two demands, the total demand, and their own earnings.

All this is described in detail here:

1. B decides whether or not he/she wants to see A's demand before B makes his/her own demand. B's decision to see A's demand or not is not seen by A
2. A makes his/her demand
3. One of the following happens:
 - a. If B in 1 decided to see A's demand, then B sees what A has demanded and B makes his/her own demand
 - b. If B in 1 decided not to see A's demand, then B makes his/her demand without knowing what demand A made
4. The two participants are informed about their choices and their own earnings

8.3 Tests

The following tests were handed out after participants had read the instructions.

8.3.1 Benchmark Treatment

Please answer the questions below. When you have finished please raise your hand and one of the experimenters will come to you and check your answers.

Question 1: Considering the roles (A and B), which of the following statements are correct? (Please indicate your answer by 'X')

It will be randomly determined at the beginning of each period which role you have:

You will have each role every second period:

You will have the same role throughout the experiment, but which one has been randomly determined:

Question 2: Considering the matching with another participant, which of the following statements are correct? (Please indicate your answer by 'X')

It will be randomly determined at the beginning of each period who you interact with:

You will interact with the same other participant throughout the experiment:

Question 3: Suppose A demands 30 points and B demands 60 points. How many points will each earn?

A earns: points

B earns: points

Question 4: Suppose A demands 60 points and B demands 45 points. How many points will each earn?

A earns: points

B earns: points

Question 5: Suppose you are A deciding how much to demand. Do you know how much B has demanded? (Please indicate your answer by 'X')

Yes:

No:

Question 6: Suppose you are B deciding how much to demand. Do you know how much A has demanded? (Please indicate your answer by 'X')

Yes:

No:

8.3.2 Unobserved Commitment and Commitment Treatment

Please answer the questions below. When you have finished please raise your hand and one of the experimenters will come to you and check your answers.

Question 1: Considering the roles (A and B), which one of the following statements is correct? (Please indicate your answer by 'X')

It will be randomly determined at the beginning of each period which role you have:

You will have each role every second period:

You will have the same role throughout the experiment, but which one has been randomly determined:

Question 2: Considering the matching with another participant, which one of the following statements is correct? (Please indicate your answer by 'X')

It will be randomly determined at the beginning of each period who you interact with:

You will interact with the same other participant throughout the experiment:

Question 3: Suppose A demands 30 points and B demands 60 points. How many points will each earn?

A earns: points

B earns: points

Question 4: Suppose A demands 60 points and B demands 45 points. How many points will each earn?

A earns: points

B earns: points

Question 5: When A makes his/her demand, does he/she know how many points B has demanded? (Please indicate your answer by 'X')

Always:

Depends on A's own decision:

Depends on B's decision:

Never:

Question 6: When B makes his/her demand, does he/she know how many points A has demanded? (Please indicate your answer by 'X')

Always:

Depends on A's decision:

Depends on B's own decision:

Never:

Question 7: Suppose A is deciding how much to demand. Does he/she know if B will see A's demand? (Please indicate your answer by 'X')

Always:

Depends on A's own decision:

Depends on B's decision:

Never:

8.3.3 Unobserved Commitment

Please answer the questions below. When you have finished, raise your hand and one of the experimenters will come and check your answers.

Question 1: Concerning the roles (A and B), which of the following statements is correct? (Please indicate your answer with a 'X')

It will be randomly determined at the beginning of each period which role you have:

You will have each role every second period:

Your role will be randomly determined at the start, and then you will keep the same role throughout the experiment:

Question 2: Concerning the matching with another participant, which of the following statements is correct? (Please indicate your answer with a 'X')

It will be randomly determined at the beginning of each period who you interact with:

You will interact with the same other participant throughout the experiment:

Question 3: Suppose A demands 30 points and B demands 60 points. How many points will each earn?

A earns: points

B earns: points

Question 4: Suppose A demands 60 points and B demands 45 points. How many points will each earn?

A earns: points

B earns: points

Question 5: When A makes his/her demand, does he/she know how many points B has demanded? (Please indicate your answer with a 'X')

Always:

Depends on A's own decision:

Depends on B's decision:

Never:

Question 6: When B makes his/her demand, does he/she know how many points A has demanded? (Please indicate your answer with a 'X')

Always:

Depends on A's decision:

Depends on B's own decision:

Never:

Question 7: Suppose A is deciding how much to demand. Does he/she know if B will see A's demand? (Please indicate your answer with a 'X')

Always:

Depends on A's own decision:

Depends on B's decision:

Never:

9 References

Aghion, P. and Bolton, P. (1987). "Contracts as a Barrier to Entry", *American Economic Review*, 77 (3), 388-401.

Andreoni, J. (2005). "Trust, Reciprocity, and Contract Enforcement: Experiments on Satisfaction Guaranteed". Working paper, available at: <http://econ.ucsd.edu/~jandreoni/>

Bagwell, K. (1995). "Commitment and Observability in Games", *Games and Economic Behavior*, 8 (2), 271-280.

Bensaid, B. and Gary-Bobo, R. (1993). "Commitment Value of Contracts Under Renegotiation Constraints", *Econometrica*, 61 (6), 1423-1429.

Ben-Porath, E. and Dekel, E. (1992). "Co-ordination and the potential for sacrifice", *Journal of Economic Theory*, 57 (1), 36-51.

Berg, J., Dickhaut, J., and McCabe, K. (1995). "Trust, Reciprocity, and Social History", *Games and Economic Behavior*, 10 (1), 122-142.

Caillaud, B., Jullien, B., and Picard, P. (1995). "Competing Vertical Structures: Precommitment and Renegotiation", *Econometrica*, 63 (3), 621-646.

Camerer, C. (2003): *Behavioral Game Theory*. Princeton, N.J.: Princeton University Press.

Van Damme, E. (1989). "Stable Equilibria and Forward Induction", *Journal of Economic Theory*, 48 (2), 476-296.

Van Damme, E. and Hurkens, S. (1997). "Games with Imperfectly Observable Commitment", *Games and Economic Behavior*, 21 (1-2), 282-308.

Fershtman, C. and Gneezy, U. (2001). "Strategic delegation: An experiment", *RAND Journal of Economics*, 32 (2), 352-368.

Fershtman, C. and Kalai, E. (1997). "Unobserved Delegation", *International Economic Review*, 38 (4), 763-774.

Fischer, S., Güth, W., Müller, W., and Stiehler, A. (2006). "From ultimatum to Nash bargaining: Theory and experimental evidence", *Experimental Economics*, 9 (1), 17-33.

Fischbacher, U. (1999). *z-Tree - Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual*, Working Paper Nr. 21, Institute for Empirical Research in Economics, University of Zurich.

Fonseca, M., Huck, S., and Normann, H. (2005). "Playing cournot although they shouldn't", *Economic Theory*, 25 (3), 669-677.

Greiner, B., (2004). *An Online Recruitment System for Economic Experiments*. In: Kurt Kremer, Volker Macho (Eds.). *Forschung und wissenschaftliches Rechnen 2003*. GWDG Bericht 63, Göttingen: Ges. fr Wiss. Datenverarbeitung, 79-93.

Güth, W., Huck, S., and Rapoport, A. (1998). "The limitations of the positional order effect: Can it support silent threats and non-equilibrium behavior?", *Journal of Economic Behavior and Organization*, 34 (2), 313-325.

Güth, W., Müller, W., and Spiegel, Y. (2006). "Noisy leadership: an experimental approach", *Games and Economic Behavior*, 57 (1), 37-62.

Hamilton, J. and Slutsky, S. (1990). "Endogenous timing in duopoly games: Stackelberg or Cournot equilibria", *Games and Economic Behavior*, 2 (1), 2946.

Hauk, E. and Hurkens, S. (2001). "Secret information acquisition in Cournot markets", *Economic Theory*, 18 (3), 661-681.

Huck, S. and Müller, W. (2000). "Perfect versus imperfect observability – an

experimental test of Bagwell's result", *Games and Economic Behavior*, 31 (2), 174-190.

Huck, S. and Müller, W. (2005). "Burning money and (pseudo) first-mover advantages: an experimental study on forward induction", *Games and Economic Behavior*, 51 (1), 109-127.

Hurkens, S. and Vulkan, N. (2006): "Endogenous private information structures", *European Economic Review*, 50 (1), 35-54.

Van Huyck, J. B., Battalio, R. C., and Walters, M.F. (1995). "Commitment versus Discretion in the Peasant-Dictator Game", *Games and Economic Behavior*, 10 (1), 134-170.

Loewenstein, G., Moore, D., and Weber, R. (2006). "Misperceiving the value of information in predicting the performance of others", *Experimental Economics*, 9 (3), 281-297.

Morgan, J. and Vardy, F. (2004): "An experimental study of commitment in Stackelberg games with observation cost", *Games and Economic Behavior*, 49 (2), 401-423.

Muthoo, A. (1996). "A Bargaining Model Based on the Commitment Tactic", *Journal of Economic Theory*, 69 (1), 134-152.

Nash, J. (1953). "Two Person Cooperative Games", *Econometrica*, 21 (1), 128-140.

Poulsen, A. and Tan, J.H.W. (2006). "Information acquisition in the ultimatum game: An experimental study", *Experimental Economics*. Available on-line at <http://www.springerlink.com/content/86q80v1743517m8m/>.

Schelling, T. (1960). *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

Tirole, J. (1988). *Industrial Organization*. Cambridge, MA: MIT Press.

Vardy, F. (2004). "The value of commitment in Stackelberg games with observation cost", *Games and Economic Behavior*, 49 (2), 374-400.

Weber, R., Camerer, C., and Knez, M. (2006). "Timing and Virtual Observability in Ultimatum Bargaining and "Weak Link" Coordination Games", *Experimental Economics*, 7 (1), 25-49.