Timing and Virtual Observability in Ultimatum Bargaining and 'Weak Link' Coordination Games

Colin F. Camerer
Division of Social Sciences 228-77
California Institute of Technology
Pasadena CA 91125
camerer@hss.caltech.edu

Roberto A. Weber
Department of Social and Decision Sciences
Carnegie Mellon University
Pittsburgh PA 15213
rweber@andrew.cmu.edu

Marc Knez
University of Chicago
Graduate School of Business
Chicago IL 60637
Marc.knez@gsb.uchicago.edu

Abstract

Previous studies have shown that simply knowing some players move first can affect behavior in games, even when the first-movers' moves are unobservable. This observation violates the game-theoretic principle that timing of unobserved moves is irrelevant. However, this previous research only shows that timing matters in games where knowledge that one player moved first can help select that player's preferred equilibrium. We extend this work by varying timing of unobservable moves in ultimatum bargaining games and "weak link" coordination games. Timing without observability affects both bargaining and coordination, but only weakly. The results are consistent with theories that allow "virtual observability" of first-mover choices, rather than theories in which timing matters <u>only</u> because first-mover advantage is used as a principle of equilibrium selection.

First Draft: 20 May 1996. Please do not circulate but comments are welcome. We thank Gary Bolton, Yuval Rottenstreich, participants in the Chicago GSB Behavioral Science brown bag workshop, the Wharton Decision Processes workshop, and the 1996 Public Choice Society meeting for ideas, and NSF grant SBR 95-11001 for financial support.

I. Introduction

Imagine a game in which somebody moved before you, but you don't know what they did. Does it matter that they moved already? The standard answer in game theory is "No." Perhaps surprisingly, the empirical answer discovered in experiments is "Yes, it can matter". This paper adds to the evidence that timing of unobserved moves can matter, and reports experiments designed to distinguish between two different theories that may explain why.

In noncooperative games of imperfect information, a player's ignorance about the move of another player who moved before her is represented in an extensive form game tree by linking all the nodes which could result from the other player's earlier moves, in an "information set". The same information set representation is used when moves are simultaneous. This graphical convention effectively implies that the time at which unobservable moves were made doesn't matter.

Timing-irrelevance is not merely the result of a graphical convention used to draw trees. It is an "invisible assumption" in game theory, which follows from the more basic principle that if the outcome of an event is unknown, it doesn't matter when the event happened-- unobservability trumps timing. Bagwell (1994) makes this point in a different way: In a class of games in which there is an advantage to moving first (e.g., in Cournot duopoly), he shows that the commitment value of moving first is severely undermined if observability of the earlier move is even slightly in doubt.

Rapoport (1997) points out that in the earliest development of game theory, von Neumann and Morgenstern (1947, p. 51) recognized the distinction between "anteriority" (priority in time) and "preliminarity" (priority in information). Preliminarity implies anteriority, but not vice versa, and thus may seem more fundamental.[1] Having recognized the distinction between timing and observability, they opted to make priority in information the basic way of characterizing strategies, defined strategies with no reference to chronological order of moves, and effectively banished timing, per se, from game theory.

Several previous experimental studies indicate that even when moves are unobservable, timing can matter. However, these previous studies confound two possible theories for the observed effect of timing. The "first mover advantage" theory is that all players agree that moving first has an advantage and expect first movers to exploit it, even when the first mover's choice will not be observed. However, a more general explanation of the timing effect, which we call "virtual observability", is that players expect first movers to choose strategies as if subsequent players observe them perfectly and respond optimally.

[1] Having information about a previous move implies the move happened earlier, but an early move need not be known.

Note that if there is a first mover advantage, the virtual observability explanation predicts that it will be exercised.

However, virtual observability also predicts timing effects in games without first mover advantages, such as the "weak-link" coordination game (Van Huyck, Battalio and Beil, 1991). Coordination problems are important in much economic activity (e.g., Cooper, 1999). As mentioned previously, experiments have shown that timing of moves (even without information) can affect outcomes. One game in which this is the case is the battle-of-the-sexes game, described below. In this game, which is a coordination game different than the weak-link game, simply having one player move first improves coordination. A natural question, therefore, is to what extent simply changing the timing of moves can improve coordination in other games. If having one person move first improves coordination in weak-link games as well, then this might present a useful insight into solving these types of problems.

Our paper explores this possiblity and adds to the literature on timing effects by reporting new evidence from weak-link coordination games and ultimatum bargaining games.. The comparison of bargaining and coordination enables us to distinguish between the two different explanations for the effect of timing. Our evidence suggests the first-mover advantage theory is not the whole story and virtual observability may be the better general explanation for previous results.

The next two sections discuss previous work on timing of moves in games and possible explanations of the discovered effects. Two sets of experiments intended to test the explanations are then reported.

II. Previous Results and Explanations

Previously reported timing effects come from games with multiple equilibria, in which one equilibrium is preferred by one player and another by a different player. An example from Cooper et al (1993) is shown in Table 1. The game is a "battle-of-the-sexes" in which the row player prefers the Nash equilibrium (2, 1) and column prefers the equilibrium (1,2). There is also a mixed-strategy equilibrium in which players believe others will choose 1 with probability .25 and 2 with probability .75. Note that if players move sequentially (with the ROW player moving first), and previous moves <u>are</u> observed, the ROW player should choose 2 and, observing that, COLUMN should choose 1. The player who moves first earns the high payoff (600) and the second-mover earns the low payoff. So there is an advantage to moving first.

Table 1: Battle-of-the-sexes (BOS) Game - Results from Cooper et al (1993)

		Column		frequency of choices	
		1	2	simul.	seq'l.
ROW	1	0, 0	200, 600	38%	12%
	2	600, 200	0, 0	62%	88%

simultaneous:	35%	65%
sequential:	70%	30%

The choice frequencies in Table 1 show that when moves were simultaneous both players approached the mixed-strategy equilibrium, choosing strategy 2 either 62% (ROW) or 65% (COLUMN) of the time. In the sequential condition, ROW moved first and the order of moves was commonly known, but ROW's move was not known to COLUMN. Then ROW players chose strategy 2 88% of the time and, more strikingly, COLUMN players chose strategy 1 70% of the time. Thus, merely knowing that ROW moved first caused players to move toward the equilibrium (2, 1), which ROW preferred, even though COLUMN did not know ROW's previous move when she chose. Similar results were conjectured by Kreps (1990, pp. 100-101) and reported in informal experiments by Amershi et al (1989), Muller and Sadanand (1998), and other investigators (see Cooper et al, 1993, footnote 6).

Rapoport (1997) studied a three-player BOS. As in the two-player studies mentioned above, players chose in a predetermined order, but did not know the moves of the previous players. About 60% of the players chose the strategy which gave the first-mover her preferred outcome. The percentages of players choosing the first-mover's preferred equilibrium did not vary much across the move order (66.3%, 62.9%, and 59.2% for first- through third-movers). An analysis of individual subjects indicated that about 40% of the subjects complied with the first-mover's preference 85% of the time, and another 20% of the subjects complied 60% of the time.

A series of papers by Amnon Rapoport (1997) and colleagues report substantial timing effects in "resource dilemma" games. In these games, players draw from a "common resource pool" of either certain or uncertain size. If the resource pool is overdrawn nobody gets anything. When players move in order and previous moves are observable, first-movers take more of the pool, leaving less for later movers. When the order of moves is commonly known but the first movers' actual draws are unknown, Budescu, Au, and Chen (1997) report that first movers do take more and later-moving players take less. In a five-player game, first movers draw 28%, third movers draw 23%, and last-movers draw 20% of the pool's expected size. These figures are close to the corresponding fractions in games where earlier moves are observable, which are 34%,25%, and 20% (Rapoport, Budescu, and Suleiman, 1993).

In two- and three-player games the results are weaker but still significant (the corresponding first- and last-mover fractions are 52% vs. 48%, and 34% vs. 3 1 %). Perhaps more importantly, players generally expect that those moving before them will have taken more than those moving after them will take. Since the only pure-strategy Nash equilibria involve various divisions of the entrie resource pool (or some certainty-equivalent, when the size of the pool is commonly-unknown), this result is much like the BOS result-- first movers get more, even when their moves are unknown.

Rapport (1997) also studied a "step-level" or threshold public goods game, in which a public good is provided if 4 (or more) out of 7 players contribute. The payoffs were chosen so that if exactly three others contribute, a player prefers to contribute; otherwise she prefers not to contribute. These games have many pure-strategy equilibria in which a subset of four players contribute. In Rapoport's study, as above, players know the order in which they choose to contribute but not the contributions of those who moved previously. The effect of timing is striking: Only 18% of the first three players contribute, but 38% of the last three players contributed. The early-movers seem to expect the later-movers to shoulder more of the burden of contributing, and the later-movers do.

Finally, Guth, Huck, and Rapport (1998) studied symmetric and asymmetric 2-player BOS games and a class of 2-player games in which a timing effect would entail disequilibrium play. In their asymmetric BOS game, the preferred equilibrium for the second mover gave each player equal payoffs, while the other pure strategy equilibrium gave the first mover a higher payoff. They report several important findings. One finding is that there is no timing effect when the prediction entails disequilibrium play. This is not too surprising, but it provides a boundary to the timing effect prediction. They also find that the timing effect is much weaker in the presence of a fair (equal-payoff) outcome. They conclude that the "first-come-first-serve" rule implied by the timing effect is weaker than fairness norm.

Explanations of timing effects

Two explanations spring to mind for why simply knowing that one player moved first matters in these games.[2]

1. <u>First mover advantage:</u> One explanation is that in games with multiple equilibria (including the BOS game in Table 1), the players solve their coordination problem by grasping for any asymmetry or "psychologically prominent" focal principle (Schelling, 1960) that selects one equilibrium over another. The fact that one player moves first is one such feature. Note that if ROW moves first and her move <u>is</u> observable, then (2, 1) is the unique subgame perfect equilibrium. We would expect ROW to choose 2 and COLUMN to begrudgingly best-respond with strategy 1. The first-mover advantage theory is that players select an equilibrium by granting ROW the first-mover advantage

---

[2] Two other possible explanations can be found in psychological principles. First, "causal illusions" occur when people think that actions they take at time t might affect actions of others at time t+1, even though there is no apparent causal mechanism for such an affect. For example, Morris, Sim & Grotto (1995) report that subjects playing a prisoners' dilemma (PD) are more likely to cooperate if they know the player they are paired with moves after them than if the other player moves before them. Players act as if their cooperative choice can magically induce reciprocal cooperation by a player who moves later.

A second psychological difference arising from timing is that players who move later know there is information they could have-- what earlier players chose-- but do not. They may feel more regret if they make a mistake in this situation (since there is something they "could have known"), then if they make an equivalent mistake moving first. Heightened regret can then cause them to act as if they are more averse to uncertainty or ambiguity (see Camerer & Weber, 1992) when they move second.

she would pin if her move was observable, even when it is not. Then first-mover advantage is a selection principle along with other principles like security, salience or focality (Mehta, Starmer & Sugden, 1994), precedent, payoff-dominance (Van Huyck, Battalio & Beil 1990), computational complexity (Ho & Weigelt, in press), and loss-avoidance (Cachon & Camerer, 1996).

The first-mover advantage theory -can be stated more formally:

> First-mover advantage: Fix a game of imperfect information in which previous moves are unobservable and in which there are multiple Nash equilibria. If one equilibrium is preferred by one player and not by others, then when the game is played sequentially, players will choose this Nash equilibrium.

2. Virtual observability: A more general theory is that players act like earlier players' moves are "virtually observable": That is, players expect first movers to choose strategies as if subsequent players observe them perfectly and respond optimally. This idea was first carefully articulated by Amershi, Sadanand, & Sadanand (1989) in a refinement of Nash equilibrium they call MAPNASH, for "manipulated Nash equilibrium". Heuristically, virtual observability means players erase information sets and act like moves will be observable; then they apply subgame perfection. If perfection selects a unique equilibrium which is a Nash equilibrium in the actual game (when the information sets are restored), they play that one. More formally,

> Virtual observability: Fix a game of imperfect information in which previous moves are unobservable. If there is a Nash equilibrium in this game which is also selected by subgame-perfection when previous moves are observable, then players will choose this equilibrium.[3] If not, timing will not affect play.

Note that virtual observability includes first-mover advantage. However, virtual observability is strictly more general than the first-mover advantage theory. The link occurs through subgame perfection. Suppose there are multiple Nash equilibria in the simultaneous-move game, the first mover prefers one of them, and that equilibrium is subgame perfect as well (e.g., the first-mover's preferred equilibrium in BOS). Then according to virtual observability, that equilibrium will played. Subgame perfection creates a natural bias toward (perfect) equilibria that the first mover prefers, since under virtual observability she can move confidently expecting the second-mover to best respond (and, thus, can "force" her most-preferred choice on the second-mover, hence the term "manipulated equilibrium" chosen by Amershi et al). So one outcome of virtual observability is that first movers will gain an advantage if they can. For example, in BOS, (2, 1) is the unique perfect equilibrium when ROW moves first and her move is

[3] This requirement will rule out equilibria, which are subgame perfect when moves are perfectly observed, but are not Nash equilibria in the simultaneous-move or imperfect-observability games. Examples include the "matching pennies" game and the Cournot duopoly game (Bagwell, 1995).

observable by COLUMN. Then virtual observability predicts that they will also play (2, 1) when her move is not observable.[4]

Therefore, whenever first-mover advantage predicts an effect of timing, virtual observability predicts the same effect. The crucial difference between the first-mover advantage theory and virtual observability lies in games where perfection refines a set of Nash equilibria, but the perfect equilibrium does not convey a simultaneous advantage to the first-mover and disadvantage to the second-mover. Table 2 shows an example, known as "stag hunt". In stag hunt both players can move L or H. The equilibrium (H,H) Pareto-dominates (L,L), but choosing L is less risky and (L,L) is an equilibrium too.

Table 2: Stag Hunt Game

		Column	
		L	H
ROW	L	200, 200	200, 0
	H	0, 200	600, 600

If the players move sequentially with observability, (H,H) is the unique subgame-perfect equilibrium in stag hunt, but it does not convey a first-mover advantage per se (since it does not benefit the first-mover at the second-mover's expense). Hence, the first-mover advantage explanation of timing effects predicts no effect of timing. Virtual observability predicts (H,H) will be chosen when one player moves first (even though (L,L) is a Nash equilibrium outcome as well). Our experiments below use an n-player version of stag hunt, the "weak-link" game, to distinguish the virtual observability and first-mover advantage explanations.

The relation between coordination and bargaining

Our experiments seek to understand timing effects by adding observations from two new types of games-- bargaining and "weak-link" coordination. There is an important relation between coordination and bargaining which partly motivates our interest in timing and bargaining.

Theorists have long recognized the inherent coordination problem present in most bargaining. For example, Schelling (1960, p. 69) points out that "the fundamental problem in tacit bargaining is that of coordination." In any game with a range of mutually acceptable outcomes, players seek to coordinate on one of those outcomes-- since failing to agree results in an outcome that is worse for both--while striving to get the most for themselves.

---

[4] Note well that (2, 1) is also a Nash equilibrium in the game with information sets restored.

One widely studied bargaining game in which the issue of coordination has not been given much attention is ultimatum bargaining. In the ultimatum game two players must divide a sum of money X. The first player (labeled the *proposer)* offers some portion x of a pie X to the second player (labeled the *responder).* If the responder accepts the offer then the responder receives x and the proposer receives X - x. If the responder rejects the offer then both players receive nothing.

A typical result from many experiments is that proposers offer around 40% of the amount being divided, and responders reject offers with high frequency if they are less than 20% or so. This basic result has been replicated dozens of times, in several countries and with large variations in stakes (see Camerer and Thaler, 1995, for a review).

In most experiments the ultimatum game is played sequentially. A proposer makes a specific offer, which is transmitted to the responder, who accepts or rejects it. If responders are rational and self-interested, and proposers know that, then the unique subgame perfect equilibrium is for proposers to offer some small amount $e > 0$ and for responders to accept it.[5] Note that in this analysis, there is no coordination problem because the subgame perfect equilibrium is unique.

An alternative experimental method, which is more informative about responders' preferences, is for the responder to precommit to a threshold "minimum acceptable offer" (MAO) which she will accept (and any lower offer will be rejected). When the game is played using the MAO method it is closely related to BOS.[6] Table 4 shows how.

The Table shows a simplified ultimatum game in which players can choose to offer, and state as MAOs, only elements of (2.50,5.00,7.50). If the MAO is less than or equal to the offer x, the payoffs are (10-x,x); otherwise they earn (0,0).

The kinship to BOS is apparent along the diagonal-- the set of pure-strategy Nash equilibria--the payoffs represent every feasible division of $10. Lower offers clearly favor proposers and higher offers favor responders. Hence, within these set of Nash equilibria players face a coordination game identical to the BOS. Both players prefer some agreement to none, but different agreements benefit them differently.

Table 3: Simplified Ultimatum Game

Responder MAO

[5] We assume for simplicity that the smallest offer must be positive.

[6] The simplified ultimatum game in Table 3 is also closely related to the "Nash demand game", in which two players propose shares for themselves of x_1 and x_2, and they earn their shares if and only if $x_1+x_2<10$ (otherwise they earn nothing). The difference between the demand game and the ultimatum game is that if the proposed shares (x, and MAO) add to less than $10 in the ultimatum game the shortfall goes to the proposer, while it is discarded in the Nash demand game. In pure-strategy equilibrium there is no shortfall, so the pure-strategy equilibria of the two games are exactly the same.

		2.50	5.00	7.50
Proposer offer	2.50	7.50, 2.50	0, 0	0, 0
	5.00	5.00, 5.00	5.00, 5.00	0, 0
	7.50	2.50, 7.50	2.50, 7.50	2.50, 7.50

The close relation with the BOS emphasizes that a coordination problem exists in the ultimatum game too. And since we know that timing seems to affect equilibrium selection in BOS and in related games with multiple equilibria (resource dilemmas and step-level public goods provision), it is natural to wonder whether timing affects ultimatum games as well.

There is a more specific reason to wonder about the effect of timing in ultimatum games. Ultimatum games using the specific-offer method seems to elicit more acceptances by responders, holding offers constant, than in games using the MAO method. This difference, to our knowledge, has not been established by a careful comparison but published data collected using the two different methods strongly suggest it (see also Schotter, Weiss, Zapater, 1994). For example, in experiments with specific offers virtually all offers of \$4.00-\$4.50 are accepted. But in experiments which elicit MAOs from responders, it is common for a quarter or more of subjects to demand \$5.00. (In one condition in Blount & Bazerman, 1995, the median MAO was \$5.)

Figure la offers a cursory picture of how the two methods seem to differ. Figure 1 compares cumulative distributions of MAOs from our simultaneous-move data (described in section III below) and rejection frequencies stitched together from several different studies which all used the specific-offer method.[7] Specific-offer rejections are less frequent for high offers-- for example, 20% of \$4.25 (or larger) offers are rejected in the specific offer method, but 60% of those offers are rejected using the MAO method-and the specific offer distribution has less variance.

However, the comparison between the specific-offer and MAO methods confounds a difference in the way responders' choices, are measured, and a difference in timing. In the specific-offer method the proposers always move first, and both players know it. In the MAO method the players (usually) move simultaneously. One hypothesis about the specific-offer vs. MAO difference (if there is one) is that timing makes the difference. In the specific-offer method, responders know they move second and may grant a first-mover advantage to the first-moving proposer by agreeing to accept less (as in the studies

[7] The studies we used were Roth et al (1991), Hoffman et al (1994, random FHSS instructions only), Forsythe et al (1994), Slonim & Roth (1995, 60sk treatment), and Cameron (1995, Teal money 2OOOrp treatment). In all these studies, there are few very low offers so estimating the rejection frequencies for low offers is difficult. Thus, we impose monotonicity on the rejection frequency function-- forcing the frequencies to rise as offers fall-- in a specific way. We begin with rejection frequencies at high offers (say X). As we move down to lower offers (say X-.5), if the rejection frequency falls we pool all the data from X and X-.5 offers, estimating a single rejection frequency for all offers in the interval (X,X-.5). This met-hod (which is similar to "isotonic regression") can be easily spotted because it produces flat spots in the frequency function.

described above). We can test this hypothesis by using only the MAO method, but altering the timing of when proposers make offers.

Figure 1b shows MAOs in proposer-first games (denoted PR) described further below, compared to the same specific offer rejection frequencies shown in Figure 1a. The two distributions are substantially closer. For example, in the specific offer method the offer which is rejected half the time is between $3.00 and $3.50, and the corresponding offer in the MAO method, with proposers moving first, is $4.00. The fact that the distributions in Figure lb are closer together than those in Figure la suggests means that timing may account for much of the difference between specific-offer behavior and MAO method behavior.

III. Ultimatum bargaining experiments

In the ultimatum game experiments, large groups of University of Chicago MBA subjects were recruited (often at the beginning or end of a class, n=284 in total). Subjects were randomly paired with one another and told they were paired with someone else in the same room, but they would not know who that person was. (Instructions are provided in the appendix.) Proposers offered a division of $10 to a responder, in increments of $.50. Responders indicated which offers they would accept by checking "accept" or "reject" from a list of all possible offers. From their acceptances we computed a minimum acceptable offer or MAO. Offers and MAOs were matched by the experimenters and subjects were paid what they earned (typically in an envelope distributed at the end of class, or at the beginning of the subsequent class).

Ultimatum games were played in three conditions. In the simultaneous (SIM) condition, proposers and responders filled out their offer and MAO forms at the same time. In the PR condition, proposers first handed in their offer forms, then responders filled out MAO forms and handed them in. Thus, in the PR condition both players know the proposer moved first, but responders did not know the proposer's offer. An RP condition is the opposite: Responders handed in MAO forms first, then proposers made offers.
The hypothesized effects of timing, predicted by both first-mover advantage and virtual observability theories, are as follows:

Offers H1:	Offer(RP)	>	Offer (SIM)	>	Offer (PR)
MAOs: H1:	MAO (RP)	>	MAO (SIM)	>	MAO (PR)

Figure 2 shows the cumulative distribution function (cdf) of MAOs in the PR, SIM, and RP conditions. Table 5 reports descriptive statistics and two statistics testing whether the three distributions appear to be drawn from the same population. There is a clear ordering in the three means, but the standard deviations are large and only MAOs from the more extreme conditions, PR and RP, are significantly different with a p-value less than .05.

Looking at Figure 2, the difference in conditions is a little more evident: For low offers, around $1-$3, the rejection rates vary from about 60% to 65% to 80% when responders move after, at the same time, or before proposers. Test statistics reported in Table 5 confirm that, using the more powerful Epps-Singleton test, the PR-RP difference is highly significant, the RP-SIM difference is only marginally significant (p=.07) and PR and SIM are hard to distinguish. Hence, we cannot reject the hypothesis of equality of MAOs in favor of the strict ordering predicted by H1, but we can reject equality in favor of the extreme prediction MAO(RP) > MAO(PR).

Table 5: Statistics testing equality of MAO distributions

		Descriptive Statistic		
condition	n	mean	median	Std dev
PR	7	2.82	4.25	2.33
SIM	3	3.39	4.50	2.30
RP	4	3.69	4.50	1.85

statistic		Epps - Singleton CF(p)		Kolmogorov - Smirnov (p)	
condition	n	RP	SIM	RP	SIM
SIM	3	9.05 (.07)		1.35 (.51)	
PR	7	19.14 001)	1.28 (.85)	2.88 (.26)	1.29 (.54)

Note: Test statistics are Epps-Singleton characteristic function (CF) tests and one-tailed large sample Kolmogorov - Smirnov tests. Test statistics are distributed chi-squared with 4 degrees of freedom under the null hypothesis of distributional equality. All p-values are one-tailed

Figures 3 shows the cumulative frequency of offers in all three conditions. The offer distributions are very similar since, as in most studies, offers are tightly clustered around $5. There is a slight tendency for lower offers in the PR condition (and a few more super-generous offers above $5 in the RP condition), but no differences in the three conditions are significant. Hence, we cannot reject the hypothesis of equality of offers in favor of the timing-based alternative H1.

Overall, the ultimatum data show a modest effect of timing.[8] MAOs are indeed lower when responders know they move second, and are higher when they move first, but not by much and only the PR-RP difference is significant. Offers do not change much at all.

[8] Two other studies report evidence of timing effects. Bagai (1992) collected MAOs when subjects were told a division was proposed "earlier this semester" or "will [be] propose[d]". The "earlier" group, corresponding to our PR, had mean MAOs of $2.34 and the "will propose" group, corresponding to RP, had mean of MAOs of $1.69. The difference is in the opposite in direction of ours, and is insignificant.

Blount (in press) also changed timing (though the change in timing was deliberately confounded with a change in subjects' knowledge of the distribution of offers from which an offer would be drawn). Her data suggest a timing effect which is in the same direction as our SIM-PR difference, but larger: Sixty percent of the responders in the PR condition, who received an offer in an envelope stapled to their response sheet, accepted $1 or less. The corresponding figure is only 28% in a separate simultaneous-move study. While this difference in MAOs is dramatic (the two cdf s are much further apart than ours in Figure 2d), the low sample size in the SIM condition limits test power (large sample KS statistic=4.74, p=.09 by a one-tailed test).

IV. Weak-fink coordination game experiments

The main purpose of this paper is to distinguish virtual observablity from the first-mover advantage theory by conducting experiments on coordination games. Table 6 shows payoffs in one such game, the "weak-link" game, first studied by Van Huyck, Battalio & Beil (1991).

In the weak-link game, groups of three subjects choose numbers from 1 to 7. The row player's payoff depends on the number she chose (shown in the rows) and on the smallest number chosen by any player in the group (hence, the term "weak link" game). The payoffs are an increasing function of the smallest number chosen, and a decreasing function of how far the row player is from the smallest number. Since everyone wants to be no higher than the minimum, and wants the minimum to be as large as possible, the game requires coordination. Every number is a Nash equilibrium and the equilibria are Pareto-ranked: coordinating on X provides a higher payoffs for everyone than X-1; choosing 7 is the highest payoff of all.

Table 5: Weak-Link Game

		MINIMUM VALUE OF X CHOSEN						
		7	6	5	4	3	2	1
YOUR CHOICE OF X	7	1.30	1.10	.90	.70	.50	.30	.10
	6		1.20	1.00	.80	.60	.40	.20
	5			1.10	.90	.70	.50	.30
	4				.100	.80	.60	.40
	3					.90	.70	.50
	2						.80	.60
	1							.70

The weak link game models situations in which group production is determined by the level of the lowest-level input. Examples include: keeping a secret, meeting a group at a restaurant which -will not seat-anyone until everyone in the group has arrived, output in "high reliability" organizations in which a single failure or low-quality component causes disaster, or submitting chapters to a book which cannot be printed until all the chapters arrive (e.g., Camerer & Knez, 1997).

Taken together, Blount's study and ours show that a weak first-mover advantage appears to drive MAOs down if responders know an offer has already been made and, more significantly in our data, to raise MAOs if responders know they move first. (Bagai's result is opposite, and puzzling.) The effect is modest in size, however, is only significant in the most extreme comparison between responder-last RP and responder-first PR conditions.

In the weak-link game-, virtual observability implies that if players move in a specified order, the first player will act as if others see his move and best-respond to it (and similarly for the second player). Denote the first and second player's choices by X_1 and X_2. Then the third player will choose $\min(X_1,X_2)$, so the second player will choose $X_2=X_1$ (choosing less would mean creating a lower minimum), and the first player will choose $X_1=7$. Since the third player "follows" the second, and the second follows the first, the first player can "create" a minimum of 7 simply by choosing 7 to begin with. Virtual observability therefore implies an efficient outcome in the weak-link game.

Notice, however, that the outcome in which everyone chooses 7 provides no first-mover advantage, per se. The first mover does prefer the equilibrum of 7s, but this equilibrium provides no special advantage for her, since others prefer it as well. So first-mover advantage has no use as a selection principle in the weak-link game and therefore predicts that timing alone shouldn't matter.

Thus, we can use weak-link games to distinguish the first-mover advantage and virtual observability theories. Virtual observability predicts that coordination will improve, to choices near seven, when we move from simultaneous moves to sequential moves which are unobservable. The first mover advantage theory predicts no difference. We can write this formally, denoting the distribution of number choices by D(.), and representing stochastic dominance relation among distributions by $A>_{sd}B$ (the distribution A stochastically dominates B). Since cumulative distributions with more choices of high numbers will be stochastically dominated by distributions with more low-number choices, we have:

First mover advantage $H2_{fma}$: $\quad D(SIM) = D(SEQ) <_{sd} D(OBS)$

Virtual observability $H2_{vo}$: $\quad D(SIM) <_{sd} D(SEQ) = D(OBS)$

Experiments were conducted with groups of Caltech undergraduates (n=60) recruited from a list of subjects who had participated in previous experiments, and UCLA undergraduates (n=54) recruited from an accounting class.

In each session 6-18 subjects sat in a room together. Subjects were randomly assigned numbers, and letters A, B, or C, and were formed into three-person groups. They did not know which other subjects were in their group. After reading the instructions (see Appendix) out loud to all subjects, the subjects answered two questions about how different choices led to different payoffs. When all subjects had answered correctly, the experiment began.

Groups participated in one of three conditions. In the SIMultaneous condition, all three subjects made their choices at the same time. After each round, forms recording the choices were collected, the experimenters recorded choices by others in the group on the same forms, and the forms were returned to the subjects. To ensure comparability with the other conditions, subjects learned the choices of each of the other two subjects in their group.

In the SEQuential condition subjects moved in a specified order--A first, then B, then C--but later-moving subjects did not know what the earlier choices were.[9] As in the SIM condition, after all subjects made choices their forms were collected, filled out to show the choices of others in each group, and returned to subjects.

In the OBServable condition, subjects knew the moves of all players in their group who moved before them. After A subjects made choices the forms were collected, and A choices were told to B and C subjects. After B subjects made choices their choices were told to C subjects.

The OBS condition provides an empirical benchmark against which the SEQ condition can be judged. The test is simple: First-mover advantage theories predict the SEQ condition will be like SIM, not like OBS. Virtual observability predicts SEQ will be like OBS, not like SIM.

To convey the data compactly, we averaged the choices of all subjects in each round, in each condition, for each subject pool. Figure 4 shows the data from Caltech subjects. In the SIM condition, choices initially averaged just above 6, and drifted slowly upward over eight rounds. By round six, three of the four groups had coordinated on 7.
In the OBS condition, choices converged to 7 sharply and immediately. The SEQ condition results are initially close to those of SIM, but they converge much more quickly to 7 and are much closer to the choices in OBS in later periods. SEQ choices in the first four periods are slow to converge upward (as in SIM) but in the last four periods 93 of 96 SEQ choices were 7s (as in OBS). Overall, the data suggest some predictive accuracy to virtual observability, since choices in SEQ more closely resemble those in OBS with experience.

Figure 5 shows similar data from UCLA undergraduates. Choices are lower and more dispersed than the Caltech choices.[10] The UCLA data also provides support for

---

[9] To help ensure that subjects could not tell when others in the room made their choices (so that subjects could not easily identify who might be in their group), as the A subjects recorded their choices the B and C subjects marked an "X" in a box on their forms. This way, all subjects made a mark on their sheets at the same time. After the A subjects had all marked choices, the B subjects made their choices (and A and C subjects marked Xs), then C subjects made choices.

[10] We felt we lacked control over subjects' incentives in the UCLA subject pool, for several reasons. Some subjects reported in debriefing that they thought some of the choices we fed back to them were faked by us. Since such false feedback is common in psychology experiments, their beliefs are not unfounded, and illustrate the credibility pollution that can harm fellow experimenters when deception is used too freely. Other subjects reported that they deliberately chose lower numbers to "do the best" and make others in their group earn less, at small expense to themselves. We excluded one group from the OBS condition because two of three subjects reported that they were deliberately choosing low numbers to harm fraternity brothers they guessed (correctly) were in their group. Their choices resulted in minima of 3,3,3,1, and 6; including

hypothesis $H2_{vo}$, since the choices in SEQ are much closer to those in OBS, particularly in later periods.

Since the minima in each group are sensitive to outliers in the left tail of the choice distributions, the minima may be a good place to look for distributional differences that are not strikingly apparent by looking at the averaged data in Figures 5. Table 7 shows the minima from all groups in all rounds, for the UCLA data (medians of the minima are shown in parentheses). The period 1 row lists minima of each group in each condition, from low to high. Subsequent rows (representing later rounds) list a group's minima in the same column as in the first row. If you want to track a particular group, read straight down vertically. For example, the rightmost group in the SIM condition that chose 4 in round 1, then chose minima of 5, 6, and 7.

Table 7: Minima from UCLA weak-link groups by rounds (medians in parentheses)

		Condition	
period	SIM	SEQ	OBS
1	13333444 (3)	11134447 (4)	34555667 (5)
2	13445345 (4)	14741557 (4.5)	36667767 (6)
3	13466346 (4)	25764577 (5.5)	43637777 (6.5)
4	13466147 (4)	16773777 (7)	55547777 (6)
5	13477147 (4)	17773677 (7)	77317177 (7)
All	(4)	(5)	(6)

In the UCLA data, the SEQ and OBS conditions reach minima of 7 a total of 40% and 43% of the time, respectively much more frequently than in the SIM condition (only 8%). The medians of the minima in each round are always strictly larger in SEQ than in SIM. Across all rounds, the median of the minima in the SEQ condition, 5, is halfway between the SIM and OBS medians of 4 and 6. With respect to minima, SEQ is squarely between SIM and OBS.

To do more formal hypothesis tests using all the data, we combined data from the earlier and later rounds of each session, excluding the last round.[11] Table 7 shows results from Epps-Singleton (ES) characteristic function tests and Kolmogorov-Smimov (KS) tests. All tests are one-tailed tests in which the alternative hypothesis is that the condition

---

those data would only bolster our conclusion that SEQ data are similar to OBS, by lowering the average OBS choices.

[11] Obviously these data are not truly independent from round to round, but no simple correction for dependence is available for better statistical control. (The extremely conservative method of examining only first-period results has too little power for our sample sizes to detect any differences at all.)

named first in the left column produces lower numbers than the second condition (the prediction of virtual observability).

Table 7: Tests for differences among weak-link treatments

	subject pool:	Caltech		UCLA	
	periods:	1-3	4-7	1-2	3-4
SIM - SEQ	n_1, n_2	72, 72	96, 96	48, 48	48, 48
	ES test (p)	3.09 (.56)	12.23 (.02)	13.20 (.02)	9.49 (.05)
	KS test (p)	.11 (.94)	4.08 (.16)	8.17 (.02)	10.67 (.01)
SIM - OBS	n_1, n_2	72, 36	96, 48	48.48	48, 48
	ES test (p)	12.52 (.02)	12.26 (.02)	13.26 (.01)	7.42 (.12)
	KS test (p)	5.35 (.03)	4.50 (.11)	10.67 (.01)	4.17 (.14)
SEQ - OBS	n_1, n_2	72, 36	96, 48	48, 48	48, 48
	ES test (p)	17.17 (.01)	1.33 (.86)	22.21 (.001)	6.27 (.18)
	KS test (p)	8.17 (.02)	.35 (.86)	2.67 (.27)	.38 (84)

The tests corroborate what is apparent from Figures 5 and 6. In earlier periods the SEQ condition is clearly different from OBS-- contrary to virtual observability-- but in later periods SEQ and OBS cannot be distinguished statistically, while SIM and SEQ can. It appears that virtual observability does not "work" right away, but begins to work as subjects learn.

V. Conclusions

This paper explored two explanations for previous observed timing effects. First-mover advantage is a kind of selection principle in games like BOS and ultimatum-bargaining (which is similar to BOS) that selects one player's preferred outcome by having that player move first. Players moving first will try to get a larger bargaining share; latermovers anticipate this, even when moves are unobservable, and acquiesce. Virtual observability predicts that players act as if they can tell what those who moved earlier did, and therefore makes (almost) the same prediction as subgame perfection when moves are observable. Virtual observability makes the same prediction as a theory based on first-mover advantage in bargaining games, but makes a special prediction in weak-link coordination games (where equilibria are Pareto-ranked).

Therefore, we tested for first-mover advantage and virtual observability in bargaining and weak-link coordination games. In ultimatum bargaining, we used the "strategy method" by requiring responders to indicate a minimum acceptable offer (MAO). Both theories predict that, compared to a simultaneous-move treatment, offers would rise and MAOs would fall when proposers move first, and the opposite patterns would result when proposers move last. In six sessions with a total of 286 subjects, offers did not change much across the three timing conditions but MAOs did change, somewhat, in the direction we predicted. In addition, when responders move first (second) they demand more (less), though the differences between demands in these two conditions and simultaneous-move demands were insignificant.

In weak-link coordination games we tested whether knowing that others had gone first would improve coordination. Generally, the results from the sequential (but unobservable) move games lie between the results from simultaneous-move games and games with observable moves. Moreover, in later periods, choices in the sequential-unobservable case are much closer to the sequential-observable choices than to those in the simultaneous move.

Our results have several implications.

Sequential coordination

The results on weak-link coordination suggest that simultaneity of choices is one important source of coordination failure. Spreading choices out in time, even when previous choices are not observed, can improve coordination (as conjectured by Bryant, 1983). When previous choices <u>are</u> observed, then subgame perfection selects the efficient outcome uniquely, but several levels of iterated rationality are needed (in the three-person game) to achieve the efficient outcome. That is, for the first player A to choose 7 she must be rational, believe that players B and C are, and believe that B believes C is rational. For B to reciprocate and choose 7 as well requires her to be rational and believe C is. For C to reciprocate requires only that she be rational. The frequency of coordination failure in the observable condition casts doubt on the willingness of players to bet on these levels of iterated rationality, as has been observed in many other games (e.g., Ho, Weigelt & Camerer, 1996; Weber, 1999).

Fairness and timing

Previous experimental findings on ultimatum bargaining are often characterized as showing that responders are willing to give up money to punish proposers they think have treated the unfairly. The timing results we report suggest this interpretation is incomplete. If distaste for unfairness drives responders to state positive MAOs, why do their MAOs fall substantially when they know proposers move first? Within the fairness framework, the obvious answer is that a low offer is more fair when proposers move first than when proposers move second. But this answer suggests that fairness means "fair exercise of advantage", and thus cannot be completely decoupled from variables that alter advantage.

On the other hand, the effect of timing in our ultimatum-bargaining experiments is also much smaller than in the previous BOS studies. The BOS results are probably much larger because timing does not compete with fairness as a selection principle in those games since an equal split outcome is not possible. Ultimatums -are more like resource dilemmas in which equal resource-use is an obvious fair point.[12] In those dilemmas the

[12] The equal split outcome is not possible in the BOS.

effect of timing should therefore be muted by the strength of equal-use as a focal principle. Indeed, the effects of timing reported by Rapoport et al, described above, are more like our ultimatum results in magnitude, and substantially weaker than in BOS. This does not mean timing effects can be ignored. Instead, the overall picture from BOS, resource dilemmas, and ultimatum games shows that many structural features of games act as selection principles. Perhaps unsurprisingly, equal-payoff is a strong principle, and has a bigger effect than the subtle effect of timing.

Elicitation methods and timing cues

Our results have some implications for how game theory experiments are conducted and their results interpreted. The first important point is that previous studies have confounded timing and "response mode" by either confronting responders with specific offers or eliciting an MAO. The MAO method has generally been used in simultaneous move games whereas the specific-offer method makes clear that proposers move first. There is a sense from this literature (substantiated in Figure 1a) that MAOs are higher than corresponding rejections of specific offers, which suggests a dynamic inconsistency in which subjects may state an MAO of $5, say (a mode in many samples), but actually accept less than $5 when faced with a specific offer.[13] We think, instead, that MAOs are inflated partly because the order of timing is ambiguous. When we made it clear that responders move second, in our PR treatment, MAOs fell substantially and almost half accepted $.50 or zero.

Once simultaneous-move ultimatum games are seen as coordination games, the possibilities that wording, methods by which roles are assigned, and other variables might affect outcomes become natural since all these features could create different focal points or act as selection principles. For example, Blount and Bazerman (1995) elicit MAOs two different ways: One method asks responders to directly record an MAO, implicitly evaluating their outcomes independently, and the second method asks them to circle which offers they would accept from a list of possible offers, implicitly evaluating outcomes comparatively. They find substantially higher MAOs in the direct method (median $5) than in the list method (median $2.50). A timing-based interpretation of their finding is that the list method contains a proposer-first timing cue which the direct method does not clearly have. Similarly, Boles & Messick (1990) found that when offers were physically presented to subjects- dollar bills were laid in front of responders-- then offers were accepted more frequently. One possible explanation is that physical presentation may again act as a timing cue.

Timing and the psychology of belief formation in games

Our findings should pique the curiosity of game theorists (and psychologists too) about how players actually form beliefs in games. As pointed out at the start, the standard game-theoretic model draws no special distinctions among a player's beliefs about what another player did, is doing, or will do. But the psychology of reasoning suggests several ways in which these thinking processes might differ. A relevant idea is

[13] Bagai (1992) tested for such reversals explicitly, and did not find any from 34 subjects.

that players may be better at reasoning backward, about events known to have already happened, than reasoning forward. Evidence from psychology experiments (e.g., Mitchell, Russo.& Pennington, 1989) shows that description of possible outcomes of previously-occurring events is often richer and more complex than description of later-occurring events: The past is easier to "imagine" than the future. In the same way, B and C subjects in the weak-link games, moving second and third, might be able to imagine that earlier-moving A subjects will chose high numbers more easily than if those A subjects move at the same time as B and C do.

A similar point can be made about other features of games which could affect belief formation, but are conventionally assumed not to. The psychological distinction between chance moves by nature and moves by another person is an example. The convention for modelling imperfect information games is to treat these two sources of uncertainty as equivalent. But players may reason about them differently (e.g., Blount, 1995). Recognizing the distinction, and exploring it both experimentally and formally, can only improve the descriptive accuracy of game theory.

References

Amershi, A., Sadanand, A., and Sadanand, V. (1989)."Manipulated Nash equilibria I: Forward induction and thought process dynamics in extensive form," Discussion Paper 1989-4, University of Minnesota, September 1989.

Blount, S. (1995). "When Social Outcomes Aren't Fair: The Effect of Causal Attributions on Preferences," Organizational Behavior and Human Decision Processes, 62(2), 131-144.

Budescu, David V., Wing Tung Au, and Xiao-Ping Chen. (1997). Effects of protocol of play and social orientation on behavior in sequential resource dilemmas. Organizational Behavior & Human Decision Processes, 69(3), 179-193.

Camerer, C. F. and M. Knez. (1997). "Coordination in organizations: A game-theoretic perspective", in Organization Decision Making, ed. Zur Shapira. Cambridge: Cambridge Series on Judgment and Decision Making.

Cooper, R. (1999). Coordination Games: Complementarities and Macroeconomics. Cambridge: Cambridge University Press.

Cooper, R., DeJong, D., Forsythe, R. and Ross, T. (1993). "Forward Induction in the Battle-of-the-Sexes Games," The American Economic Review, 83, 1303-1316

Cooper, Russell; DeJong, Douglas; Forsythe, Robert; and Ross, Thomas. (1994) "Alternative Institutions for Resolving Coordination Problems: Experimental Evidence," in J. Friedman (Ed.), Problems of Coordination in Economic Activity, Norwell, Mass: Kluwer Academic Publishers, 1994, pp. 129-146.

Ho, Teck-Hua and Keith Weigelt. (1996). "Task complexity, equilibrium selection, and learning : An experimental study," Management Science, 42(5), 659-679.

Knez, M. and Camerer, C (1995). " Social Comparison and Outside Options in 3-Player Ultimatum Game Experiments, forthcoming in Games and Economic Behavior.

Knez, M. and Camerer, C (in press). "Creating expectational assets in the laboratory: Coordination in `weakest-link' games," Strategic Management Journal.

Kreps, David M. (1990). Game Theory and Economic Modelling. Oxford: Oxford University Press.

Miller, Dale T. and S. Gunasegram. (1990). "Temporal order and the mutability of events: Implications for blame assignment," Journal of Personality and Social Psychology, 59, 1111-1119.

Mitzkewitz, M. and Nagel, R. (1993). "Experimental Results on Ultimatum Games with Incomplete Information," International Journal of Game Theory, 22, 171-198.

Morris, Michael W.; Damien Sim; and Vittorio Grotto. (1995). "Time of decision, ethical obligation, and causal illusion," in XXX, pp. 209-239

Muller, R. A. and A. Sadanand. (1998) "Virtual observability in two person games." Unpublished mauscript.

Rapoport, A. (1997). "Order of play in strategically equivalent games in extensive form," International Journal of Game Theory, 26(1), 113-136.

Rapoport, Amnon; Budescu, David; and Ramzi Suleiman. (1993). "Sequential requests from randomly distributed shared resources," Journal of Mathematical Psychology, 37, 241-265.

Roth, A. (1995). "Bargaining Experiments," in Handbook of Experimental Economics, J. Kagel and A. Roth, editors, Princeton University Press.

Roth, A. and Malouf, M. (1979). "Game-Theoretic Models and the Role of Information in Bargaining," Psychological Review, 86, 574-594.

Roth, A. (1987). "Bargaining Phenomena and Bargaining Theory," in Laboratory Experimentation in Economics: Six Points of View, edited by A. Roth, Cambridge: Cambridge University Press, 14-41.

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

Schotter, A., Weigelt, K., and Wilson, C. (1994). "Laboratory investigation of mutiperson rationality and presentation effects," Games and Economic Behavior 6, 445-468.

Van Huyck, J. B., Gillette, A.B., and Battalio, R. B. (1992). "Credible Assignments in Coordination Games," Games and Economic Behavior 4, 606-626.

Van Huyck, J. B., Battalio, R. B. and R. O. Beil. (1991. "Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games," Quarterly Journal of Economics, 106(2), 885-910.

Weber, R. A. (1999). "Uncommon knowledge: Learning and behavior in the 'dirty faces' game." Unpublished manuscript.

Appendix: Instructions

In this experiment you will either be a proposer or a responder. The proposer has to decide how to divide up a ten dollar bill between him or herself and the responder The proposer makes an offer of X dollars, where X is divisible by fifty cents. If the responder accepts then the responder receives $X and the proposer receives $10 - $X. If the responder rejects the offer then both the responder and the proposer receive zero.

You have been randomly assigned the role of proposer or responder. If you are a proposer you have an OFFER sheet, if you are a responder you have an ACCEPTANCE sheet. At the top of your sheet you are given an ID number. For example, if your ID number is P12 then you are proposer number 12, while if your ID number is R12 then you are responder number 12. Finally, proposer P12 will be making an offer to responder R12. Your ID number is strictly confidential. Moreover, at no time during or after the experiment will you know the identity of the person you are paired with.

Treatment RP

There will be two steps in the experiment. First, each responder will record the minimum offer that he or she is willing to accept from the proposer. The number the responder writes down is binding. That is, if the offer the responder receives is greater than or equal to this number, then the offer is accepted. However, if the offer is less than this number then the offer is rejected. Also, the number the responders record should be in increments of fifty cents. Once all the acceptance sheets have been collected, each proposer will then record their offer to the responder on their OFFER sheet, where the offer should be in increments of fifty cents.

Treatment SIM

Each responder will record the minimum offer that he or she is willing to accept from the proposer. The number the responder writes down is binding. That is, if the offer the responder receives is greater than or equal to this number, then the offer is accepted. However, if the offer is less than this number then the offer is rejected. Also, the number the responders record should be in increments of fifty cents. While the responders are making their decisions, each proposer will record their offer to the responder on their OFFER sheet, where the offer should be in increments of fifty cents.

Treatment PR

There will be two steps in the experiment. First, each proposer will record their offer on their OFFER sheet, where the offer should be in increments of fifty cents. Once all the offers have been collected, each responder should record the minimum offer that he or she is willing to accept from the proposer. The number the responder writes down is binding. That is, if the offer the responder receives is greater than or equal to this number, then the offer is accepted. However, if the offer is less than this number then the offer is rejected. Again, the number the responders record should be in increments of fifty cents.

Once all of the record sheets have been collected, your record sheet will then be paired with your partner's as indicated by your identification number. The experimenter will then determine whether or not offers are accepted or rejected and how much each participant receives. Your cash payoff will then be placed in an envelope and at the end of the session you should pick up the envelope corresponding to your identification number.

Figure 1a. MAO(SIM) vs. Specific Offer Rejection Frequencies

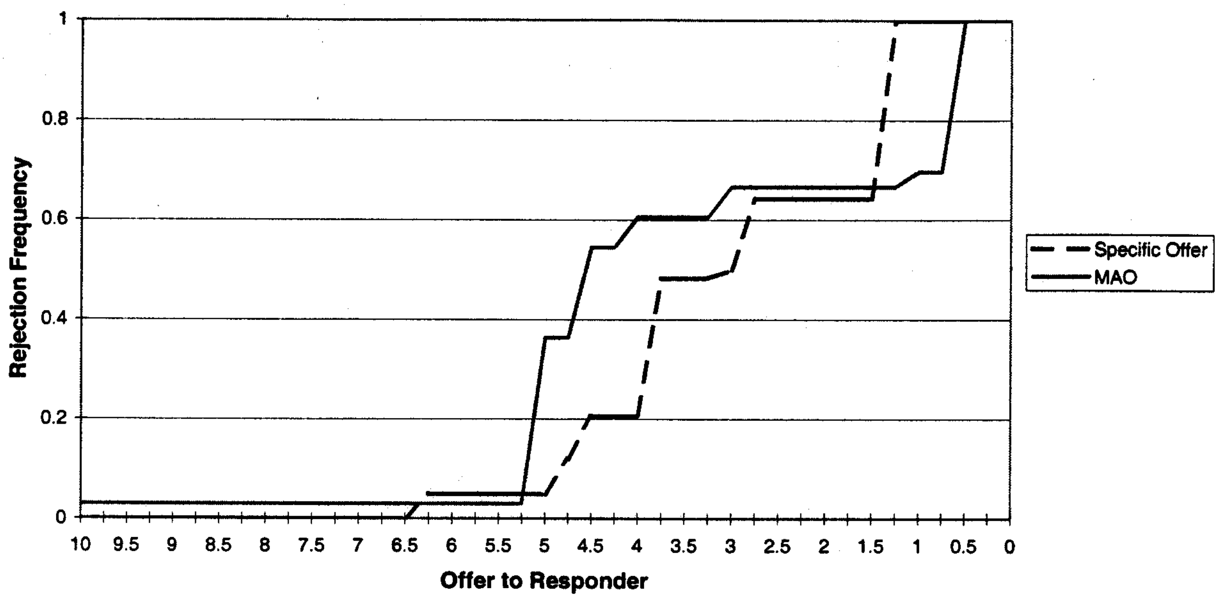

Figure 1b. MAO(PR) vs. Specific Offer Rejection Frequencies

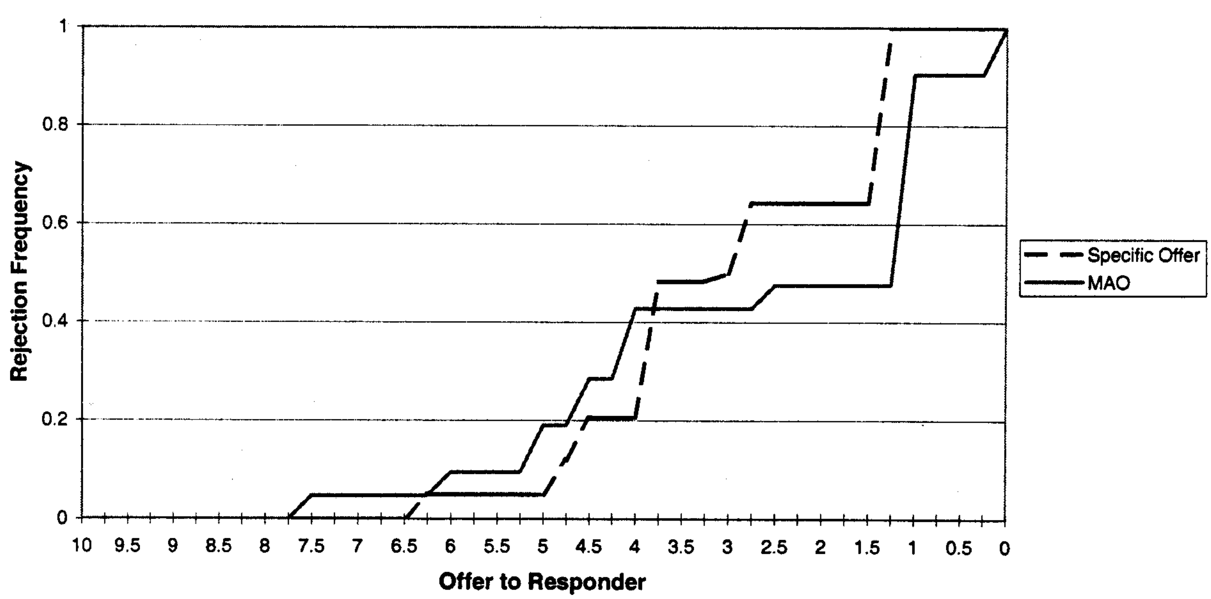

Figure 2. Cumulative MAO Distributions Across Conditions

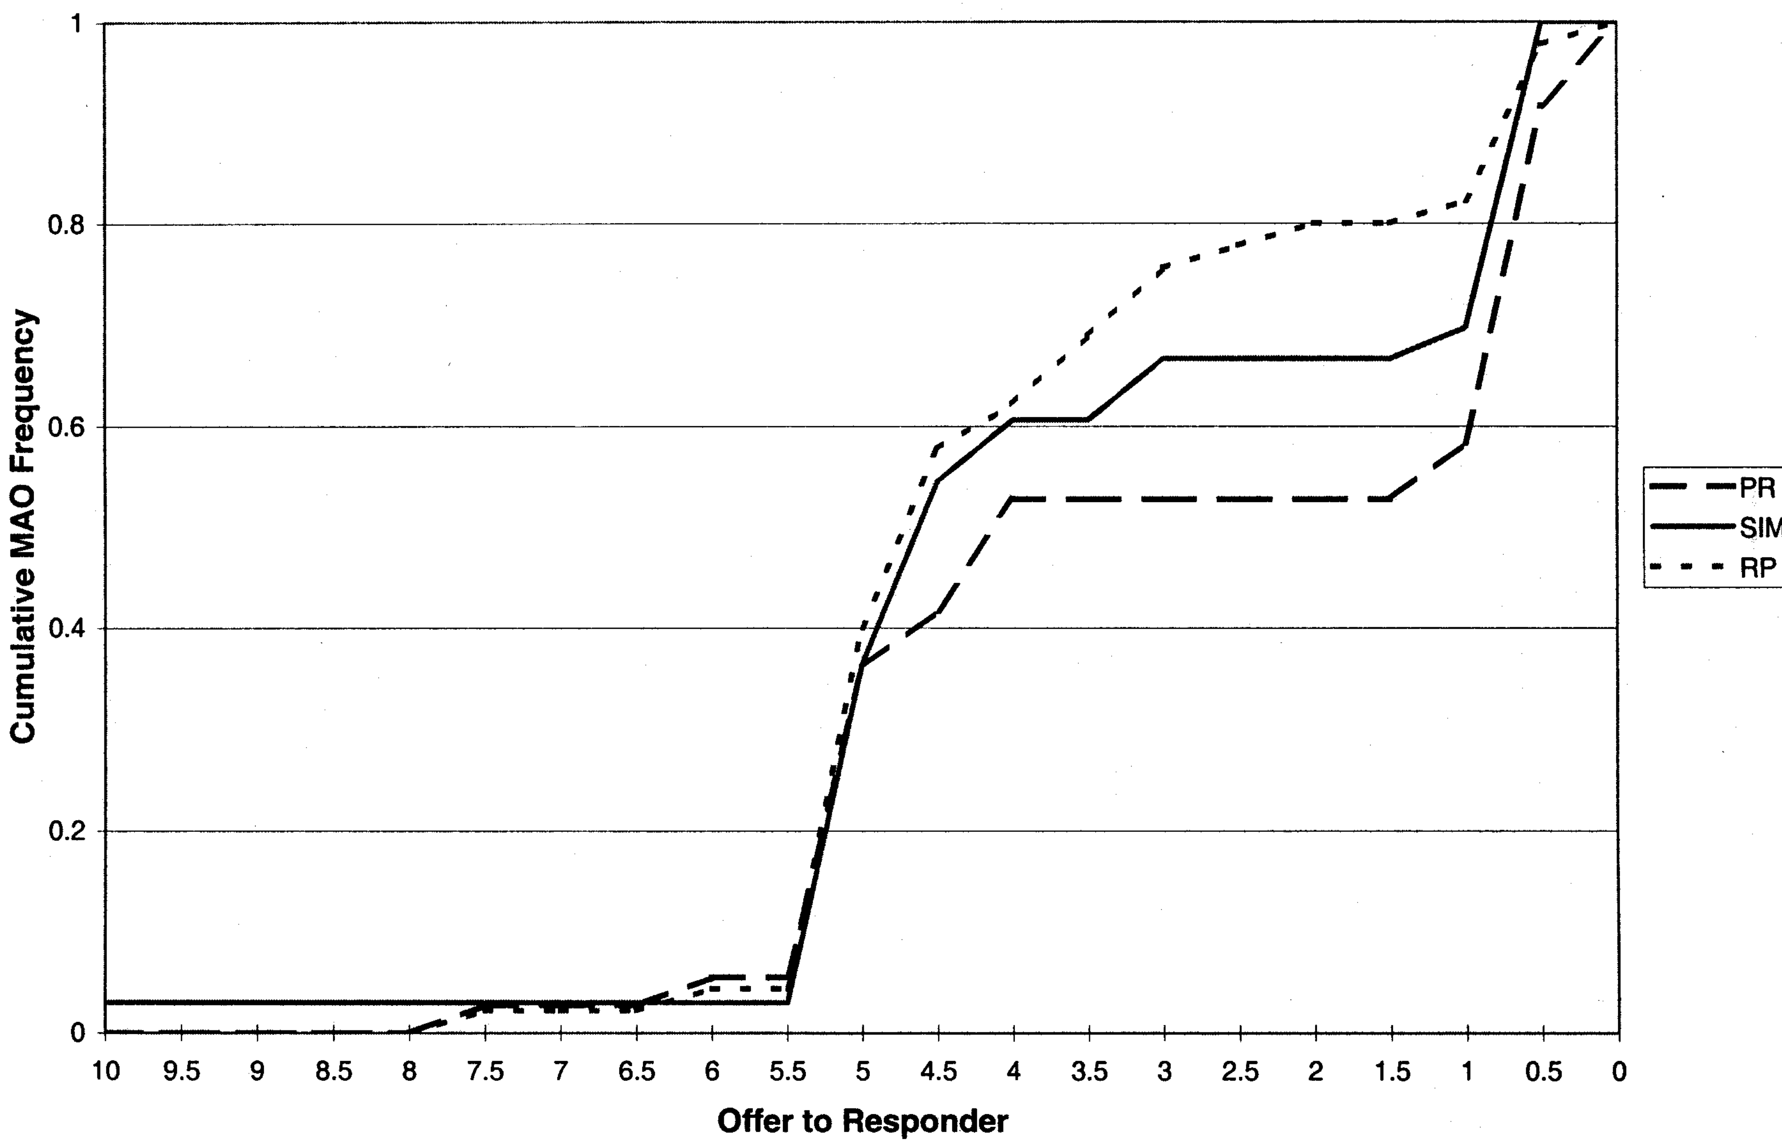

Figure 3. Cumulative Offer Frequencies Across Conditions

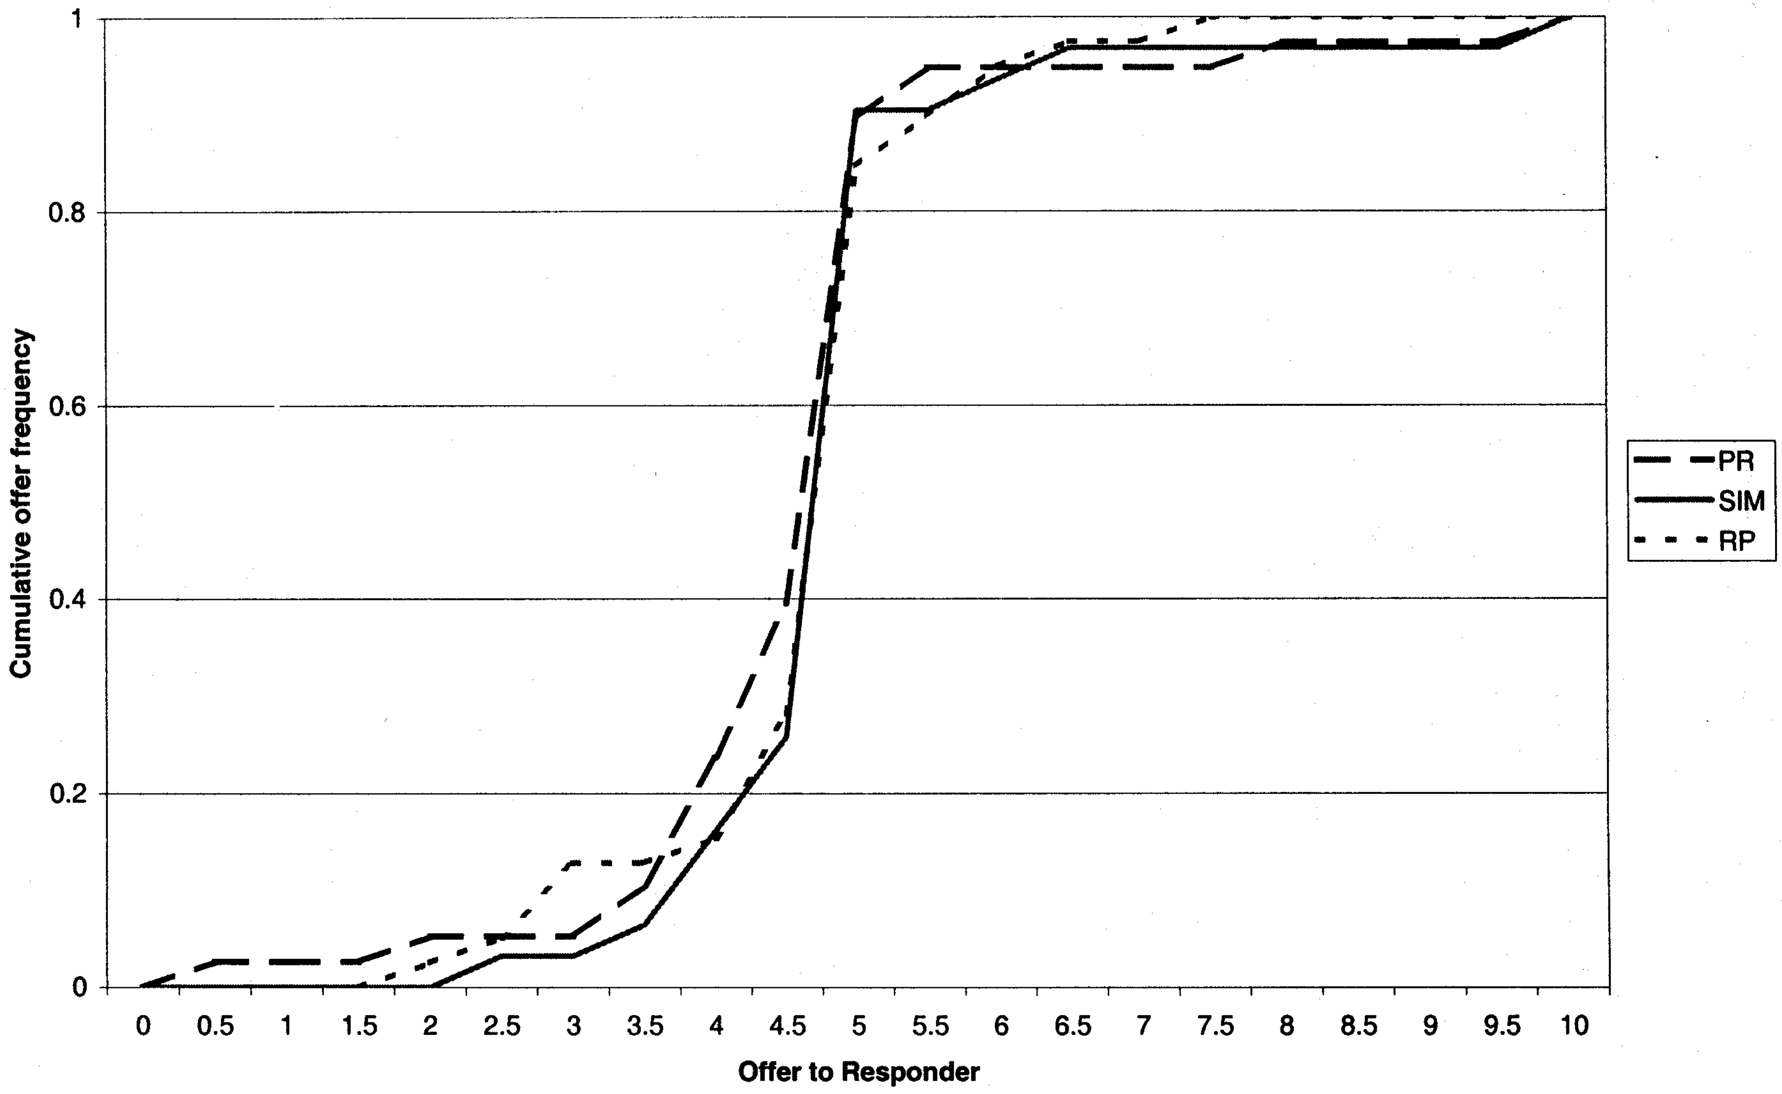

Fig. 4. Average weak-link choices by condition (Caltech)

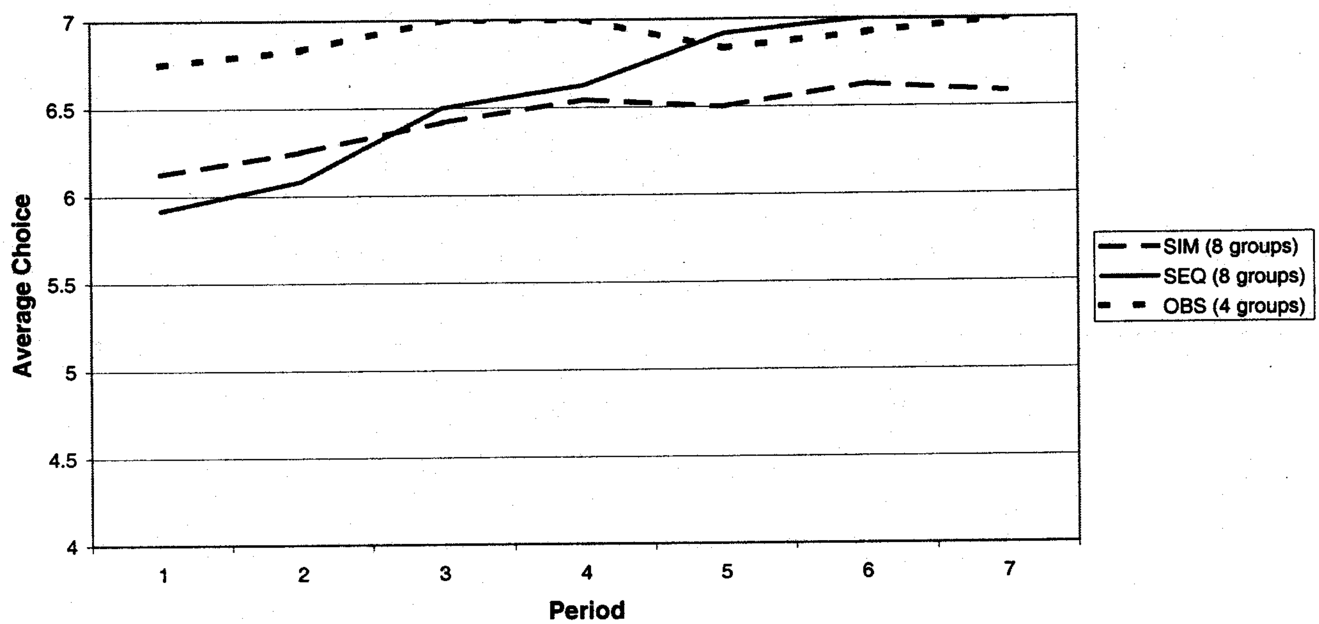

Fig. 5. Average weak-link choices by condition (UCLA)