The Deadline Effect in Bargaining: Some Experimental Evidence

By ALVIN E. ROTH, J. KEITH MURNIGHAN, AND FRANCOISE SCHOUMAKER*

This paper reports the distribution of agreements over time in four bargaining experiments. There is a striking concentration of agreements reached in the very last seconds before the deadline. This "deadline effect" appears to be quite robust, in that the distribution of agreements over time appears to be much less sensitive to the experimental manipulations than is the distribution of the terms of agreement. Each of these experiments also exhibited a substantial frequency of disagreement.

Since last-minute agreements are widely believed to occur frequently in naturally occurring negotiations,[1] it may be helpful to state clearly just what it is that laboratory investigations have to contribute to the study of deadline phenomena. First, while there is a great deal of anecdotal information about the frequency of "eleventh-hour" agreements in naturally occurring negotiations, it has proved difficult to collect reliable data. Second, being able to study deadline phenomena in the laboratory will enable us to distinguish between alternative hypotheses in a way that the study of field data does not permit.[2] Third, while the distribution of agreements over time is one of the clearest phenomena observed in bargaining experiments to date, none of the presently available theoretical models of bargaining is able to account simultaneously for the distribution of agreements over time together with the observed patterns of agreements and substantial observed frequency of disagreements, so these results suggest clear directions for further theoretical work.

The data on the distribution of agreements over time, although analyzed here for the first time, were collected in a series of experiments in which other variables were of primary interest. In each of these experiments, in which bargaining was permitted to proceed for between nine and twelve minutes, bargaining was conducted via terminals in a computer-controlled laboratory that automatically recorded the time of each agreement. We first conducted a formal analysis of this part of the data for the experiment reported in Keith Murnighan, Alvin Roth, and Francoise Schoumaker (1988). Those re-

*Roth: Department of Economics, University of Pittsburgh, Pittsburgh, PA 15260. Murnighan: Department of Business Administration, University of Illinois, 1206 S Sixth Street, Champaign, IL 61820. Schoumaker: Department de Sciences Economique, Universite de Montréal, Montréal, Quebec H3C 3J7. This work has been supported by grants from the National Science Foundation, the Office of Naval Research, and by an Alfred P. Sloan research fellowship. We have received helpful comments from Douglas Gale, Robert Forsythe, Maria Herrero, John Kennan, Roger Myerson, and Shmuel Zamir, as well as from two referees, and assistance in preparing the graphics from John Gallagher.

[1]Some representative quotations:

> Disputes over the terms of collective bargaining agreements are frequently settled only at the last minute. The last-ditch all-night parleys are as familiar to newspaper readers as they are wearing on reporters... The frequency of these photo-finishes suggests that something may be involved fundamental to the process of collective bargaining.
> [John Dunlop and James Healy, 1955, p. 57]

> The question is repeatedly asked as to why negotiations are not settled until a deadline—often at midnight or in the wee hours of the morning—even after a symbolic stopping of the clock. [Dunlop, 1984, p. 16]

> Few of us, even those least involved in bargaining activities, are unfamiliar with the 'eleventh hour' effect widely publicized in mass media accounts of collective bargaining.
> [Jeffrey Rubin and Bert Brown, 1975, p. 120]

[2]For example, labor negotiators often attribute a tendency to reach agreements just before contracts expire to the difficulty of selling any agreement to a diverse constituency if there is still time for continuing negotiations. However the deadline effect observed in our laboratory environment cannot be attributed to this, since each bargainer is bargaining strictly on his own behalf.

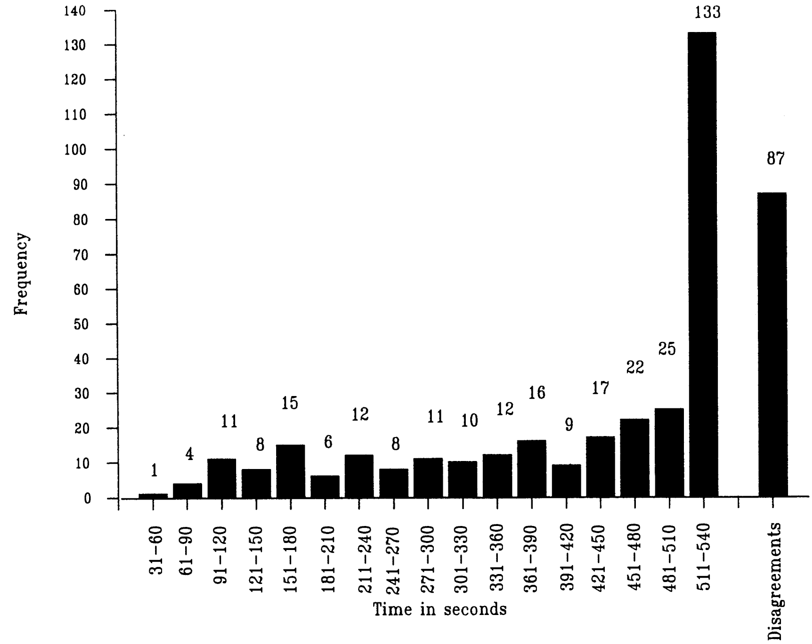

FIGURE 1A. THE FREQUENCY OF AGREEMENTS AND DISAGREEMENTS FROM THE NEW EXPERIMENT

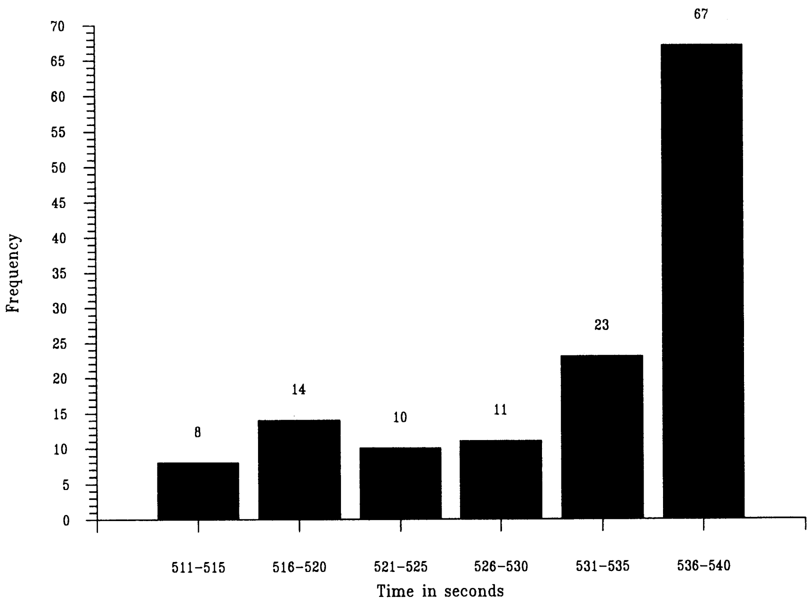

FIGURE 1B. THE FREQUENCY OF AGREEMENTS REACHED IN THE LAST 30 SECONDS OF BARGAINING IN THE NEW EXPERIMENT sults, reported here, prompted a retrospective analysis of two earlier experiments (Roth, Michael Malouf, and Murnighan, 1981; Roth and Murnighan, 1982).[3] We also present here new experimental data which permit us to test specific hypotheses about the relationship between initial bargaining positions, bargaining outcomes, and time of agreement. Although all this data comes from laboratory experiments, there is a sense in which it is not all fully "experimental data," since the experiments were mostly designed for purposes other than to test specific hypotheses about agreement times.

The distribution of agreements over time proved to be one of the most consistent features of the results of these otherwise diverse experiments. A typical distribution, from the new experiment is shown by the histograms in Figure 1. Figure 1A graphs the number of agreements reached in each 30-second interval in the (9 minutes of) bargaining; Figure 1B breaks the last 30 seconds into 5-second intervals. Across all four experiments, 41 percent of the observed agreements were reached in the last 30 seconds of bargaining, and 53 percent of these were reached in the last 5 seconds. (Of agreements reached in the last 5 seconds, 52 percent were reached in the last second; see Table 4.)

I. The Bargaining Environment

The experiments were initially designed to test axiomatic models of bargaining of the kind proposed by John Nash (1950; see Roth, 1979), which are stated in terms of the von Neumann-Morgenstern expected utilities of the bargainers. To experimentally test these theories, an experiment must permit these utilities to be determined.

A class of games that permits this was first employed in the experiment of Roth and Malouf (1979), and is employed in three of the four experiments discussed here. In these *binary lottery games*, each agent i can win only one of two monetary prizes, a large prize λ_i or a small prize σ_i (with $\lambda_i > \sigma_i$). The players bargain over the distribution of "lottery tickets" that determine the probability of receiving the large prize: for example, an agent i who receives 40 percent of the lottery tickets has a 40 percent chance of receiving the amount λ_i and a 60 percent chance of receiving the amount σ_i. Players who do not reach agreement in the allotted time each receive σ_i. Since the information about preferences conveyed by an expected utility function is meaningfully represented only up to the arbitrary choice of origin and scale, there is no loss of generality in normalizing each agent's utility so that $u_i(\lambda_i)=1$ and $u_i(\sigma_i)=0$. The utility of agent i for any agreement is then precisely equal to his probability of receiving the amount λ_i, that is, equal to the percentage of lottery tickets he has received.[4] (In the specific binary lottery games discussed in this paper, σ_i always was equal to \$0.00.)

The following is typical of the procedures in each experiment.

> Each participant was seated at a visually isolated terminal of a computer system. Participants were seated at scattered terminals, and received all of their instructions and conducted all communication via the terminal.

[3] Data were no longer available for time of agreements in the experiment of Roth and Malouf (1979), which is the first study in the series reported here. Another related experiment (Roth and Schoumaker, 1983) is not included: negotiations in that experiment were conducted by simultaneous bids in fixed periods.

[4] We do not need to suppose in our interpretation of these experiments that the use of binary lottery games has controlled for the behavior of the experimental subjects, who may or may not be utility maximizers. Rather, the purpose of using binary lottery games is to control the predictions of the theory, specifically Nash's model of bargaining. The question of whether the various predictions of a theory like Nash's are good descriptors of behavior is independent of whether utility theory is a good description of individual choice. (That is, some of the predictions could be correct even if individuals aren't utility maximizers, and vice versa.) But since Nash's model is stated in terms of the expected utility available to the bargainers, it is necessary to control for what the utility of utility maximizers would be, to know what the predictions of the theory in any particular situation are.

Background information including a brief review of probability theory was presented first. The procedures for sending messages and proposals were then introduced. A proposal was a pair of numbers, the first being the sender's probability of receiving his prize and the second the receiver's probability. In each information condition, the terminal displayed the expected monetary value which the player would receive from any proposal he made or received. The opponent's expected value was only displayed in those conditions in which the player knew his opponent's prize. An agreement was reached whenever one of the bargainers returned a proposal identical to the one he had just received.

Bargainers could send any messages they wished, with one exception. To ensure anonymity, the monitor intercepted any messages that revealed the identity of the players. Each bargainer could make proposals at any time (regardless of what the other bargainer was doing).

The procedures varied slightly from one experiment to the next, as noted below. The time allowed for bargaining varied from nine to twelve minutes, with the time remaining coming on the screen to mark the last three minutes.

II. The Experiments

The following brief descriptions, and Table 1, provide the context needed to discuss the distribution of agreement times. Note that the various experimental manipulations accounted for significant differences in bargaining outcomes (in contrast to the uniformity we will see when we look at distributions of agreement times).

A. *The Initial Experiment of Roth and Malouf (1979)*

The most important variable in this first study manipulated the information available to the two bargainers. In the partial information condition, players knew only their own prize. In the full-information condition, players also knew one another's prize. (Theories of bargaining such as Nash's predict that the outcome of bargaining should be unaffected by the change in the information available to bargainers in the two conditions.) During the 12 minutes of bargaining, bargainers could not reveal information about their prizes in the partial information condition, and could not reveal their identity in either condition. The results were very clear: agreements were almost always equal divisions in the partial information condition, and agreements shifted significantly toward the equal expected value outcome in the full-information condition.

In this study the expected values were automatically computed for subjects to remove the effects of arithmetic ability in determining the relative success of the bargainers. The experiment discussed next was designed in part to verify that these results were not simply an artifact of making this computation available.

B. *The Experiment of Roth, Malouf, and Murnighan (1981)*

This experiment studied four games under three conditions, one in which prizes were stated in terms of an intermediate commodity, "chips." In a low-information condition, players knew only their own prizes; in an intermediate condition, players also knew their opponents' prize as valued in chips, but not money; in the full-information condition, players knew their opponents' prize in chips and in money. In the intermediate-information condition, the expected value of the prizes *in chips* was presented in exactly the same way as the expected value of the prizes in money in the previous experiment. Thus if the previous results were an artifact of providing expected values, the intermediate-information agreements would shift toward giving the player with fewer chips a higher probability of winning his prize.

The results (see Table 1) showed significant variations among outcomes. But the information about chips had no significant effect on the terms of agreement: significant shifts from equal divisions toward equal expected value agreements occurred only after

TABLE 1—A BRIEF SUMMARY OF THE FOUR EXPERIMENTS

Paper	Brief Description of Conditions/Studies	Differences Found Among: Terms of Agreement	Differences Found Among: Disagreement Frequencies	Differences Found Among: Agreement Times
Roth, Malouf, and Murnighan (1981)	3 information conditions 4 games	YES: A main effect for games and a games-by-information interaction	NO	NO
Roth and Murnighan (1982)	2 common knowledge conditions 4 information conditions	YES: Unimodal distributions in two info conditions; bimodal in two others; and stronger differences in common knowledge conditions	YES: One condition with higher frequencies than all others; two combined conditions different from other six	NO
Murnighan, Roth, and Schoumaker (1988)	3 studies, with different disagreement outcomes 2 games Bargainers classified as more or less risk averse	YES: (1) Disagreement outcomes, (2) the risk aversion of the bargainers, and (3) the 3 studies led to significant differences	YES: Higher disagreement outcomes led to more disagreements	YES: Study 1 agreements were faster than Study 2's or Study 3's
The New Experiment	4 games with different focal points	YES: Significant differences among the games	NO	NO

information about the monetary value of the prizes was presented.

C. *The Experiment of Roth and Murnighan (1982)*

This experiment expanded the investigation of different information conditions. The two players in each binary lottery game were assigned either a \$5 or a \$20 prize. As in the previous experiments, there was a condition in which neither bargainer was informed of the other's prize value, and a condition in which both bargainers were informed.

Two additional information conditions were included, one in which only the \$5 player knew both prizes, the other in which only the \$20 player knew both prizes. Finally, half of the conditions made this information common knowledge for both players; the other half left both parties uninformed about the other's information.

When the \$5 player knew both prizes, the terms of agreement were distributed bimodally, with the modes near equal division or equal expected value outcomes. When the \$5 player did not know both prizes, the terms of agreement approximated equal divisions, and were unimodal. More disagreements occurred when the \$5 player knew both prizes and this was not common knowledge (see Table 1).

D. *The Experiment of Murnighan, Roth, and Schoumaker (1988)*

Theoretical work (Roth, 1979; Richard Kihlstrom, Roth, and David Schmeidler, 1981; Roth and Uriel Rothblum, 1982) predicted that risk aversion should be disadvantageous in bargaining, except in situations in which the terms of the predicted agreement had a positive probability of being worse than disagreement. Three experimental studies were conducted to test these predictions. They used ternary rather than binary lottery games, in which player i has three monetary prizes, a_i, b_i, and c_i [$a_i > b_i$ and $a_i > c_i$]. The players bargain over probabilities p_1 and p_2 (with $p_2 = 1 - p_1$) such that player i receives a_i with probability p_i, and receives b_i with probability $1 - p_i$. If the players fail to reach agreement in the allotted time, then players 1 and 2 receive c_1 and c_2, respectively.

Prior to negotiating, the risk aversion of the players was assessed by having them make choices among lotteries.

In the first study, bargainers with different risk aversion were paired in two ternary lottery games, where $a_i = \$10$, $b_i = \$5$, and $c_i = \$2$ (for both bargainers) in one game, and $a_i = \$10$, $b_i = \$2$, and $c_i = \$5$ in the other. The players played each of these games twice with each of two opponents, although the fact that their two opponents were the same was disguised by the procedure. The second study changed the value of a_i, b_i, and c_i to \$10, \$5, and \$4. In the third study, one player had a high prize of \$8, and the other a high prize of \$16, with low and disagreement prizes shifting between \$4 and \$3.

The results supported the predictions: the more risk-averse bargainers did better in the high-disagreement prize game than in the low-disagreement prize game. The frequency of disagreement varied across studies and games; it was highest when the disagreement prize was high. The terms of agreement also varied significantly, with the third study (with unequal prizes for the two bargainers) exhibiting the familiar shift toward equal expected value agreements.

E. *The New Experiment*

This experiment[5] investigated different prize values in binary lottery games. The procedures were identical to those in Roth and Murnighan (1982). Each member of a bargaining pair was assigned a low prize or a high prize: prize pairs in the four conditions were (\$10 and \$15), (\$6 and \$14), (\$5 and \$20), and (\$4 and \$36). Both bargainers were informed about the value of both prizes, and this was common knowledge. So the low-prize players could argue for equal expected value divisions of lottery tickets of 60–40, 70–30, 80–20, and 90–10, respectively. Based on previous observations, we predicted that both the mean payoff to the low-prize player and the variance of the payoff distribution would increase as the prize differences increased. We also conjectured that the concentration of agreements near the deadline would be more pronounced in conditions with higher variance of agreement terms.

Unlike the Roth and Murnighan (1982) results, the data were not bimodal; distributions were approximately normal. The low-prize player's percentage of lottery tickets increased significantly as his prize decreased $[F(3,322) = 3.76,\ p < .011]$, with the \$10–\$15 condition (55–45) different from the \$4–\$36 condition (61–39). The \$6–\$14 and \$5–\$20 conditions were not significantly different from each other or from the other conditions. The variance of agreement terms also differed: Bartlett's $F = 24.48$, $p < .01$, highest in the \$4–\$36 condition (std. dev. = 15.7) and lowest in the \$10–\$15 condition (std. dev. = 6.36). Disagreement frequencies were almost identical from one condition to the next: of 96 bargaining interactions in each condition, disagreements were 22, 23, 21, and 21 (23 percent overall). Thus this new data showed significant, predicted effects for prize values on the terms of agreement, but not on disagreement frequencies.

In summary, these four experiments have investigated a diverse set of experimental conditions, and both the terms of agreement and the frequency of disagreement varied significantly across these conditions (see Table 1). As we will now see, the timing of these agreements changed very little across the different conditions or experiments.

III. The Distribution of Agreements Over Time

Figure 1 presented the data from our most recent experiment. Figures 2, 3, and 4 present the data from the three earlier experiments in similar fashion. All of the figures use arbitrary time blocks of 30 and 5 seconds. Table 2 provides an additional breakdown of the data with different (but still arbitrary) time cutoffs. All four studies show a substantial number of late agreements, even as late as the last second. Although the choice of final intervals is arbitrary, there is a clear concentration of agreements near the deadline regardless of the intervals chosen to measure "nearness" to the deadline. This is evident from the cumulative distributions of agreements over time (Figures 5–8), which are all nearly vertical near the deadline.

[5]We hope to give a full account of other aspects of the results of this experiment in Roth, Murnighan, and Schoumaker (in preparation).

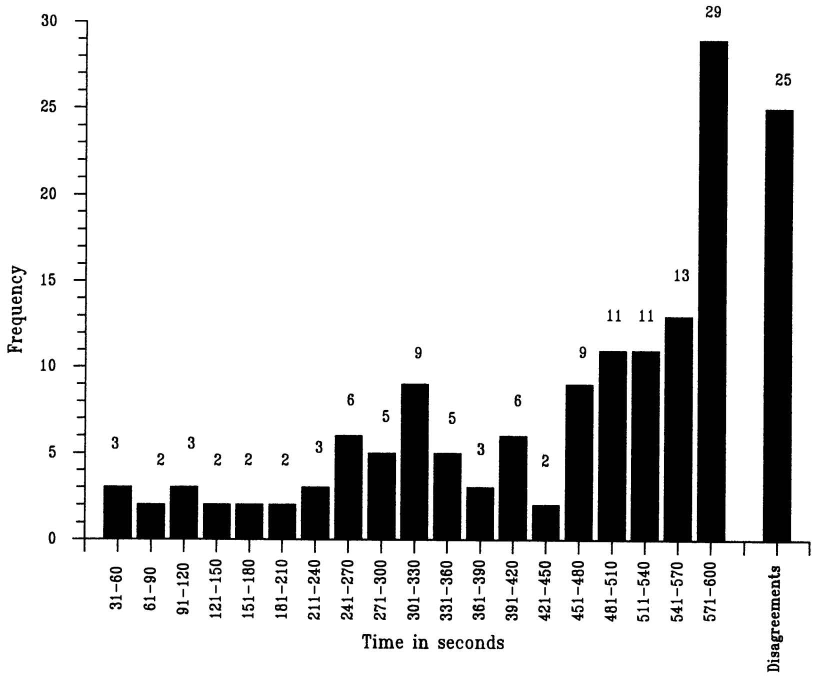

FIGURE 2A. THE FREQUENCY OF AGREEMENTS AND DISAGREEMENTS FROM ROTH, MALOUF, AND MURNIGHAN (1981)

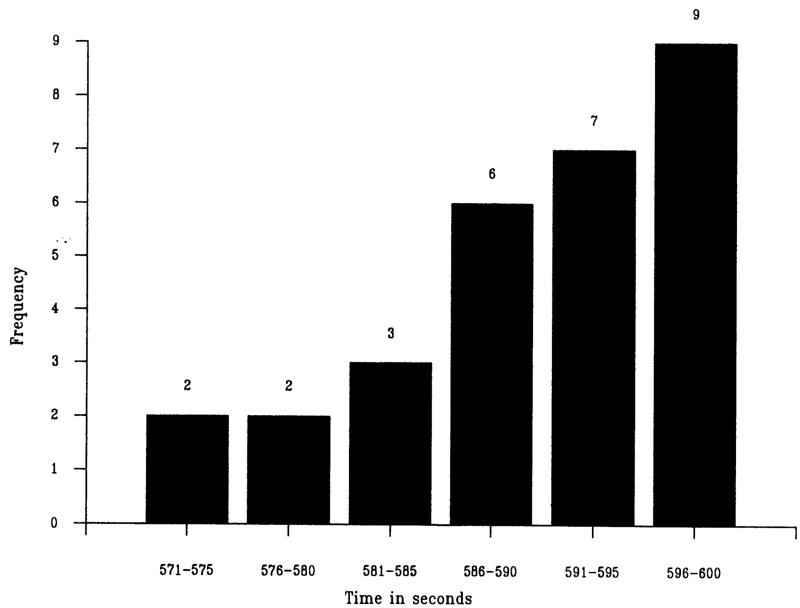

FIGURE 2B. THE FREQUENCY OF AGREEMENTS REACHED IN THE LAST 30 SECONDS OF BARGAINING IN ROTH, MALOUF, AND MURNIGHAN (1981)

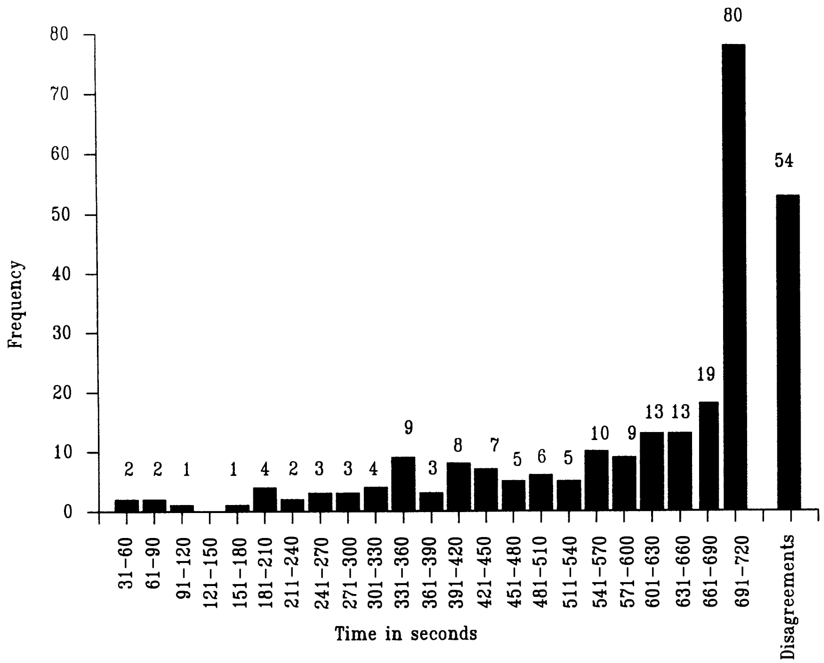

FIGURE 3A. THE FREQUENCY OF AGREEMENTS AND DISAGREEMENTS FROM ROTH AND MURNIGHAN (1982)

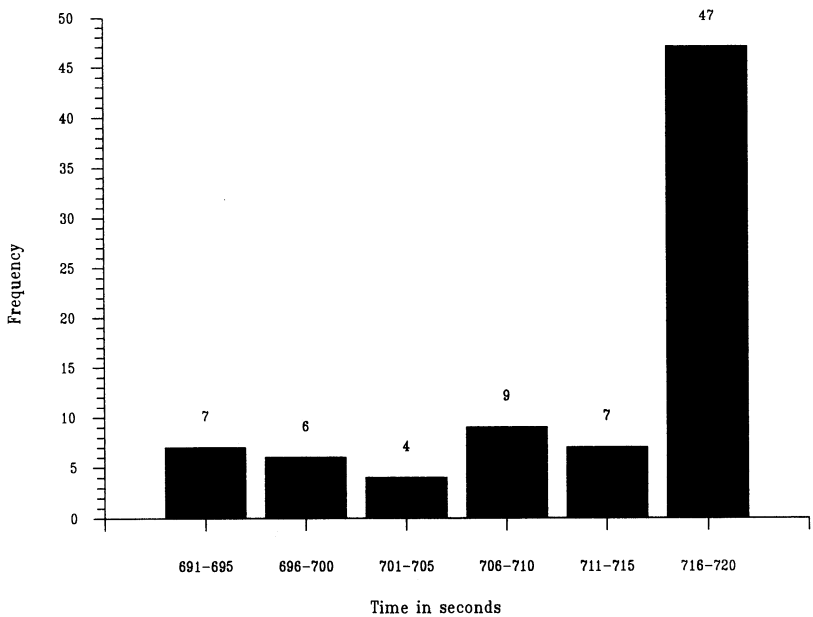

FIGURE 3B. THE FREQUENCY OF AGREEMENTS REACHED IN THE LAST 30 SECONDS OF BARGAINING IN ROTH AND MURNIGHAN (1982)

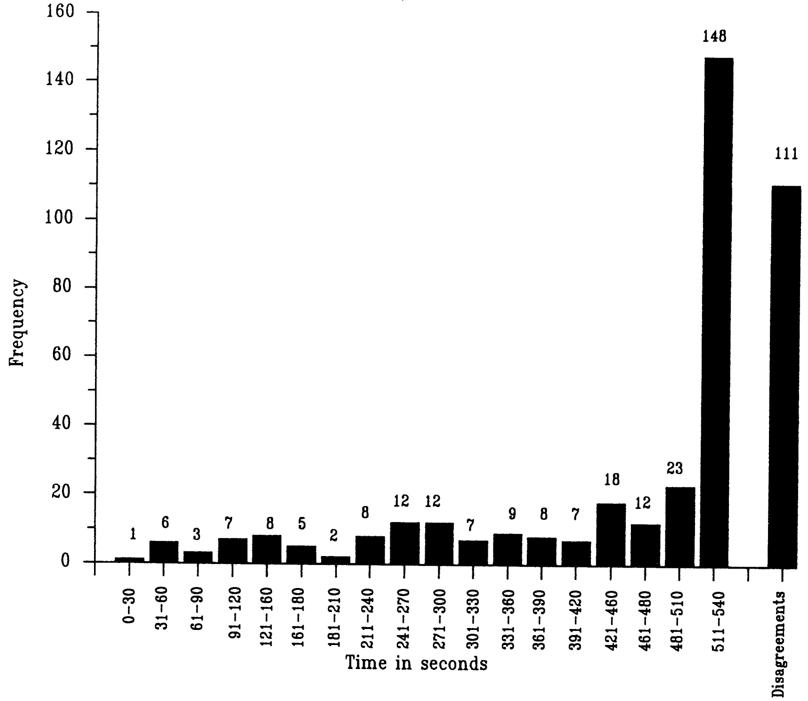

FIGURE 4A. THE FREQUENCY OF AGREEMENTS AND DISAGREEMENTS FROM MURNIGHAN, ROTH, AND SCHOUMAKER (1988)

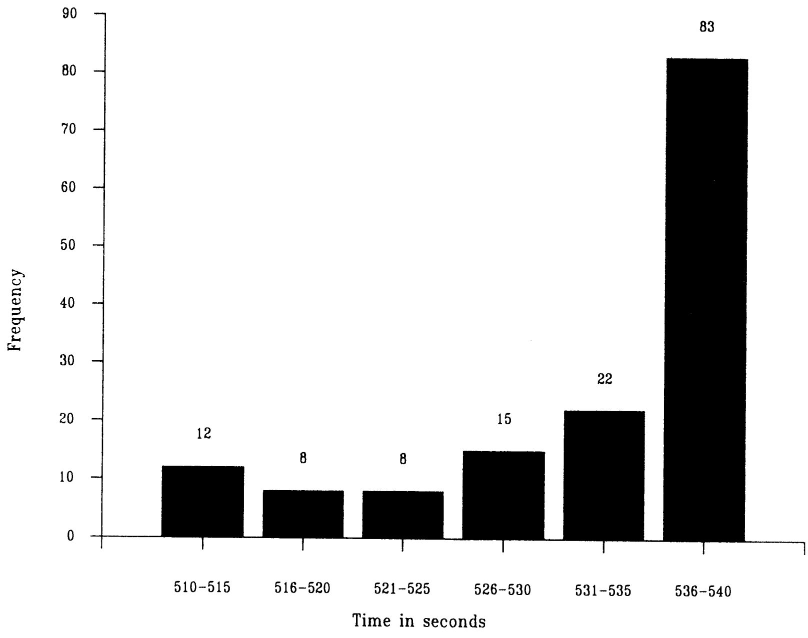

FIGURE 4B. THE FREQUENCY OF AGREEMENTS REACHED IN THE LAST 30 SECONDS OF BARGAINING IN MURNIGHAN, ROTH, AND SCHOUMAKER (1988)

TABLE 2—THE FREQUENCY OF AGREEMENT TIMES IN FOUR EXPERIMENTS

Experiment/Condition	Total Time (seconds)	Agreements	Disagreements	Agreements in Last ___ seconds 60	45	30	15	10	5	1
Roth, Malouf, and Murnighan (1981)	600	126	25(17) (in percent)	42	37	29(19) (in percent)	23	16	9	2
Roth and Murnighan (1982)	720	209	54(21)	99	90	80(30)	63	56	46	37
Common Knowledge		93	20(18)	45	42	39(35)	30	27	18	14
Neither Knows		23	4(15)	11	11	10(37)	7	7	4	4
$20 Knows		24	6(20)	12	11	11(36)	10	8	4	4
$ 5 Knows		21	5(19)	10	10	8(32)	7	7	6	5
Both Know		25	5(17)	12	10	10(33)	6	5	4	1
Not Common Knowledge		116	34(23)	54	48	41(27)	33	29	28	23
Neither Knows		33	3(8)	16	15	12(33)	12	11	11	8
$20 Knows		20	4(17)	6	4	4(16)	4	2	1	1
$ 5 Knows		37	18(33)	23	20	17(31)	12	11	11	10
Both Know		26	9(26)	9	9	8(23)	5	5	5	4
Murnighan, Roth, and Schoumaker (1988)	540	296	111(27)	170	160	148(36)	121	108	83	30
Low-Disagreement Payoff		177	27(13)	99	94	86(42)	65	55	41	14
High-Disagreement Payoff		119	84(41)	71	66	62(31)	56	53	42	16
The New Experiment (1987)	540	325	91(22)	159	149	133(32)	103	91	67	37
10–15 Game		82	22(21)	45	44	41(39)	28	26	18	9
6–14 Game		81	23(22)	36	35	28(27)	20	18	12	7
5–20 Game		80	24(23)	40	35	33(32)	28	22	18	11
4–36 Game		82	22(21)	38	35	31(30)	27	25	19	10
Totals		956	281(23)	462	428	390(32)	310	271	205	106

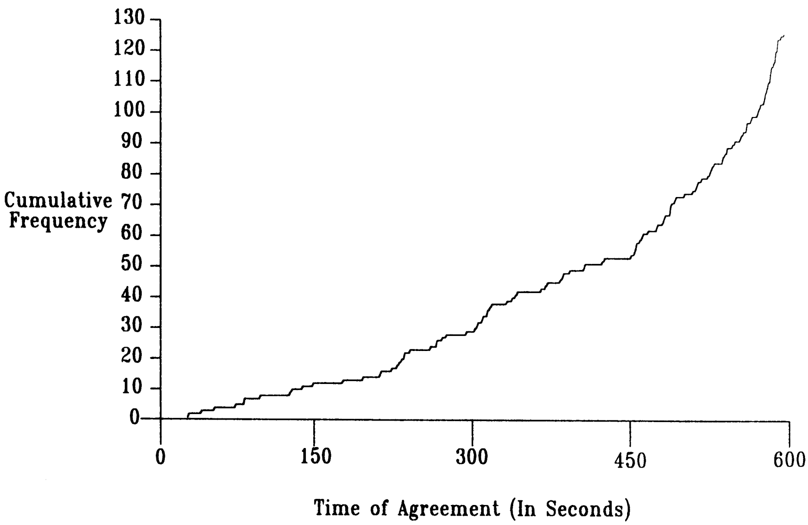

FIGURE 5. CUMULATIVE FREQUENCIES FROM ROTH, MALOUF, AND MURNIGHAN (1981)

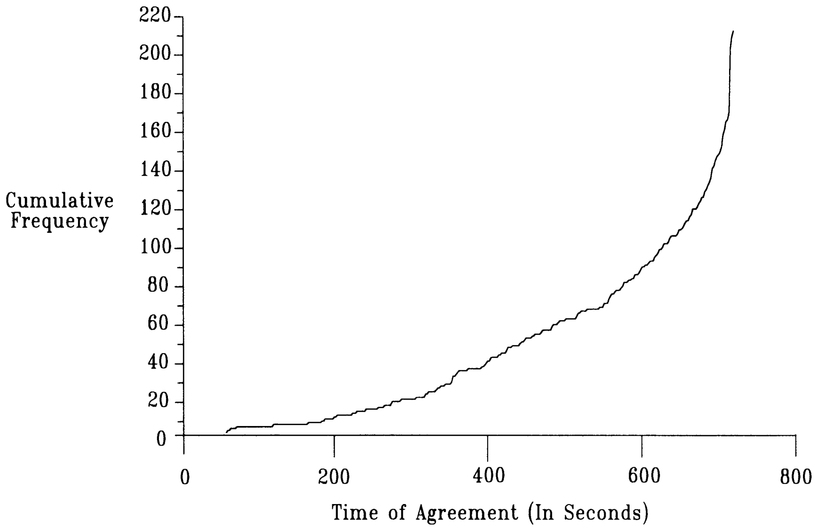

FIGURE 6. CUMULATIVE FREQUENCIES FROM ROTH AND MURNIGHAN (1982)

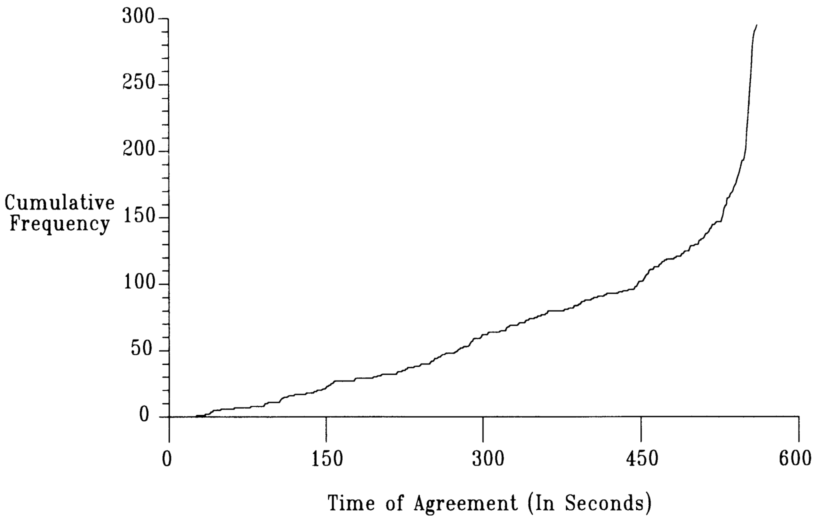

FIGURE 7. CUMULATIVE FREQUENCIES FROM MURNIGHAN, ROTH, AND SCHOUMAKER (1988)

Agreement times were compared across conditions within each experiment. Unlike the tests of agreement terms, both nonparametric and parametric tests showed no significant differences among conditions in Roth, Malouf, and Murnighan (1981), Roth and Murnighan (1982), or the new experiment. In the most stringent test, in the new experiment, neither the timing of agreements nor the variance of agreement times responded to the differences in the bargainers' prize values and the associated changes in the distribution of agreement terms.

There was a significant difference in agreement times in comparisons of the three studies in Murnighan, Roth, and Schoumaker

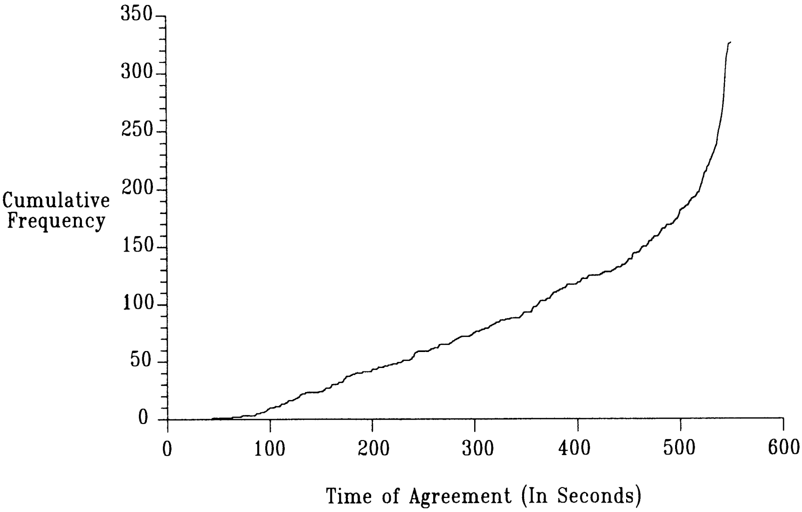

FIGURE 8. CUMULATIVE FREQUENCIES FROM THE NEW EXPERIMENT

(1988): $F(2,289) = 8.32$, $p < .01$. Study 1 agreements ($\bar{X} = 352$ seconds) were made more quickly than study 2 ($\bar{X} = 463$) or study 3's ($\bar{X} = 426$). Almost all the agreements in study 1 of this experiment were reached in the low-disagreement prize condition. In those games the players had equal prizes, and there was a preponderance of equal division agreements. The games in the other two studies were more complicated.

Analyses were also conducted within the conditions of each experiment to assess the relationship between agreement outcomes and agreement time.[6] Table 3 summarizes the correlational analyses for each of the studies. Although a few positive correlations were significant, suggesting that better outcomes were obtained by the low prize player when agreements were later, the most interesting results are in the Roth and Murnighan (1982) experiment, when the $5 player knew both prizes and the $20 player knew only his own. Here a strong negative correlation resulted when the information was common knowledge, and a strong positive correlation resulted when the information was not common knowledge. These correlations disappear when they are calculated for only the agreements reached in the last 30 seconds ($r(3) = -.06$ and $r(4) = -.22$, ns). Nevertheless, these findings provide the basis for speculation about the different strategies that may have been utilized in these two conditions. Overall, however, the correlations suggest little in the way of a consistent relationship between agreement outcomes and agreement times.

Table 4 provides information about the frequency of agreements in the last 30 seconds of negotiating, and how many of those agreements occurred in the last 5 or 1 second, and how many agreements in the last five seconds went to the very last second. There is clearly a concentration of agreements right before the deadline, most dramatically in Roth and Murnighan (1982), where almost half of the agreements reached in the last 30 seconds were reached in the last second.[7] These numbers are reflected in

[6]Analyses across conditions of the experiments also investigated the relationship between average agreement times and disagreement frequencies: the correlation was not significant.

[7]Recall that in the last three minutes of bargaining, the amount of time remaining appeared on the screen. Because of the cycle time of the software controlling the

TABLE 3—CORRELATIONS BETWEEN AGREEMENT TIME AND AGREEMENT OUTCOME FOR THE LOWER-PRIZE PLAYER (FREQUENCIES IN PARENTHESES) FOR ALL FOUR EXPERIMENTS

Roth, Malouf, and Murnighan (1981)	Game 1	2	3	4
Full Info	.33(12)	.07(12)	.46(12)	.11(10)
Partial Info	−.19(14)	.16(13)	−.21(10)	.33(14)
No Info	−.08(7)	−.49(8)	no variance (all 50–50's)	.83*(7)
Roth and Murnighan (1982)	Common Knowledge		Not Common Knowledge	
Neither Knows	−.15(23)		.20(33)	
$20 Knows	−.06(24)		−.16(20)	
$ 5 Knows	−.61**(21)		.49**(37)	
Both Know	−.59**(25)		−.15(26)	
Murnighan, Roth, and Schoumaker (1987)	Disagreement Prize: Low		High	
Study 1	−.07(32)[a]		−.28(7)[a]	
Study 2	−.06(59)[a]		−.04(32)[a]	
Study 3	.26*(84)[b]		.21(81)[b]	
The New Experiment	Game 10–15	6–14	5–20	4–36
	.20(82)	.03(81)	.27*(80)	−.08(82)

*$p<.05$
**$p<.01$

[a]In these conditions, both players had the same prize. The correlations were computed using the outcomes for the less risk averse of the two negotiators.
[b]In these two conditions, correlations were computed for the low-prize player. Comparable correlations for the less risk-averse players were not significant: $r(84) = -.11$ and $r(81) = -.15$, ns.

the steepness of the cumulative distributions of agreements in the final seconds of bargaining (Figures 5–8).

In summary, in each of these experiments the data exhibit a very clear concentration of agreements in the final seconds before the deadline. The distribution of times at which agreements are reached appears to be much more constant among conditions within each experiment and among experiments than are the terms of agreement (see Table 1).

IV. Possible Theoretical Models

We turn now to consider what kinds of theoretical bargaining models offer explanations for a concentration of agreements at the deadline. If we were looking for an explanation of this phenomenon in isolation, straightforward modifications of models in the literature would seem quite promising. However, the situation is more complicated when we look for an explanation of the deadline effect that can also account for the patterns of agreements and disagreements in these experiments.

experiments (which were conducted using a mainframe computer under time-sharing), any agreement may have been concluded a little before or after the recorded time. Similarly, the last recorded agreement was as late as 724 seconds in one "12-minute" bargaining encounter, but appears in this data as an agreement at 720 seconds.

TABLE 4—THE PERCENTAGE OF AGREEMENTS REACHED IN THE LAST 30 SECONDS AND, OF THOSE, THE PERCENTAGE REACHED IN THE LAST 5 OR 1 SECOND

	Percent Agreements in Last 30 Seconds	Percent of Last 30 Agreements in Last 5	Percent Last 30 Second Agreements in Last 1	Percent Last 5 Second Agreements in Last 1
Roth, Malouf, and Murnighan (1981)	23	31	7	22
Roth and Murnighan (1982)	38	59	46	80
Common Knowledge	42	46	36	78
Not Common Knowledge	35	66	56	85
Murnighan, Roth, and Schoumaker (1987)	50	56	20	36
Low-Disagreement Payoff	49	47	16	34
High-Disagreement Payoff	52	68	26	38
Roth, Murnighan, and Schoumaker (1987)	41	50	28	55
10–15 Game	59	43	22	50
6–14 Game	35	43	25	58
5–20 Game	41	55	33	61
4–36 Game	38	61	32	53
Totals	41	53	27	52

The family of axiomatic models of bargaining do not explicitly consider time, and do not say anything directly about the deadline effect. To consider what kinds of strategic models might be appropriate, recall the following features of the experimental environment.

To reach an agreement, a bargainer needed to return a proposal identical to the last proposal received. Since entering and sending a proposal requires a short series of key strokes, some time is needed to accept an offer, and still more time is needed to make a new proposal and have it accepted. Consequently, in the last seconds before the deadline there may be some doubt about whether enough time remains for a given transaction. Furthermore, because the computer programs that controlled these experiments had a cycle time of approximately 2 seconds, bargainers could not be completely certain how many seconds remained. Thus the exact timing of the (effective) deadline is a little uncertain. The bargainers bargained for a while with no risk that the bargaining would be exogenously terminated, but as the deadline became quite close, there was an increasing probability that the bargaining would end in the next instant, before additional proposals could be made.[8]

A. *Complete Information Models of Sequential Bargaining*

A growing body of literature concerning strategic models of bargaining explicitly models the passage of time. (Much of the recent work takes its inspiration from the infinite horizon model of Rubinstein, 1982; see Roth, 1985, for a collection of representative work.) There are two key elements to this approach. First, bargaining is modeled as being costly, in that negotiation takes time, consumption does not occur until agreement takes place, and future payoffs are discounted relative to present payoffs.

[8]Even apart from the precise details of sending and accepting proposals and keeping time in these experiments, it seems reasonable to postulate that in a variety of negotiating environments with a fixed deadline, as the deadline approaches, the probability rises that essentially exogenous factors will prevent an agreement from being reached.

Specifically, time is divided into periods in which bargainers alternately have the option to make offers. If an offer is accepted, bargaining ends and the bargainers receive their agreed payoffs. If an offer is declined, bargaining continues to the next period, and the value of potential agreement shrinks according to some discount factor. Disagreement results if no offer is accepted at any period. The second element is that the predicted outcome is a subgame-perfect equilibrium. Thus neither bargainer can commit himself to a contingent course of action that he would not want to carry out if the contingency should arise.

These two assumptions allow a unique subgame-perfect equilibrium to be computed. Intuitively, the shrinking value of the pie, together with the requirement of subgame perfection, allows the solution to be determined by "working backward" (compare Ingolf Stahl, 1972). The shrinking pie gives the bargainers an incentive to reach an early agreement.

Although most of the discussions in the literature speak of this shrinkage as being due to discounting, when players are expected utility maximizers it is equivalent to think of this "shrinkage" as arising from a probability that each period will be the last, so that future payoffs are discounted by the probability that the bargaining will continue to that point. That is, if a player turns down an offer in any period, there is some probability that the opportunity to continue bargaining will disappear before a counteroffer can be made, and in this case both bargainers will receive nothing (compare Binmore, Rubinstein, and Asher Wollinsky, 1986).

So the bargaining environment of our experiments can be thought of as consisting of a very large number of very short periods, each having zero probability of being the last, until we reach some period near the deadline. After that, each period has some probability of being the last, and this probability increases as the deadline approaches, until the last possible period is reached. (A finite horizon model of just this kind is considered in Joseph Harrington, 1986. See also Oliver Hart, 1986.)

Of course, in our experiments bargainers are not constrained to alternate in making offers. But the remarks we have made about these simple models, while sensitive to the rules about how offers can be made, extend, for example, to models in which each bargainer has an equal chance of being able to make the next offer at each period (see Binmore, 1987). When periods are very short, this may not be such a bad approximation to our experimental environment. Nevertheless, to develop a specific model along these lines, it would be necessary to specify exactly the rules of proposals in a way that precludes simultaneous proposals by the two bargainers,[9] but is faithful to the rules of these experiments. This could involve some subtleties which we will not explore here, since our purpose in this section is to discuss certain qualitative features of equilibrium in a model of this kind.

At a subgame-perfect equilibrium of a model of this sort, the *terms* of agreement, but not the time of agreement, are uniquely defined. (The time of agreement would be uniquely defined if the first period of bargaining had a positive probability of being the last, but we are now speaking of models in which there is no cost to delaying agreement until near the end.) At equilibrium, agreement can be reached at any period up to and including the first period that has a positive probability of being the last (i.e., at equilibrium no agreements are reached after the end of this first probabilistic period). Finally there are no disagreements at equilibrium.

While such a model does not mandate a concentration of agreements near the very end, it is consistent with it, and also consistent with some agreements being reached well before the deadline, which is what we observe. Nevertheless the predictions of such a model do not fit the data we observe. On the one hand, we observe a substantial percentage of disagreements. On the other, it is hard to argue that the observed agreements

[9] In order to preserve the uniqueness of the subgame-perfect equilibrium.

correspond to a unique equilibrium agreement, particularly in those experiments that yield a bimodal distribution of agreements.

We can find models that predict a positive probability of disagreement at equilibrium, by looking at models of incomplete information.

B. *Incomplete Information Models of Sequential Bargaining*

In the last few years great strides have been made in exploring models of bargaining under incomplete information.[10] Whereas the complete information literature suggests that all agreements will be efficient (i.e., no disagreements will occur when joint profits are possible), various forms of inefficiency emerge as equilibrium behavior in incomplete information bargaining models. In single-period models, this inefficiency takes the form of a positive probability of disagreement at equilibrium (see, for example, Roger Myerson and Mark Satterthwaite, 1983, or Kalyan Chatterjee, 1985). In multiperiod models in which time is discounted, it takes the form of delays in the time at which agreement is reached (see, for example, Cramton, 1985; or Chatterjee, 1985). In multiperiod models in which there is some probability that each period will be the last, it also takes the form of delays, which now also imply a positive probability of disagreement, since there is a positive probability that the bargaining will terminate before agreement is reached.

To get some intuition about these models, consider a model of sequential bargaining in which (at least) one of the bargainers has some private information about something germane to the outcome of bargaining.[11] Suppose for the moment that the information is binary, and means that the informed bargainer is either in a "strong" or "weak" position, such that if the information were to become common knowledge, the informed bargainer would obtain a more favorable agreement if he is in the strong position. Even though only one bargainer is informed about whether his (own) position is strong or weak, there may still be equilibria at which a bargainer in a strong position will obtain a more favorable agreement than one in a weak position. These equilibria, which are called "separating equilibria" (since they separate strong from weak bargainers), must involve a cost to a bargainer in a strong position that he would be unwilling to pay if he were in a weak position. (Otherwise a bargainer in a weak position would simply do whatever the equilibrium would have called for him to do had he been strong.) In multiperiod models in which delay is costly, the cost to the strong bargainer of demanding the more favorable agreements is that these are reached later than agreements reached by weak bargainers.

When the costliness of delay arises from a positive probability that each period will be the last, this means that at a separating equilibrium strong bargainers face a higher percentage of disagreements than weak bargainers. Not only the existence of disagreements and delays in these models, but also the tradeoffs predicted between the frequency of disagreement and the terms of agreement, are highly evocative of the results of the experiments discussed above.[12] Nevertheless there are aspects of the data that appear difficult to reconcile with such a model. The substantial fraction of the agreements observed well before the deadline (when the

[10] See, for example, Chatterjee, 1985; Peter Cramton, 1985; Drew Fudenberg and Jean Tirole, 1983; Fudenberg, David Levine, and Tirole, 1985; Myerson and Satterthwaite, 1983; Myerson, 1985; Rubinstein 1985a,b; Joel Sobel and Ichiro Takahashi, 1983.

[11] In the literature this is often taken to be information about the size of the potential profits to be divided, but this is not the only kind of information that could enter the model in this way. For our present purposes, it might be useful to think of the private information as concerning the bargainer's own subjective probability distribution over the time at which a final proposal can be made.

[12] In the noncommon knowledge conditions of Roth and Murnighan (1982), the tradeoff between frequency of disagreement and the terms of agreement is virtually identical to what would be predicted at a strategic equilibrium: see the analyses presented in that paper or in Roth (1987).

probability of sudden termination is still zero) may be difficult to explain with this kind of model, particularly since the terms of early agreements do not seem to be distributed differently from late agreements.[13]

V. Further Avenues of Investigation

Further experimental and theoretical work is needed to better identify the roles that deadlines play in determining the outcome of negotiations. The kinds of theoretical models discussed above suggest some directions for experiments intended to serve this purpose.[14]

Another line of investigation concerns the robustness of these results not merely to the kinds of experimental manipulations reported here, (and not only under the particular mechanics of bargaining shared by these experiments), but also to different kinds of "deadlines." In these experiments, the deadline was very sharp: after the deadline, no further bargaining could take place. In many natural bargaining situations in which eleventh-hour agreements are thought to be common (see fn. 1), the deadline is an event like the expiration of a contract. In such cases bargaining may still continue after the deadline, but at greater cost, or subject to greater uncertainties. Of course, under laboratory conditions the nature of the deadline can be freely varied, and so this too is a natural target of further experimental investigation.

In conclusion, one of the clearest phenomena to emerge from the experiments discussed here is a "deadline effect": a very high percentage of agreements are reached very close to the end, just before the deadline. This phenomenon is robust, in the sense that the distribution of agreements over time does not respond to changes in the bargaining environment nearly as much as do other features of the outcome of bargaining, such as the terms on which agreement is reached. Because deadlines are widely believed to often have similar effects in natural bargaining situations, understanding this phenomenon is likely to have practical implications about the design and conduct of negotiations, and may also shed further light on the causes of bargaining inefficiencies, which appear with substantial frequency in all the data considered here, in the form of costly disagreements. Presently available theoretical models do not appear to be able to adequately account simultaneously for the observed patterns of agreement terms, agreement times, and frequency of disagreement, although separately each of these can be "explained." It seems likely to us that many of these features of the data are related, and can neither be observed nor adequately understood in isolation. Consequently, the task of further understanding the effects of deadlines presents a challenge both to theorists and experimenters.

[13] The idea is that, if bargainers in strong positions generally get better agreements by holding out longer, then late agreements should look different from early agreements. However, the absence of systematic correlations between time of agreement and terms of agreement in our data is not to surprising, in view of the fact that the experimental design makes all the observable features of the bargaining common knowledge between the bargainers in all the experiments discussed here except that of Murnighan, Roth, and Schoumaker (1987). (And in that experiment, there is no significant correlation between the time of agreement and the terms achieved by the less risk-averse bargainer; see Table 3.) Maybe if we could observe players' priors about exactly how much time was left we would see some correlation, with players who were more relaxed about the deadline doing better in last-minute agreements. This is, of course, an idea that can be investigated with an appropriately designed experiment.

[14] Of course quite different theoretical models from the ones discussed here may also suggest directions to further explore deadline phenomena. Shmuel Zamir has pointed out to us that both delays and disagreements can occur in games of timing (with complete information), in which there is some irreducible uncertainty about when a given decision to accept a standing offer will be executed.

REFERENCES

Binmore, Ken "Perfect Equilibria in Bargaining Models," *The Economics of Bargaining*, K. Binmore and P. Dasgupta, eds., Oxford: Basil Blackwell, 1987, 77–105.

______, **Rubinstein, Ariel and Wollinsky, Asher,** "The Nash Bargaining Solution in Economic Modelling," *Rand Journal of Economics*, Summer 1986, *17*, 176–88.

Chatterjee, Kalyan, "Disagreement in Bargaining: Models with Incomplete Information," *Game-Theoretic Models of Bargaining*, A. E. Roth, ed., Cambridge: Cambridge University Press, 1985, 9–26.

Cramton, Peter C., "Sequential Bargaining Mechanisms," *Game-Theoretic Models of Bargaining*, A. E. Roth, ed., Cambridge: Cambridge University Press, 1985, 149–79.

Dunlop, John T., *Dispute Resolution: Negotiation and Consensus Building*, Dover, MA: Auburn House, 1984.

______ **and Healy, James J.,** *Collective Bargaining: Principles and Cases*, Homewood, IL: Richard D. Irwin, 1955.

Fudenberg, Drew and Tirole, Jean, "Sequential Bargaining with Incomplete Information," *Review of Economic Studies*, July 1983, *50*, 221–47.

______, **Levine, David and** ______, "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information," *Game-Theoretic Models of Bargaining*, A. E. Roth, ed., Cambridge: Cambridge University Press, 1985, 73–98.

Harrington, Joseph E., Jr., "A Non-Cooperative Bargaining Game with Risk Averse Players and an Uncertain Finite Horizon," *Economics Letters*, 1986, *20*, 9–13.

Hart, Oliver, "Bargaining and Strikes," Working Paper, MIT 1986.

Kihlstrom, Richard, Roth, Alvin E. and Schmeidler, David, "Risk Aversion and Solutions to Nash's Bargaining Problem," *Game Theory and Mathematical Economics*, O. Moeschlin and D. Pallaschke, eds., Amsterdam: North-Holland, 1981, 65–71.

Murnighan, J. Keith, ______ **and Schoumaker, Francoise,** "Risk Aversion in Bargaining: An Experimental Study," *Journal of Risk and Uncertainty*, 1988, *1*, 101–24.

Myerson, Roger B., "Analysis of Two Bargaining Problems with Incomplete Information," *Game-Theoretic Models of Bargaining*, A. E. Roth, ed., Cambridge: Cambridge University Press, 1985, 115–47.

______ **and Satterthwaite, Mark A.,** "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, April 1983, *29*, 265–81.

Nash, John, "The Bargaining Problem," *Econometrica*, April 1950, *18*, 155–62.

Roth, Alvin E., *Axiomatic Models of Bargaining*, Berlin: Springer, 1979.

______, **ed.,** *Game-Theoretic Models of Bargaining*, Cambridge: Cambridge University Press, 1985.

______, "Bargaining Phenomena and Bargaining Theory," *Laboratory Experimentation in Bargaining: Six Points of View*, A. E. Roth, ed., Cambridge: Cambridge University Press, 1987, 14–41.

______ **and Malouf, Michael W. K.,** "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Review*, 1979, *86*, 574–94.

______, ______ **and Murnighan, J. Keith,** "Sociological Versus Strategic Factors in Bargaining," *Journal of Economic Behavior and Organization*, June 1981, *2*, 153–77.

______ **and Murnighan, J. Keith,** "The Role of Information in Bargaining: An Experimental Study," *Econometrica*, September 1982, *50*, 1123–42.

______ **and Schoumaker, Francoise,** "Bargaining and Coordination: Some Experiments," in preparation.

______ **and Rothblum, Uriel G.,** "Risk Aversion and Nash's Solution for Bargaining Games with Risky Outcomes, *Econometrica*, May 1982, *50*, 639–47.

______ **and Schoumaker, Francoise,** "Expectations and Reputations in Bargaining: An Experimental Study," *American Economic Review*, June 1983, *73*, 362–72.

Rubin, Jeffrey Z. and Brown, Bert R., *The Social Psychology of Bargaining and Negotiation*, New York: Academic Press, 1975.

Rubinstein, Ariel, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, May 1982, *50*, 97–109.

______, (1985a) "Choice of Conjectures in a Bargaining Game with Incomplete Information," *Game-Theoretic Models of Bargaining*, A. E. Roth, ed., Cambridge: Cambridge University Press, 1985, 99–114.

______, (1985b) "A Bargaining Model with Incomplete Information about Time Preferences," *Econometrica*, September 1985, *53*, 1151–72.

Sobel, Joel and Takahashi, Ichiro, "A Multi-Stage Model of Bargaining," *Review of Economic Studies*, July 1983, *50*, 411–26.

Stahl, Ingolf, *Bargaining Theory*, E.F.I., Stockholm, 1972.