

RES Easter School: Behavioural Economics Brasenose College Oxford, 22-25 March 2015

Overview of Behavioral Game Theory

Vincent P. Crawford

University of Oxford, All Souls College,
and University of California, San Diego



UC San Diego



European Research Council

Established by the European Commission

Behavioral game theory

Behavioral game theory combines theory and empirical (mainly experimental) evidence to better understand strategic behavior.

Because it involves individual decisions, behavioral game theory in principle includes all topics in behavioral decision theory; but in practice it has developed differently from “behavioral decision theory in games”.

Behavioral game theory

Behavioral game theory combines theory and empirical (mainly experimental) evidence to better understand strategic behavior.

Because it involves individual decisions, behavioral game theory in principle includes all topics in behavioral decision theory; but in practice it has developed differently from “behavioral decision theory in games”.

It has focused mainly on an issue that is unique to games: how people form beliefs about others’ decisions, mostly taking decisions as rational.

(It also considers social preferences, as influenced by reciprocity.)

The implicit hope is for an eventual synthesis of behavioral decision theory and game theory, with the former “plug and play” in the latter.

Learning and strategic thinking

Models of how people form beliefs about others' decisions fall naturally into two categories:

- Learning: When people are rational (in the decision-theoretic sense) and have enough prior experience with analogous games, adaptive learning (whereby others are expected to respond as others have to analogous games in the past) has a strong tendency to converge to a Nash equilibrium, in which people accurately predict others' decisions.

Learning and strategic thinking

Models of how people form beliefs about others' decisions fall naturally into two categories:

- Learning: When people are rational (in the decision-theoretic sense) and have enough prior experience with analogous games, adaptive learning (whereby others are expected to respond as others have to analogous games in the past) has a strong tendency to converge to a Nash equilibrium, in which people accurately predict others' decisions.
- Strategic thinking: In theory, even when people lack enough prior experience, they can reason their way to Nash equilibrium beliefs in their initial responses to a game.

Although models in both categories are often used to justify assuming Nash equilibrium, the categories are relevant more generally.

How learning and thinking interact to determine outcomes

First consider Van Huyck, Cook, and Battalio's (1997 *JEBO*) "continental divide" experiment (named in Camerer, *Behavioral Game Theory*, ch. 1).

Seven subjects chose simultaneously and anonymously among efforts 1 to 14, with payoffs determined by own efforts and the group median effort, as publicly announced via a table (next slide).

The group median was then publicly announced, subjects chose new efforts, and the process continued for 15 rounds.

There were 10 sessions, each with its own, separate subject group.

How learning and thinking interact to determine outcomes

First consider Van Huyck, Cook, and Battalio's (1997 *JEBO*) "continental divide" experiment (named in Camerer, *Behavioral Game Theory*, ch. 1).

Seven subjects chose simultaneously and anonymously among efforts 1 to 14, with payoffs determined by own efforts and the group median effort, as publicly announced via a table (next slide).

The group median was then publicly announced, subjects chose new efforts, and the process continued for 15 rounds.

There were 10 sessions, each with its own, separate subject group.

The best response to a median ≤ 7 was less than the median, and the best response to a median ≥ 8 was greater than the median (the "divide").

There were two symmetric pure-strategy equilibria "all-3" and "all-12".

(There were also mixed-strategy equilibria, behaviorally irrelevant here.)

Your Choice

Payoff Table G
Median Choice

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	45	49	52	55	56	55	46	-59	-88	-105	-117	-127	-135	-142
2	48	53	58	62	65	66	61	-27	-52	-67	-77	-86	-92	-98
3	48	54	60	66	70	74	72	1	-20	-32	-41	-48	-53	-58
4	43	51	58	65	71	77	80	26	8	-2	-9	-14	-19	-22
5	35	44	52	60	69	77	83	46	32	25	19	15	12	10
6	23	33	42	52	62	72	82	62	53	47	43	41	39	38
7	7	18	28	40	51	64	78	75	69	66	64	63	62	62
8	-13	-1	11	23	37	51	69	83	81	80	80	80	81	82
9	-37	-24	-11	3	18	35	57	88	89	91	92	94	96	98
10	-65	-51	-37	-21	-4	15	40	89	94	98	101	104	107	110
11	-97	-82	-66	-49	-31	-9	20	85	94	100	105	110	114	119
12	-133	-117	-100	-82	-61	-37	-5	78	91	99	106	112	118	123
13	-173	-156	-137	-118	-96	-69	-33	67	83	94	103	110	117	123
14	-217	-198	-179	-158	-134	-105	-65	52	72	85	95	104	112	120

Half of the 10 groups happened to have an initial median ≥ 8 , and half had an initial median ≤ 7 .

- The median ≥ 8 groups converged close to the all-12 equilibrium.
- The median ≤ 7 groups converged close to the all-3 equilibrium.




Fig. 3. Median choice in sessions 1 to 10 by period

Because the best response to a median ≤ 7 was less than the median and the best response to a median ≥ 8 was greater than the median, these dynamics and limiting outcomes are predicted by any plausible adaptive learning model in which people myopically adjust in the direction of higher payoffs, given the most recent value of the group median.

As a result, equilibrium selection in the limit is highly history-dependent.

Because the best response to a median ≤ 7 was less than the median and the best response to a median ≥ 8 was greater than the median, these dynamics and limiting outcomes are predicted by any plausible adaptive learning model in which people myopically adjust in the direction of higher payoffs, given the most recent value of the group median.

As a result, equilibrium selection in the limit is highly history-dependent.

Use the myopic best responses described above to define a “basin of attraction” for each pure-strategy equilibrium—the states of the system that lead the dynamics to converge to that equilibrium.

Equilibrium selection can then be described as completely determined by which equilibrium’s basin of attraction subjects’ initial responses fell into.

Thus, learning and strategic thinking interact to determine outcomes:

Even if we care only about limiting outcomes, and learning is sure to converge to an equilibrium, predictions depend on the prior probability distribution of the median initial response, and may be inherently uncertain.

(In some applications people's initial responses may also be important for their own sakes and convergence to equilibrium may not be assured.)

Thus, learning and strategic thinking interact to determine outcomes:

Even if we care only about limiting outcomes, and learning is sure to converge to an equilibrium, predictions depend on the prior probability distribution of the median initial response, and may be inherently uncertain.

(In some applications people's initial responses may also be important for their own sakes and convergence to equilibrium may not be assured.)

This motivates our separate study of strategic thinking later on.

Plainly, with multiple equilibria, strategic thinking must go beyond the logic of equilibrium (e.g. Harsanyi and Selten 1988).

And experiments that evoke initial responses show that subjects often deviate from equilibrium systematically, in modelable ways, especially when it requires fixed-point or indefinitely iterated dominance reasoning.

History-dependent equilibrium selection and the structure of adaptive learning

The dynamics in Van Huyck, Cook, and Battalio's experiment are very informative about equilibrium selection via history-dependent learning.

But because the dynamics and limiting outcomes are as predicted by almost any plausible adaptive learning model, Van Huyck et al.'s results are not very informative about the structure of adaptive learning.

History-dependent equilibrium selection and the structure of adaptive learning

The dynamics in Van Huyck, Cook, and Battalio's experiment are very informative about equilibrium selection via history-dependent learning.

But because the dynamics and limiting outcomes are as predicted by almost any plausible adaptive learning model, Van Huyck et al.'s results are not very informative about the structure of adaptive learning.

Van Huyck, Battalio, and Beil's (1990 *AER*, 1991 *QJE*, 1993 *GEB*) experiments, in which subjects repeatedly played symmetric coordination games, in treatments that varied how payoffs were determined, interaction pattern, and group size, are more informative about learning.

See also:

Crawford, “Learning Dynamics, Lock-in, and Equilibrium Selection in Experimental Coordination Games,” in Pagano and Nicita, editors, *The Evolution of Economic Diversity*, 2001; manuscript at <http://econweb.ucsd.edu/%7Evcrawfor/9719.pdf>)

Crawford, “Adaptive Dynamics in Coordination Games,” 1995 *Econometrica*; <http://econweb.ucsd.edu/%7Evcrawfor/Crawford95EMT.pdf>

Crawford and Broseta, “What Price Coordination? The Efficiency-enhancing Effect of Auctioning the Right to Play,” 1998 *AER*; <http://econweb.ucsd.edu/%7Evcrawfor/CrawBro98AER.pdf>

Camerer and Ho, “Experienced-Weighted Attraction Learning in Normal Form Gamesm” 1999 *Econometrica*; <http://www.jstor.org/stable/2999459>

Camerer and Ho, “Experience-Weighted Attraction Learning in Coordination Games: Probability Rules, Heterogeneity, and Time-Variation,” 1998 *Journal of Mathematical Psychology*; <http://dx.doi.org/10.1006/jmps.1998.1217>

Van Huyck et al.'s (1990 *AER*, 1991 *QJE*) subjects chose simultaneously among 7 efforts, with payoffs and optimal choices determined by their own efforts and the median or minimum effort in large groups (9 to 16 subjects), or the current pair's minimum with random pairing.

(I focus now on the pairing treatment with random re-pairing, but later compare that treatment with initially random but then fixed pairing.)

The median or minimum was publicly announced after each play, and the structure was publicly announced at the start.

Van Huyck et al.'s (1990 *AER*, 1991 *QJE*) subjects chose simultaneously among 7 efforts, with payoffs and optimal choices determined by their own efforts and the median or minimum effort in large groups (9 to 16 subjects), or the current pair's minimum with random pairing.

(I focus now on the pairing treatment with random re-pairing, but later compare that treatment with initially random but then fixed pairing.)

The median or minimum was publicly announced after each play, and the structure was publicly announced at the start.

The stage games (next slide) have seven strict, symmetric, Pareto-ranked equilibria: like a meeting that all would prefer to start on time, but which can't start until a given quorum is achieved, with waiting costly.

(They are multiple-effort versions of two-effort Stag Hunt or Bank Runs games; the larger strategy spaces make the dynamics more informative.)

There is an “obviously” right way to play, but the Pareto-dominant equilibrium is intuitively more fragile, the larger the quorum or the group.

PAYOFF TABLE Γ

		Median value of X chosen						
		7	6	5	4	3	2	1
Your choice of X	7	1.30	1.15	0.90	0.55	0.10	-0.45	-1.10
	6	1.25	1.20	1.05	0.80	0.45	0.00	-0.55
	5	1.10	1.15	1.10	0.95	0.70	0.35	-0.10
	4	0.85	1.00	1.05	1.00	0.85	0.60	0.25
	3	0.50	0.75	0.90	0.95	0.90	0.75	0.50
	2	0.05	0.40	0.65	0.80	0.85	0.80	0.65
	1	-0.50	-0.05	0.30	0.55	0.70	0.75	0.70

PAYOFF TABLE A

		Smallest Value of X Chosen						
		7	6	5	4	3	2	1
Your Choice of X	7	1.30	1.10	0.90	0.70	0.50	0.30	0.10
	6	-	1.20	1.00	0.80	0.60	0.40	0.20
	5	-	-	1.10	0.90	0.70	0.50	0.30
	4	-	-	-	1.00	0.80	0.60	0.40
	3	-	-	-	-	0.90	0.70	0.50
	2	-	-	-	-	-	0.80	0.60
	1	-	-	-	-	-	-	0.70

The five leading treatments all evoked similar initial responses, and subjects almost always converged to *some* stage-game equilibrium.

But the dynamics and limiting outcomes varied with the order statistic, interaction pattern, and group size, with very large differences in drift, history-dependence, and rate of convergence.

The five leading treatments all evoked similar initial responses, and subjects almost always converged to *some* stage-game equilibrium.

But the dynamics and limiting outcomes varied with the order statistic, interaction pattern, and group size, with very large differences in drift, history-dependence, and rate of convergence.

The limiting coordination outcome was less efficient, the more fragile were the efficient equilibria.

And holding fragility (as measured by the order statistic) constant, the limiting coordination outcome was less efficient in larger groups.

The five leading treatments all evoked similar initial responses, and subjects almost always converged to *some* stage-game equilibrium.

The five leading treatments all evoked similar initial responses, and subjects almost always converged to *some* stage-game equilibrium.

But the dynamics and limiting outcomes varied with the order statistic, interaction pattern, and group size, with very large differences in drift, history-dependence, and rate of convergence.

The limiting coordination outcome was less efficient, the more fragile were the efficient equilibria.

And holding fragility (as measured by the order statistic) constant, the limiting coordination outcome was less efficient in larger groups.

The five leading treatments all evoked similar initial responses, and subjects almost always converged to *some* stage-game equilibrium.

As explained in Crawford (1995, 2001), the dynamics sharply distinguish among equilibrium, however refined; “rational learning”; evolution; long-run equilibrium; quantal response equilibrium, and history-dependent adaptive learning; and they strongly favor the latter class of models.

Strategic teaching

Strategic teaching (Camerer, Ho, and Chong 2002 *JET*) means deviating from the decision that maximizes current-period expected payoffs, with the goal of beneficially influencing others' future decisions.

In Van Huyck et al.'s (1990 *AER*) random re-pairing treatment, teaching is costly but pointless, and there was no trend toward efficient equilibria.

TABLE 5—DISTRIBUTION OF ACTIONS FOR TREATMENT C:
RANDOM PAIRINGS

	Period				
	21	22	23	24	25
Experiment 6					
No. of 7's	5	5	4	10	8
No. of 6's	0	1	3	0	0
No. of 5's	2	5	3	3	4
No. of 4's	3	1	1	1	1
No. of 3's	1	1	1	0	0
No. of 2's	1	1	2	2	2
No. of 1's	4	2	2	0	1
Experiment 7					
No. of 7's	—	—	6	5	5
No. of 6's	—	—	1	0	1
No. of 5's	—	—	0	3	0
No. of 4's	—	—	2	1	4
No. of 3's	—	—	2	0	0
No. of 2's	—	—	0	0	1
No. of 1's	—	—	3	5	3

But in Van Huyck et al.'s parallel treatment with initially random but then fixed pairing, relationships have a future, and 12 of 14 pairs reached the most efficient equilibrium (via various repeated-game strategies).

TABLE 4—EXPERIMENTAL RESULTS FOR TREATMENT C:
FIXED PAIRINGS

	Period						
	21	22	23	24	25	26	27
Experiment 5							
Pair 1							
Subject 1	7	7	7	7	7	7	7
Subject 16	7	7	7	7	7	7	7
Minimum	7*	7*	7*	7*	7*	7*	7*
Pair 2							
Subject 2	7	2	7	7	7	7	7
Subject 15	1	7	3	6	7	7	7
Minimum	1	2	7	7	7	7	7
Pair 3							
Subject 3	1	1	1	1	1	1	1
Subject 14	1	1	7	1	1	1	7
Minimum	1*	1*	1	1*	1*	1*	1
Pair 4							
Subject 4	1	7	7	7	7	7	7
Subject 13	7	2	5	7	7	7	7
Minimum	1	2	5	7*	7*	7*	7*
Pair 5							
Subject 5	1	7	4	7	7	7	7
Subject 12	1	4	7	7	7	7	7
Minimum	1	4	4	7*	7*	7*	7*
Pair 6							
Subject 6	5	7	7	7	7	7	7
Subject 11	7	7	7	7	7	7	7
Minimum	5	7*	7*	7*	7*	7*	7*
Pair 7							
Subject 7	1	7	6	7	7	7	7
Subject 10	5	3	6	7	7	7	7
Minimum	1	3	6*	7*	7*	7*	7*

Pair 8							
Subject 8	7	6	6	7	7	7	7
Subject 9	3	5	7	7	7	7	7
Minimum	3	5	6	7*	7*	7*	7*
Experiment 6							
Pair 1							
Subject 2	7	7	4	5	6	6	7
Subject 15	2	3	6	6	7	7	7
Minimum	2	3	4	5	6	6	7*
Pair 2							
Subject 3	5	7	7	7	7	7	7
Subject 14	7	7	7	7	7	7	7
Minimum	5	7*	7*	7*	7*	7*	7*
Pair 3							
Subject 4	1	1	1	1	4	4	1
Subject 13	7	1	1	3	1	1	2
Minimum	1	1*	1*	1	1	1	1
Pair 4							
Subject 5	5	7	7	7	7	7	7
Subject 12	7	7	7	7	7	7	7
Minimum	5	7*	7*	7*	7*	7*	7*

TABLE 4—FIXED PAIRINGS, Continued

	Period						
	21	22	23	24	25	26	27
Pair 5							
Subject 6	4	5	7	7	7	7	7
Subject 11	4	5	7	7	7	7	7
Minimum	4*	5*	7*	7*	7*	7*	7*
Pair 6							
Subject 7	5	7	7	7	7	7	7
Subject 10	5	7	7	7	7	7	7
Minimum	5*	7*	7*	7*	7*	7*	7*

* — Denotes a mutual best-response outcome.

Camerer et al. (2002) modeled strategic teaching in beauty contest and borrower-lender trust games by allowing heterogeneity in behavior:

Player roles were filled by myopic adaptive learners or sophisticated and forward-looking players, each with positive prior probability.

Sophisticated players play an equilibrium in a game among themselves, with the adaptive learners treated as a mechanical part of the game.

In Van Huyck et al.'s treatment with initially random, then fixed pairing, modeling subjects' behavior remains a challenge (Crawford 2002 *JET*).

Learning from imperfect analogies

To date almost all theoretical or experimental analyses of learning in games have concerned learning to play a single fixed game, with past plays and decisions perfectly analogous to present plays and decisions.

Real analogies are far less perfect, yet people seem to learn from them.

Learning from imperfect analogies

To date almost all theoretical or experimental analyses of learning in games have concerned learning to play a single fixed game, with past plays and decisions perfectly analogous to present plays and decisions.

Real analogies are far less perfect, yet people seem to learn from them.

Such learning requires that people interpret their experience using general principles that bridge imperfect analogies.

Equilibrium and refinements do just that, but are not reliable descriptions of behavior, and beg the questions that motivate learning analyses.

Rankin, Van Huyck, and Battalio (2000 *GEB*) studied subjects' limiting behavior in repeated play of similar two-player Stag Hunt games with two symmetric pure-strategy equilibria, "all-Stag" and "all-Rabbit".

All-Stag is Pareto-superior to all-Rabbit, but Stag is riskier in that unless all others play Stag, a player does better with Rabbit.

	Stag	Rabbit
Stag	2 2	1 0
Rabbit	0 1	1 1

2-person Stag Hunt

Rankin, Van Huyck, and Battalio (2000 *GEB*) studied subjects' limiting behavior in repeated play of similar two-player Stag Hunt games with two symmetric pure-strategy equilibria, "all-Stag" and "all-Rabbit".

All-Stag is Pareto-superior to all-Rabbit, but Stag is riskier in that unless all others play Stag, a player does better with Rabbit.

	Stag	Rabbit
Stag	2 2	1 0
Rabbit	0 1	1 1

2-person Stag Hunt

Rankin et al. disabled mechanical learning by perturbing the payoffs and action labels each period, so that subjects could use their experience in previous plays only via deductive analogies between games.

Previous results for repeated play of identical Stag Hunt games strongly favored the Pareto-inferior, risk-dominant equilibrium.

But seven of Rankin et al.'s seven subject groups either converged or appeared to be converging to the payoff-dominant equilibrium.

Perhaps disabling mechanical learning made Pareto-dominance salient.

Previous results for repeated play of identical Stag Hunt games strongly favored the Pareto-inferior, risk-dominant equilibrium.

But seven of Rankin et al.'s seven subject groups either converged or appeared to be converging to the payoff-dominant equilibrium.

Perhaps disabling mechanical learning made Pareto-dominance salient.

Van Huyck and Battalio (*JET* 2002) continue this investigation in a class of discrete 2x2 bargaining games with two Pareto-efficient equilibria, one "utilitarian" and one "Rawlsian", again disabling mechanical learning.

5 of Van Huyck and Battalio's 26 subject groups appeared to be converging to an equilibrium, 4 utilitarian and 1 Rawlsian.

The nature of the analogies across games that guided Rankin *et al.*'s and Van Huyck and Battalio's subjects' behavior is still an open question.

Supporting cooperation in long-term relationships

Long-term relationships are often modeled as repeated games, in which a given stage game is played over and over again by the same players.

Analysis focuses on whether repeated interaction can overcome short-run incentive problems (usually with no coordination problems) via credible (subgame-perfect equilibrium) threats to end a relationship if behavior does not follow the “implicit contract”.

Supporting cooperation in long-term relationships

Long-term relationships are often modeled as repeated games, in which a given stage game is played over and over again by the same players.

Analysis focuses on whether repeated interaction can overcome short-run incentive problems (usually with no coordination problems) via credible (subgame-perfect equilibrium) threats to end a relationship if behavior does not follow the “implicit contract”.

There is usually a huge number of subgame-perfect equilibria/implicit contracts, and players have opposing preferences over many of them.

Yet the standard assumption is that players’ beliefs are perfectly coordinated on one of them; and equilibrium threats severely punish small deviations, even though they could be due to misunderstanding.

Behaviorally optimal implicit contracts and their enforcement is an important and largely open topic for theorists and experimentalists.

Bargaining

Bargaining theory includes Nash's (1950, 1953 *Econometrica*) unstructured/cooperative theory and Rubinstein's (1982 *Econometrica*) and others' structured/noncooperative subgame-perfect equilibrium analysis of ultimatum and alternating-offers bargaining games.

Bargaining experiments include many on structured bargaining, and some on unstructured bargaining by Roth and collaborators (see Roth, "Bargaining Phenomena and Bargaining Theory," in Roth (ed.), *Laboratory Experimentation in Economics: Six Points of View*, 1987

In my view the theory of unstructured bargaining and experiments on unstructured bargaining are more fundamental and more interesting: Bargaining in the world is usually close to unstructured, and such designs allow more informative tests of theories of bargaining.

In Roth's and collaborators' unstructured bargaining experiments, pairs of subjects bargained over 100 lottery tickets, with each subject's share determining his probability of winning the larger of two monetary prizes, specific to him (Roth and Malouf's 1979 "binary lottery procedure").

In Roth's and collaborators' unstructured bargaining experiments, pairs of subjects bargained over 100 lottery tickets, with each subject's share determining his probability of winning the larger of two monetary prizes, specific to him (Roth and Malouf's 1979 "binary lottery procedure").

Subjects could make any binding proposal they wished, or accept their partner's latest proposal, at any time.

They could also send nonbinding messages at any time, except that they could not identify themselves or, in some treatments, reveal their prizes.

If subjects could agree how to share the lottery tickets by an announced deadline the agreement was enforced; otherwise they got nothing.

In Roth's and collaborators' unstructured bargaining experiments, pairs of subjects bargained over 100 lottery tickets, with each subject's share determining his probability of winning the larger of two monetary prizes, specific to him (Roth and Malouf's 1979 "binary lottery procedure").

Subjects could make any binding proposal they wished, or accept their partner's latest proposal, at any time.

They could also send nonbinding messages at any time, except that they could not identify themselves or, in some treatments, reveal their prizes.

If subjects could agree how to share the lottery tickets by an announced deadline the agreement was enforced; otherwise they got nothing.

The environment was public knowledge, with exceptions noted below.

The binary lottery procedure controls subjects' unobservable vN-M utility functions, making them risk-neutral in lottery tickets even if not in money.

With the structure publicly announced except sizes of and information about prizes, this made subjects' preferences over the lottery tickets they bargained about public knowledge, and allowed direct tests of theories that assume common knowledge of the structure of the bargaining game.

The designs exploit invariances created by the binary lottery procedure to test both cooperative and noncooperative theories of bargaining.

Bargaining over lottery tickets is theoretically equivalent to a complete-information divide the dollar game with risk-neutral players, whose symmetry leads cooperative theories to predict equal division.

The designs exploit invariances created by the binary lottery procedure to test both cooperative and noncooperative theories of bargaining.

Bargaining over lottery tickets is theoretically equivalent to a complete-information divide the dollar game with risk-neutral players, whose symmetry leads cooperative theories to predict equal division.

These conclusions are independent of subjects' risk preferences, prizes, or information about prizes, so that cooperative theories can be tested by observing the effects of varying those factors.

The designs exploit invariances created by the binary lottery procedure to test both cooperative and noncooperative theories of bargaining.

Bargaining over lottery tickets is theoretically equivalent to a complete-information divide the dollar game with risk-neutral players, whose symmetry leads cooperative theories to predict equal division.

These conclusions are independent of subjects' risk preferences, prizes, or information about prizes, so that cooperative theories can be tested by observing the effects of varying those factors.

Noncooperative theories are harder to test this way because their predictions may depend on the details of the structure, the binary lottery procedure also yields invariances that allow such tests.

The results firmly reject both Nash's unstructured/cooperative theory and Rubinstein's structured/noncooperative theories of bargaining.

Subjects' agreements roughly followed a common hierarchy of equal-sharing norms in which subjects implemented the most "relevant" norm their public knowledge allowed.

The results firmly reject both Nash's unstructured/cooperative theory and Rubinstein's structured/noncooperative theories of bargaining.

Subjects' agreements roughly followed a common hierarchy of equal-sharing norms in which subjects implemented the most "relevant" norm their public knowledge allowed.

When subjects had a choice to reveal or withhold private information that was relevant to such norms, they did so strategically.

Their agreements were then determined by the most relevant norm they could implement, using public knowledge plus the private information they had incentives to reveal, given it would be used this way.

The results firmly reject both Nash's unstructured/cooperative theory and Rubinstein's structured/noncooperative theories of bargaining.

Subjects' agreements roughly followed a common hierarchy of equal-sharing norms in which subjects implemented the most "relevant" norm their public knowledge allowed.

When subjects had a choice to reveal or withhold private information that was relevant to such norms, they did so strategically.

Their agreements were then determined by the most relevant norm they could implement, using public knowledge plus the private information they had incentives to reveal, given it would be used this way.

The manipulation of norms by withholding information is inconsistent with nonstrategic explanations in which subjects "try to be fair".

But their agreements can be understood using a simple strategic model, with shared ideas about fairness as coordinating principles.

Disagreements were most common when both subjects knew enough to implement more than one relevant norm, or when information was not public knowledge.

With the feasible divisions of lottery tickets and subjects' preferences public knowledge, it is natural to assume complete information.

Disagreements are then incompatible with Nash's bargaining solution or the subgame-perfect equilibrium of Rubinstein's alternating-offers model.

Roth (1985), "Toward a Focal-Point Theory of Bargaining," in Roth (ed.), *Game-Theoretic Models of Bargaining* proposes a simple miscoordination explanation of the disagreements, via a mixed-strategy equilibrium in a Nash (1953 *Econometrica*) "demand game" in which players' beliefs are focused on the norms subjects' public knowledge allowed them to implement.