

A Survey of Experiments on Communication via Cheap Talk*

Vincent Crawford

University of California, San Diego, La Jolla, California 92093-0508

Received July 28, 1997; revised August 31, 1997

This paper is a survey of experimental evidence on behavior in games with communication, focusing on environments where “talk is cheap” in the sense that players’ messages have no direct payoff implications. Also considered are some environments in which communication was permitted throughout the game, in addition to those environments in which only preplay communication was allowed. *Journal of Economic Literature* Classification Numbers: C72, C92. © 1998 Academic Press

1. INTRODUCTION

This paper is a survey of experimental evidence on behavior in games with communication, focusing on environments where “talk is cheap” in the sense that players’ messages have no direct payoff implications, and considering some environments in which communication was permitted throughout, in addition to those environments in which only preplay communication was allowed. Games with communication are just extensive-form games with special restrictions on payoffs, but evidence on such games is worth considering separately because of their economic importance and the special issues raised by the possibility that a pre-existing common language influences the interpretations of messages.

In the rest of this section, I review the theories of behavior in games with communication that are tested in the experiments discussed below.¹

* This paper is a revised and expanded version of Section 7, “Preplay Communication,” of [14], prepared for the Invited Symposium on Experimental Economics at the Econometric Society Seventh World Congress. Section 7 was omitted from [15], the published version of [14], due to space restrictions. I am grateful to John Conlisk, Miguel Costa Gomes, Mark Machina, John McMillan, Alvin Roth, and Joel Sobel for helpful advice; to Miguel Costa Gomes for able research assistance; and to the National Science Foundation for research support.

¹ Banks [2], Farrell [21], Farrell and Rabin [22], and Gibbons [24, Chapter 4] are more comprehensive introductions to the theoretical literature. Kagel and Roth [26] is an excellent overview of the literature on game experiments, and Crawford [13] is an earlier discussion of experiments that consider communication via cheap talk.

Despite the intuition suggested by the saying “talk is cheap,” messages with no direct payoff implications can be informative when players’ preferences are not too far apart. Crawford and Sobel [16] and Green and Stokey [25] introduced a simple model in which cheap talk messages can convey private information. There are two players, a Sender who observes a private information variable that is relevant to both players, and a Receiver who makes a decision that affects both players’ welfare. The Sender observes his private information and sends a message to the Receiver, who then makes his decision. The message is cheap talk in that players care only about its information content, not about the message per se.² Thus the Sender does not find lying personally costly, nor can he be punished for lying after the fact. By the same token, the Sender cannot commit himself to tell the truth if the Receiver’s anticipated responses to the various messages he might send make it advantageous to lie. (As [16] shows, this fact can be disadvantageous, as the Sender sometimes does worse than if he could commit to telling the truth, and thereby have the Receiver believe him.) The goal is to determine how much of the Sender’s private information his message will convey to the Receiver, and to study the welfare properties of the relationship between the realization of the Sender’s private information variable and the Receiver’s decision, which I shall call the *outcome*.

In a sequential equilibrium, the Sender’s message to the Receiver means, in effect, “Given the realization of my private information variable, I like what you will do when I send this message at least as much as anything I could get you to do by sending a different message.” When players’ preferences are perfectly opposed, such a message cannot convey any useful information. Then the only equilibria are “babbling” equilibria, in which the Sender’s message is uninformative and is ignored by the Receiver. (If the Receiver could do better by responding to the Sender’s message, his response would (by definition) make the Sender worse off.) This is the grain of truth in the cheap-talk intuition. However, when players’ preferences are sufficiently close, communication via cheap talk can be informative. In a pure coordination game, for instance, it is plainly consistent with equilibrium for the Sender to reveal his private information truthfully and for the Receiver to believe him.

[16] studied the informativeness of communication via cheap talk by characterizing the game’s sequential equilibria when players have similar but not identical preferences, in the sense that the realization of the Sender’s private information variable influences their ideal decisions, given perfect information, in similar ways, but their ideal decisions differ systematically.³

² Except for this special restriction on preferences, the model is a standard signaling game as in Cho and Kreps [7].

³ [25] placed less structure on preferences and gave a more complex and in some respects more general characterization of equilibria.

These assumptions include a tractable class of intermediate cases between the two extremes of identical or perfectly opposed preferences. In such cases one suspects that it will seldom be advantageous for the Sender to share all of his information with the Receiver, but it does seem likely to be beneficial for him to share some of his information, because the Receiver might thereby be motivated to make a decision that is better for both.⁴

The characterization in [16] yields several conclusions. Because there is nothing in the model to favor one language for messages over another, there is always an *inessential* multiplicity of equilibria, with many different ways to use messages to support a given equilibrium outcome. Babbling equilibria always exist, and they are the only equilibria whenever players have sufficiently different preferences. However, whenever players' preferences are similar enough there are also equilibria in which the Sender's messages reflect a kind of intentional vagueness, but are still informative. Thus there is an *essential* multiplicity of equilibria, which, given the ubiquitous babbling equilibria, is a necessary feature of interesting models of strategic communication.

In equilibrium, the Sender in effect partitions the space of possible values of his (scalar) private information variable and tells the Receiver only which partition element it falls in. The Receiver then chooses his best decision given that information. The maximum possible number of partition elements in these *partition equilibria* increases with the similarity of players' preferences. For given preferences, there is a unique equilibrium partition with any given number of elements, ranging from the one-element partition of the babbling equilibria to the equilibrium partition with the maximum possible number of elements. For each of these partitions there is a unique equilibrium outcome; and the equilibrium with the most partition elements is the most informative, in that it yields the Receiver a higher *ex ante* (that is, before he hears the Sender's message) expected payoff than any other equilibrium. Thus, the comparative statics of changing players' preferences establishes a sense (somewhat weaker than usual) in which more information can be transmitted in equilibrium when players' preferences are more similar. However, perfect information transmission can occur in equilibrium only in the limit when players' preferences are identical, and babbling equilibria persist even in this case.

The essential multiplicity of equilibria in communication games is more unyielding than multiple-equilibrium problems elsewhere in economics. Traditional equilibrium refinements, in particular, do not help to reduce it.⁵

⁴ This is certainly the case when players' preferences are identical, and seems likely to remain true when their preferences are nearly identical.

⁵ See however Blume [3], who shows that traditional refinements can be effective in a class of closely related communication games with restricted message spaces.

The problem is mainly that they are defined entirely in terms of payoffs, and therefore have no basis on which to distinguish among cheap talk messages, which all have the same direct payoff implications. This fact increases the importance of empirical information regarding the principles that govern behavior in communication games, and particularly of such information from experiments, which allow the precise control of the strategic environment needed to test theories of strategic behavior.

The ineffectiveness of traditional refinements also highlights the importance of another distinction, whether there is a preexisting common language that ties down the literal meanings of cheap-talk messages or not. This distinction is plainly of great importance in practice, but is assumed irrelevant in traditional game theory because it affects neither the payoff structure nor the theoretical possibilities for signaling. For instance, the characterization of the set of equilibria in [16] is valid whether or not there is a common language, but does not address the issue of how it might influence equilibrium selection.

Farrell [20] (written in the early middle 1980s) studied communication via cheap talk with a common language, in a model like those of [16, 25] in which the Sender has private information. A common language eliminates the inessential multiplicity of equilibria in cheap-talk games by fiat. More importantly, and less obviously, a common language also reduces the essential multiplicity of equilibria by restricting the plausible interpretations of out-of-equilibrium messages, which could otherwise be anything needed to support an equilibrium.⁶ [20] formalized the idea of a plausible interpretation in a refinement called *neologism-proofness*, which rules out equilibria for which there is an out-of-equilibrium message and a set of the Sender's types such that it is *exactly* those types who would benefit by deviating to that message, if it were taken literally. Neologism-proofness is often effective in reducing the essential multiplicity of equilibria and obtaining precise, plausible predictions (sometimes too effective, as neologism-proof equilibria may fail to exist).

Because neologism-proofness depends on the language, it differs from traditional refinements in an essential way. It is however an equilibrium notion, which uses the language to restrict the interpretations of cheap-talk messages in ways that are reminiscent of restrictions of the interpretations of costly signals like Cho and Kreps' [7] intuitive criterion. Rabin [29] defined a nonequilibrium notion for incomplete-information games called *credible message rationalizability*. Credible message rationalizability combines a credibility restriction in the spirit of neologism-proofness with the assumption that players' decisions maximize their expected payoffs given their beliefs (without assuming the coordination of beliefs that would be required

⁶ These restrictions are limited to out-of-equilibrium messages, because an equilibrium implicitly assigns meanings to any message used in it with positive prior probability.

for equilibrium). This also yields plausible restrictions on behavior, some different from those that neologism-proofness implies.

Another important role for communication is signaling players' *intentions* for future decisions. It has long been conjectured that "enough" preplay communication yields an effective (though nonbinding) agreement to play an equilibrium in the underlying game, and that such agreements tend to be efficient relative to the game's incentives (in the sense that the agreement is undominated in the set of Nash equilibria). It is of interest to consider to what extent formal analysis confirms these conjectures.

Farrell [18, 19] studied signaling intentions in simple sequential games of complete information, where an underlying game is preceded by one or more rounds of structured preplay communication in which players make nonbinding (and therefore cheap-talk) announcements about their intended decisions in the underlying game. Players again have a pre-existing common language, which is used to restrict their beliefs about the meanings of these announcements in plausible ways that bear a family resemblance to neologism-proofness, but reflects the differences between private information and intentions.

In [18] the underlying game is Battle of the Sexes. In this game achieving a desirable (both "fair," in some *ex ante* sense, and efficient) outcome requires that the players, who initially have no way to distinguish their roles in the game, find a way to break the symmetry of their roles.⁷ [18] asks to what extent players can do this when Battle of the Sexes is preceded by one or more communication rounds in which they make simultaneous announcements about their (pure) intended decisions. The difficulty of symmetry-breaking is modeled by assuming that (with or without communication) players play a symmetric equilibrium. The effectiveness of communication in a common language is modeled by assuming that players take the first pair of announcements that would be in equilibrium in the Battle of the Sexes game as the literal truth, and ignore all other announcements. (This highlights the role of a common language, without which there would be no "literal truth.")

Without communication, players play the mixed-strategy equilibrium of Battle of the Sexes, which is worse for each player than his least preferred pure-strategy equilibrium. With communication, they play a symmetric

⁷ This need for symmetry-breaking is the essential difficulty of Battle of the Sexes. Note that it is obscured by the "Fights" and "Ballet" decision labeling with which this game is usually presented, because those labels identify decisions with different payoff meanings for each player. By contrast, the labels used in evolutionary game theory, "Hawk" ("Play for your favorite equilibrium") and "Dove" ("Play for the other player's favorite equilibrium"), identify decisions with the same payoff meanings and reflect the symmetry of players' roles, allowing the difficulty of symmetry-breaking to be modeled by restricting players to symmetric equilibria.

equilibrium in which they randomize their announcements in each round until either some round yields announcements consistent with one of the pure-strategy equilibria in Battle of the Sexes, or the communication phase ends. In the former case players play the equilibrium identified by their announcements. In the latter case their history provides no basis for symmetry-breaking, so symmetry again restricts them to the mixed-strategy equilibrium of Battle of the Sexes. Thus, communication gives players a chance to break the symmetry of their roles, and thereby to do better than without communication. In the symmetric equilibrium, they share the gains equally *ex ante*. Their expected payoffs increase with the number of rounds of communication. However, even an unlimited number of rounds of communication does not yield full efficiency, so the usefulness of communication is ultimately limited by the differences in players' preferences, as with signaling private information.

In [19] the underlying game is a finite n -person matrix game. There is one round of structured preplay communication, in which one (exogenously specified) player makes a nonbinding suggestion about all players' decisions in the game. As in [18], suggestions that are consistent with rationality for all players when their beliefs are consistent with the suggestion are believed, and all other suggestions are ignored. A precise suggestion, to be believed, must be an equilibrium. [19] also allows vague suggestions, which to be believed must be consistent with rationalizability, given the restrictions the suggestion implies. It is important to consider vague suggestions because precise suggestions may not be optimal.

Farrell noted that it would be begging the question to *assume* equilibrium in the game with communication, and assumed rationalizability instead. He showed that if the underlying game is such that the largest rationalizable payoff for the player who makes suggestions is achievable in some equilibrium of the underlying game, then he must get that payoff. Thus, if that payoff can only be achieved in an equilibrium, players must play an equilibrium; and if the game is one of pure coordination, they must play an efficient equilibrium. However, games that satisfy the above condition are very special, and in other games communication may yield neither equilibrium nor efficiency.

Rabin [30, 31] extended the analyses of [18, 19], defining a nonequilibrium notion called *negotiated rationalizability* that makes it possible to combine the assumption that players maximize expected payoffs given their beliefs (without assuming coordination of beliefs) with a variety of behaviorally motivated restrictions on beliefs. He then used this notion to derive useful bounds on players' expected payoffs when they have unlimited opportunities to communicate, which generalize Farrell's [18] result for an unlimited number of rounds of communication and yield considerable additional insight.

These analyses also show that models in which cheap-talk messages signal intentions, like their private-information counterparts, have both inessential and essential multiplicities of equilibria, neither of which is reduced by traditional refinements.

So far, most experimental work on signaling via cheap talk has considered signaling intentions, and has assumed the existence of a common language. This concentration seems sensible, because in most applications signaling intentions appears to be more important than signaling private information, and deciding how a given cheap-talk message will be used strategically seems a harder problem than identifying its literal meaning. Nonetheless, signaling via cheap talk without a common language is also of interest, and may be more descriptive in applications such as negotiations between agents who share a common language but have different cultures, as well as those where agents literally have no common language. Such models have also attracted the attention of experimentalists.⁸

The rest of the paper is organized as follows. Section 2 gives some general background on the design of game experiments, following [14, Section 3]. Section 3 discusses experiments in which subjects use a common language entirely or primarily to signal intentions. Section 4 discusses experiments in which subjects use a common language to signal private information. Section 5 discusses experiments that test “evolutionary” refinements in a model of signaling private information with no common language.

2. EXPERIMENTAL DESIGNS

This section gives some general background on the designs of most game experiments in economics, including those discussed below.

A successful design must control the environment so that the results can be interpreted as responses to a clearly identified game. A typical design has one or more subject populations repeatedly, randomly, and anonymously paired to play a stage game. The goal is to test theories of behavior in the stage game, and accordingly the effects of repeated interaction are minimized by making repeated encounters unlikely or impossible. Subjects are usually told the outcome after each play, including their current partners’ strategies. To maintain control, communication and correlation are allowed only as the stage game permits them. The stage game is otherwise completely free to vary, which allows a wide range of strategic questions to be posed in tractable ways.

⁸ As will be seen in the discussion of Blume, DeJong, Kim, and Sprinkle [4] and Blume, Kim, and Sprinkle [5], it can be difficult to prevent subjects from smuggling a common language into the laboratory.

Subjects' unfamiliarity with such environments is overcome by using simple stage games, explaining them in written instructions and question and answer sessions, and providing enough experience via practice rounds or repeated play to assure meaningful responses and reveal the effects, if any, of learning.

To focus as sharply as possible on the principles that govern strategic behavior, players' uncertainty about things other than each other's private information and strategy choices are usually kept to a minimum. Control over information is achieved by publicly announcing the structure of the environment at the start. The resulting condition, called *public knowledge*, comes as close as possible in the laboratory to inducing the theoretical condition of common knowledge of the structure.

Control over preferences is achieved by paying subjects according to their payoffs. Nonpecuniary influences on preferences are usually suppressed by avoiding face-to-face or nonanonymous interactions and frames with psychological associations. (Face-to-face or nonanonymous interactions are avoided by allowing subjects to communicate through a computer interface, or by passing written messages through the experimenters.) Subjects' payments are normally linear functions of their game payoffs, with the results analyzed assuming risk-neutrality. Sometimes, as in the "binary lottery" procedure of Roth and Malouf [35], each subject is rewarded with a probability, again a linear function of his payoff, of winning a given amount of money (or the larger of two possible amounts). Under standard assumptions subjects then maximize the probability of winning the prize (or the larger prize), and are therefore risk-neutral in a variable under experimental control.

The designs of the experiments discussed below all involve one or two populations of subjects repeatedly, randomly paired (or in some cases grouped) to play a stage game.⁹ Subjects are motivated by one of the above methods, and the structure is public knowledge.

Because behavior in the laboratory often takes time to stabilize, theories that assume equilibrium are usually tested by comparing their predictions to observed behavior in the last period or periods of a treatment. By contrast, theories that assume rationality without requiring equilibrium can be, and sometimes are, tested using the entire sample.

3. SIGNALING INTENTIONS WITH A COMMON LANGUAGE

Experimental work on signaling intentions with a common language has focused on simple coordination games or bargaining games in which

⁹ With two populations, each population fills one of the player roles in the stage game.

coordination plays a leading role (with “bargaining” construed broadly to include games like Battle of the Sexes). In some designs communication has been structured, with one or more rounds in which subjects make simultaneous or one-sided announcements about their intentions (Clark, Kay, and Sefton [8], Cooper, DeJong, Forsythe, and Ross [9–11], Palfrey and Rosenthal [28]). In others communication has been unstructured, with the timing and content of messages largely unrestricted (Forsythe, Kennan, and Sopher [23], Roth [32, 33], Valley, Thompson, Gibbons, and Bazerman [37]). The common-language assumption has been implemented either by requiring subjects to communicate only in a restricted language that consists of the labels with which their decisions in the underlying game are presented, or simply by allowing them to discuss their alternatives using natural language.

There are two important determinants of the effects of structured communication in these experiments: whether communication is simultaneous or one-sided; and whether players can achieve a desirable outcome without breaking the symmetry of their roles.

When a desirable outcome requires symmetry-breaking, one-sided communication is much more effective than two-sided communication, with coordination frequencies of 95% in Battle of the Sexes (almost always on the Sender’s favorite equilibrium), versus 55% with two-sided communication and 41% with no communication [9, 10, 11]. With two-sided communication, multiple rounds are more effective than one round, with coordination frequencies of 63% versus 55% [9], but these stop well short of yielding efficiency. These results are very close to predictions based on Farrell’s [18, 19] analyses.¹⁰

When desirable outcomes do not require symmetry-breaking, as in the Stag Hunt games used in the studies discussed here, communication appears to play an important *reassurance* role, allowing subjects to coordinate on more efficient equilibria by reducing their uncertainty about each other’s decisions. By contrast with symmetry-breaking, reassurance is best accomplished via two-sided communication, which yielded frequencies of coordination on the efficient equilibrium of 90% , versus 53% with one-sided communication and virtually zero without communication [11].

¹⁰ Palfrey and Rosenthal [28] found support for an incomplete-information generalization of the analysis in [18] in a three-person public goods experiment in which subjects privately observed their contribution costs, made one round of simultaneous, nonbinding announcements, and then simultaneously chose whether or not to contribute. Subjects were very successful in using their opportunities to communicate to achieve the coordination of contribution decisions required for efficient provision of the public good, and even had some success, despite the incentives to lie, in eliciting contributions from those whose costs were lowest.

These results appear to contradict the theoretical argument in Aumann [1] (see also Farrell [19]) that communication via cheap talk must be ineffective in such games. (Aumann's argument is presented for a Stag Hunt game in which both players have strict preferences over their partners' decisions, independent of their own intentions, but his logic seems to extend to the game in [11], in which one of those preferences is weak.) However, recent experiments reported in Clark, Kay, and Sefton [8] observe large, significant differences in signals and decisions between games that differ only in the strictness of subjects' preferences over one of their partners' decisions. Clark, Kay, and Sefton [8] replicated Cooper, DeJong, Forsythe, and Ross's findings [11], both with and without communication, for a Stag Hunt game in which subjects had weak preferences over one of their partners' decisions; but they found that when subjects had strict preferences over one of their partners' decisions, messages announcing the intention to play the efficient equilibrium were far less credible (although they were made just as often). Overall, Clark, Kay, and Sefton's results tend to confirm the reassurance interpretation of the role of communication offered above, while suggesting that the logic behind it is delicate [8].

It is intriguing to compare the results of [9, 10, 11] with those of the experimental study of "0-sided" communication in Van Huyck, Gillette, and Battalio [38], in which subjects were not allowed to communicate but were given a public, nonbinding suggestion about which equilibrium to play. Suggestions of one of three isomorphic, symmetric, and efficient equilibria were very effective, with a coordination rate of 98% versus 40% without a suggestion; but suggestions of inefficient or "unfair" equilibria were far less effective, with coordination rates of 17–25% and 16% respectively.

In unstructured bargaining, Roth [32, 33] found that communication via cheap talk tended to focus subjects' beliefs sharply on the small number (one or two) of fairness norms their public knowledge allowed them to implement. Otherwise its effects were mainly confined to the strategic revelation of private information relevant to norms (see also Forsythe, Kennan, and Sopher [23] and the discussion in [15, Section 5.3]).

Costa Gomes [12] reconsidered the evidence from the experiments of [32, 33] and the evidence on structured communication from [9, 11], examining data from the entire samples and using Rabin's [30, 31] notion of *negotiated rationalizability*. The power of Rabin's approach depends on imposing empirically sensible restrictions on players' strategies. Costa Gomes limited subjects to insisting on one of the equilibria that is both efficient and "fair" in one of the senses their information allowed them to implement, which is approximately true in the samples. He found that most of the results of those experiments satisfy the restrictions implied by negotiated rationalizability.

4. SIGNALING PRIVATE INFORMATION WITH A COMMON LANGUAGE

There is a small amount of experimental evidence on signaling private information with a common language. Studying games in which partition equilibria take simple forms, Dickhaut, McCabe, and Mukherji [17] found some support for Crawford and Sobel's [16] comparative statics result that increasing the difference between the Sender's and the Receiver's preferences reduces the informativeness of communication. In similar settings, Sopher and Zapater [36] found support for Rabin's [29] credible message rationalizability.

Valley, Thompson, Gibbons, and Bazerman [37] studied a double oral auction both without communication and with communication via cheap talk, either written or face-to-face. They found that either written or face-to-face communication often allowed subjects to coordinate on single price or split-the-difference outcomes that revealed enough of their private information to take them outside the incentive-efficiency bound for unrestricted mechanisms in this setting derived by Myerson and Satterthwaite [27]. They concluded that talk, even when it is anonymous and not face-to-face, may not be cheap (see also Palfrey and Rosenthal [28], discussed above). They also offered some fascinating speculations on when written communication is better than face-to-face communication (see also Roth [32, pp. 294–302]).

5. SIGNALING PRIVATE INFORMATION WITH NO COMMON LANGUAGE

Although the theoretical literature on communication includes several analyses, both traditional and adaptive, in which subjects had no common language, Blume, DeJong, Kim, and Sprinkel [4] and Blume, Kim, and Sprinkel [5] appear to be the only such experimental studies. Both studied the nonequilibrium, "evolutionary" refinements of Blume, Kim, and Sobel [6] and others, which predict that informative signaling will emerge when players' preferences are sufficiently close.¹¹ In some treatments they used abstract message spaces and a label-scrambling technique to eliminate any possible effects of a common language. The results from those treatments

¹¹ These refinements seem most appropriate for environments in which the meanings of messages and their strategic uses evolve at comparable speeds. By contrast, notions like neologism-proofness and credible message rationalizability, which assume a preexisting common language, seem most appropriate for environments in which the meanings of messages evolve much more slowly than their uses, so that the dynamics of the uses of messages can be analyzed with their meanings held constant.

were then compared with those for the analogous treatments with a common language. With or without a common language, informative communication usually emerged, though more slowly without a common language. There was significant support for evolutionary refinements in games with closely aligned preferences, but less support in games with divergent preferences, where such refinements are less restrictive.

REFERENCES

1. R. Aumann, Nash equilibria are not self-enforcing, in "Economic Decision-Making: Games, Econometrics and Optimization" (J.-J. Gabszewicz, J.-F. Richard, and L. A. Wolsey, Eds.), Elsevier, Amsterdam/New York, 1990.
2. J. Banks, "Signaling Games in Political Science," Harwood Academic, Reading, UK, 1991.
3. A. Blume, Equilibrium refinements in sender-receiver games, *J. Econ. Theory* **64** (1994), 66–77.
4. A. Blume, D. DeJong, Y.-G. Kim, and G. Sprinkle, "Evolution of the Meaning of Messages in Sender-Receiver Games: An Experiment," Working Paper 94-03, Department of Economics, University of Iowa, 1994.
5. A. Blume, Y.-G. Kim, and G. Sprinkle, "Evolution of Communication with Partial Common Interest," manuscript, Department of Economics, University of Iowa, 1997.
6. A. Blume, Y.-G. Kim, and J. Sobel, Evolutionary stability in games of communication, *Games Econom. Behav.* **5** (1993), 547–575.
7. I.-K. Cho and D. Kreps, Signaling games and stable equilibria, *Quart. J. Econ.* **102** (1987), 179–221.
8. K. Clark, S. Kay, and M. Sefton, When are Nash equilibria self-enforcing? An experimental analysis, *Internat. J. Game Theory*, in press.
9. R. Cooper, D. DeJong, R. Forsythe, and T. Ross, Communication in the Battle of the Sexes game: Some experimental results, *Rand J. Econ.* **20** (1989), 568–587.
10. R. Cooper, D. DeJong, R. Forsythe, and T. Ross, Communication in coordination games, *Quart. J. Econ.* **107** (1992), 739–771.
11. R. Cooper, D. DeJong, R. Forsythe, and T. Ross, Alternative institutions for resolving coordination problems: Experimental evidence on forward induction and preplay communication, in "Problems of Coordination in Economic Activity" (J. Friedman, Ed.), Kluwer Academic, Dordrecht/Norwell, MA, 1994.
12. M. Costa Gomes, "A Suggested Interpretation of Some Experimental Results on Preplay Communication," manuscript, Department of Economics, UCSD, 1997.
13. V. Crawford, Explicit communication and bargaining outcomes, *Amer. Econ. Rev., Papers and Proceedings* **80** (1990), 213–219.
14. V. Crawford, "Theory and Experiment in the Analysis of Strategic Interaction," Discussion Paper 95-37, Department of Economics, UCSD, 1995.
15. V. Crawford, Theory and experiment in the analysis of strategic interaction, in "Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress" (D. Kreps and K. Wallis, Eds.), Vol. I, Cambridge Univ. Press, Cambridge, UK, 1997.
16. V. Crawford and J. Sobel, Strategic information transmission, *Econometrica* **50** (1982), 1431–1452.
17. J. Dickhaut, K. McCabe, and A. Mukherji, An experimental study of strategic information transmission, *Econom. Theory* **6** (1995), 389–403.
18. J. Farrell, Cheap talk, coordination, and entry, *Rand J. Econ.* **18** (1987), 34–39.

19. J. Farrell, Communication, coordination and Nash equilibrium, *Econ. Lett.* **27** (1988), 209–214; see also “Erratum,” **33** (1990), 299.
20. J. Farrell, Meaning and credibility in cheap-talk games, *Games Econom. Behav.* **5** (1993), 514–531.
21. J. Farrell, Talk is cheap, *Amer. Econ. Rev., Papers and Proceedings* **85** (1995), 186–190.
22. J. Farrell and M. Rabin, Cheap talk, *J. Econ. Persp.* **10** (1996), 103–118.
23. R. Forsythe, J. Kennan, and B. Sopher, An experimental analysis of strikes in bargaining games with one-sided private information, *Amer. Econ. Rev.* **81** (1991), 253–270.
24. R. Gibbons, “Game Theory for Applied Economists,” Princeton Univ. Press, Princeton, NJ, 1992.
25. J. Green and N. Stokey, “A Two-Person Game of Information Transmission,” Discussion Paper 751, Harvard Institute of Economic Research, 1980.
26. J. Kagel and A. Roth, Eds., “Handbook of Experimental Economics,” Princeton Univ. Press, Princeton, NJ, 1995.
27. R. Myerson and M. Satterthwaite, Efficient mechanisms for bilateral trading, *J. Econ. Theory* **29** (1983), 265–281.
28. T. Palfrey and H. Rosenthal, Testing for effects of cheap talk in a public goods game with private information, *Games Econom. Behav.* **3** (1991), 183–220.
29. M. Rabin, Communication between rational agents, *J. Econ. Theory* **51** (1990), 144–170.
30. M. Rabin, “Focal Points in Pre-Game Communication,” Working Paper 91-179, Department of Economics, UC Berkeley, 1991.
31. M. Rabin, A model of pre-game communication, *J. Econ. Theory* **63** (1994), 370–391.
32. A. Roth, Toward a focal-point theory of bargaining, in “Game-Theoretic Models of Bargaining” (A. Roth, Ed.), Cambridge Univ. Press, Cambridge, UK, 1985.
33. A. Roth, Bargaining phenomena and bargaining theory, in “Laboratory Experimentation in Economics: Six Points of View” (A. Roth, Ed.), Cambridge Univ. Press, Cambridge, UK, 1987.
34. A. Roth, Bargaining experiments, in “Handbook of Experimental Economics” (J. Kagel and A. Roth, Eds.), Princeton Univ. Press, Princeton, NJ, 1995.
35. A. Roth and M. Malouf, Game-theoretic models and the role of information in bargaining, *Psychol. Rev.* **86** (1979), 574–594.
36. B. Sopher and I. Zapater, “Communication and Coordination in Signalling Games: An Experimental Study,” manuscript, Department of Economics, Brown University, 1993.
37. K. Valley, L. Thompson, R. Gibbons, and M. Bazerman, “Is Talk Really Cheap? Outperforming Equilibrium Models of Communication in Bargaining Games,” Working Paper 96-022, Harvard Business School, 1996.
38. J. Van Huyck, A. Gillette, and R. Battalio, Credible assignments in coordination games, *Games Econom. Behav.* **4** (1992), 606–626.