

Leaders, Followers, and the Institutional Problem of Trust

Carl M. Rhodes

Department of Government
Franklin and Marshall College
Lancaster, PA 17604
C_RHODES@acad.fandm.edu

Department of Political Science
University of Cincinnati
Cincinnati, Ohio 45221-0375

and

Rick K. Wilson
Department of Political Science
Rice University
Houston, TX 77251-1892
rkw@rice.edu

This paper was prepared for presentation for the Trust Working Group Meeting and sponsored by the Russell Sage Foundation, February 19-20, 1999, New York, NY. Valuable assistance for the experiments came from Jane Sell and Sean Bolks.

Abstract

In this paper we ask whether leaders can serve as institutionally imposed solutions to the problem of trust. After briefly reviewing the problematic and fragile nature of trust and cooperation, we argue that leaders can encourage the formation of trust among a group of followers. Their ability to do so, however, is contingent on the willingness of followers to offer their trust to leaders. We then examine the dynamic of leadership and trust in a series of experiments. Results suggests that the ability of leaders to foster trust and cooperation is dependent on 1) the nature of the underlying context, and 2) the reputation and behavior of leaders themselves.

Introduction

The political dynamic in the United States at the end of 1998 was unprecedented. Articles of impeachment were voted for a sitting president. One Speaker and one Speaker-Elect resigned after losing the support of the majority party. And although President Clinton, Speaker Gingrich, and Congressmen Livingston were far apart on the ideological spectrum, belonged to different political parties, and held very different political views, each struggled with same dilemma as the year drew to its close: how to lead a group of people who no longer trusted them.

Whether these individuals are trustworthy, of course, is a matter of considerable importance for the public, but the public's reaction has not been the source of difficulty facing these men. Much more important is the suspicion that emanates from their colleagues in government and those who must work with them on a day-to-day basis. What we wish to argue here is that this problem is not unique: trust is an important in almost any institution, but especially in settings with leadership.

Congress is a particularly interesting institution in which to study leadership, in that leaders are merely agents of their parties. Leaders cannot fire their followers (who are elected by separate constituencies), and they are granted almost no sanctions and very few enticements to wield. However, Congress is by no means the only institution in which leaders hold few powers and serve largely at the pleasure (even amusement) of the followers. In countless city councils, neighborhood associations, academic departments, legislatures, and social groups, leading means convincing followers to first trust the leader, and then trust each other. Without the cooperation of followers, the leader's task is impossible, and the institution is likely to fail.

Leadership is hardly the only means for inducing cooperation (Ostrom 1998). Cooperation can emerge with the aid of many institutional "rules", including direct communication, affordable monitoring or self-enforcing agreements). But like leadership, all these institutional solutions to the problem of trust rely on at least two levels of trust being created: trust in the institution (e.g., trust in the leader) and trust in one's peers (e.g., trust among followers). This paper explores the

impact of leadership on trust and cooperation in social dilemmas, but it will also attempt to shed light on the impact of institutional rules and structure on the problem of trust more generally.

We proceed as follows. In the next section, we discuss the problem of leadership and its relation to trust. The second section discusses ways in which trusting reputations might emerge from leaders. The third section details a set of experiments that disentangle the impact of these factors on trust. The fourth section summarizes our experimental results. The final section concludes with some thoughts on the nature of leaders, followers, and the problem of trust.

Problems of Leadership

In many settings, the key dilemma is how to ensure cooperative behavior when individual self-interest recommends non-cooperation. Social dilemmas (Dawes, 1980), collective action problems (Olson, 1965; Hardin, 1982) and common property resource problems (Ostrom, 1990) are all examples of instances where individual actors have strong incentives to play strategies that yield a Nash equilibrium which is pareto dominated by cooperative acts.

Given the abundance and difficulty of social dilemmas, many scholars have turned their attention to leadership as an institutionally imposed solution, although explicit models and tests of the leadership process remain rare (Calvert 1992; Cox and McCubbins 1993; Kiewiet and McCubbins 1991; Miller 1992; Rohde and Shepsle 1987; Sinclair 1995; 1992; Salisbury 1969; Rohde 1991; Jillson and Wilson 1994). With the exception of our own work, the most explicit models of leadership argue that leaders “lead” by dispensing sanctions and reward (Bianco and Bates 1990, Frohlich, Oppenheimer, and Young 1971; Bendor and Mookherjee 1987; Alchian and Demsetz 1972; Holmstrom 1982; Calvert 1987; and Alt, Calvert, and Humes 1988; Bianco and Bates 1990). While the ability of leaders to sanction and reward is no doubt important in some settings, its ultimate value (and availability) in explaining the process of leadership, especially leadership within decentralized institutions, remains open to serious question (Miller 1992; Ellickson 1991; Ostrom, Gardner, and Walker 1994; Rohde 1991; Oliver 1980; Wilson 1995). In countless political and social settings, leaders are given few if any rewards to dispense or sanctions to impose, and instead

must rely on persuasion to affect follower behavior.

The introduction of leadership, although perhaps intended to solve social dilemmas, brings with it an important problem of its own. Miller (1992) refers to this as the “Madisonian Dilemma,” which covers a well-known set of principal-agent problems. If a leader has specialized knowledge and power, then how can followers be confident that the leader will not abuse that advantage? Problems of agency in decentralized settings, in which leaders are given few formal powers, are especially problematic. In these settings, leaders cannot “force” (sanction or reward) their followers into compliance, but must instead rely on persuasion. Followers, however, have only incomplete information on both the motivation and competence of the individual trying to persuade them. Given this uncertainty, should they offer the leader trust? If they do, and the leader’s interests are congruent with their own while possessing a reasonable level of competence, everyone is better off. But if they doubt the leader’s motivations or competence and withhold their trust, this institutional “fix” for social dilemmas fails. As Miller (1992) notes, leading in a complicated setting such as this comes down to trust: followers trusting both the leader and each other.

The same problem is equally apparent in other, less intuitive, locales. The standard model for primates – particularly the Great Apes and Chimpanzees – is that they are hierarchically organized, dominated by an alpha male, and that social deviance is quickly punished by that leader, much like the traditional account of the economic firm (for a recounting of the standard primate model and a wonderful critique of science, see Haraway 1989). In much the same way that Miller (1992) has called into question the strictly hierarchical model of the firm, this recent work points to the problems with trying to understand primate interactions as solely a function of hierarchy (in particular, see de Waal, 1996).

These lessons have not gone unnoticed by political scientists, some of whom have taken advantage of rich observational data on primates to try to understand small group, kinship-based politics. Glendon Schubert (1991, p. 37) makes the obvious point that the study of primates provides a useful evolutionary alternative to what we understand concerning humans. While the

study of primates cannot substitute for the study of humans (the environmental differences alone make this point), several primatologists have fruitfully turned to using concepts from political science to understand leadership among Chimpanzees. Frans de Waal (1982) paints a picture of coalition politics among a group of chimps. He was especially interested in patterns of leadership and the manner in which leaders gained support from followers. While his initial focus was on hierarchy among males, he quickly recognized that leadership was highly dependent on reciprocal linkages with females. Remaining the clear alpha-male (the leader) required building a considerable stock of favors. He finds that the alpha-male spends a great deal of time engaged in grooming, food-sharing and other behaviors that builds a reputation.

A more recent book de Waal (1996) explores the sources of cooperation, finding that reciprocity (even among unrelated members) builds concepts of trust among members. Reciprocity is built through a complex set of repeated actions. These involve grooming pairs, food sharing and alliance formation. In his book, de Waal (see particularly Chapter 4, 1996) discusses a large number of food sharing experiments in which it is clear that reciprocal norms are quite well defined among chimpanzees. Moreover, there are clear patterns in which some pairings share a great deal with one another and other pairings share very little. Most surprising, perhaps, is the asymmetric sharing by the dominant male. The dominant male tends to share more with others than he gets in return. In part this can be seen as building trust, in which others recognize the dominant male's willingness to share and an implicit expectation that such behavior will be reciprocated when the dominant male needs support. As de Waal puts it when writing about Ntologi, an alpha male in one of the groups studied,

Ntologi chiefly shared with females, as well as with those males who were unlikely to threaten his position; that is, influential aging males and prime adult males in stable midranking positions. While these males lacked the social status and/or physical capacity to rise to the top, they could be effective allies. (1996, p. 143)

From this brief overview of a number of perspectives, it seems clear that if leaders are to help followers solve social dilemmas, they are unlikely to do so through the exercise of raw power alone. More often than not, the link between leaders and followers critically depends on trust.

From Congress, to social and political groups, to even Chimpanzees, leadership is about inspiring trust in both the leader and among followers. In the next section, we consider the problem of trust, and how a leader might encourage it, more directly.

Embedding the Problem of Trust in Leadership

In thinking about trust, we explicitly follow Ostrom (1998). Trust involves one person forming expectations about the likely actions of others, which then affects the individual's strategic choice (1998, p. 12). In a wide variety of settings this means that if Actor A *trusts* Actor B, then A will choose a strategy over which B can take advantage, thereby leaving A worse off. However, if B is *trustworthy* (reciprocates that trust), then A does better by engaging in a trusting action (and so too does B). In many settings, actors can do far better if they can achieve a cooperative equilibrium path, trusting and being trustworthy, rather than playing some form of a Nash equilibrium (or sub-game perfection).

A large number of empirical papers show that laboratory subjects often fail to play a Nash best-response, and eschew playing even dominant strategies. Many of these papers are centrally concerned with actors looking beyond themselves and to their partners. A common finding is the nexus of trusting and reciprocal behavior. This other-regarding behavior occurs in many settings including ultimatum and dictator games (Camerer, 1997; Eckel and Grossman, 1996, 1998; Forsythe, et al., 1994; Hoffman, et al., 1994), public goods games (Ledyard, 1995), investment trust games (Berg et al., 1995) and gift exchange experiments (Fehr, et al., 1993). These results are neither random nor haphazard. Behavior inconsistent with game theoretic predictions is routine and patterned.

In a similar sense Ostrom (1998) is concerned with the ways in which small groups, largely self-organized, rely on trust, reciprocity and reputation to enhance levels of cooperation. As she argues, "at the core of a behavioral explanation are the links between the trust that individuals have in others, the investment others make in trustworthy reputations, and the probability that participants will use reciprocity norms." (1998, p. 12). She further argues that

this core set of relationships is critically affected by institutional variables that mediate the interaction of actors. In her own words, “the really big puzzle in the social sciences is the development of a consistent theory to explain why cooperation levels vary so much and why specific configurations of situational conditions increase or decrease cooperation in first- or second-level dilemmas.” (1998, p. 9)

We agree that this core set of relationships, and the institutions that surround them, is critically important. We are particularly interested in the set of institutional arrangements that encourage the formation of trust. While there may be trusting people, we doubt that evolution has created individuals who are always cooperators. Instead, we think that cooperation is very likely conditional. Humans are extraordinarily sensitive to their context and to the institutional mechanisms that enhance (or discourage) both trust and trustworthiness. To explore this dynamic, we dissect a very simple institution involving leadership.

To date, much of the concern with cooperative behavior and the locus of trust has focused on very minimal institutions (see discussions of minimal rule sets in Ostrom, Gardner and Walker, 1994). Ordinarily those institutions are egalitarian, where no one has a special position. However, in many (if not most) non-hierarchical institutions, there are individuals who are granted a distinct role with special endowments. These individuals, who we call leaders, may serve to hear complaints, make judgments, dispense learning, or recommend particular courses of action.

Clearly, there are many institutional features and rules that might encourage or discourage the formation of trust, but we have chosen to examine the institution of leadership for three reasons. First, the position of “leader” is nearly ubiquitous. Indeed, it is difficult to think of a social or political group, organization, or institution that does not have a leadership position of some kind. Second, and as noted above, leadership is widely hailed as an institutional savior of sorts, despite the fact that few explicit models or tests of the leadership dynamic have been conducted. Finally, we suspect that studying the impact of leaders on trust will shed light on other institutional solutions to social dilemmas as well, in that before an institution can encourage the formation of trust, individuals must accept and trust the institution. Below, we argue that before

leaders can encourage the formation of trust among followers, they must first convince followers to trust in leaders themselves.

In our research to date (Wilson and Rhodes 1997; Rhodes and Wilson forthcoming; Rhodes 1999a; Rhodes 1999b), we examine the impact of a very simply institution in which a leader has no sanctions or rewards. Instead, the leader holds a set of “cheap-talk” signals that can be relayed to followers. The key problem for followers is to determine whether they can trust the leader. Specifying the institution in a very simple manner allows us to explore how a weak leadership mechanism is related to trust.

For a leader to have a positive impact on trust, two things must occur. First, the leader must increase the level of intra-follower trust and cooperation. A leader can create intra-follower trust by serving as a focal point to coordinate joint strategic play (Calvert 1992; Wilson and Rhodes 1997; Rhodes and Wilson forthcoming; Rhodes 1999a; Rhodes 1999b). By suggesting a particular course of action, leaders help followers draw inferences about the likely behavior of other followers. Armed with this information, followers can engage in trusting behavior under the expectation that it will be reciprocated by others.

Second, a leader must convince followers to trust the leader herself. This is critically dependent on an individual leader’s reputation for competence (i.e., the extent to which followers believe that the leader is capable of assisting them) as well as an understanding of the leader’s reputation for motive (i.e., the extent to which followers believe that the leader is motivated to assist them). In most instances, followers have only uncertain levels of knowledge about their leader. This is especially true for a new leader in which followers may only have priors about their leader’s type. Subsequent actions by the leader builds a reputation, which in turn enables followers to update their priors. At the same time, depending on the context, the leader’s interest may or may not coincide with the followers and this is something the followers may know with uncertainty. Where there is a deeply entrenched, long-term set of relations between leaders and followers, then the task for followers is quite simple. Leader reputation may tell all. At the same

time this means that a leader lives or dies by her reputation. The penalty for taking an action inconsistent with a reputation is often far reaching.

Trust and cooperation are critically dependent on the institutional configuration that surrounds it. One institution in particular, leadership, holds considerable promise in encouraging the formation of trust among followers, if followers are first willing to trust in the institution of leadership itself. We now turn to an empirical examination of the relationship between leaders, leadership, and the problem of trust.

An Overview of the Experimental Design

This paper draws from a number of our previous studies (Wilson and Rhodes 1997; Rhodes and Wilson forthcoming; Rhodes 1999a; Rhodes 1999b). Because all of this work shares a common experimental design, a brief overview is provided here. Where appropriate, we provide additional detail in later sections, and the interested reader is encouraged to consult the original studies for a fuller discussion.

All of the experiments reported on in this paper required groups of eight subjects seated at personal computers in a single room. Subjects participated in highly structured decision settings in which instruction and interactions were computer-based and handled by a local-area network. Subjects were recruited from the student bodies of Rice and Texas A&M Universities, as well as the larger Houston community (mostly students of other Houston-area schools, such as the University of Houston). Subjects only were told that the experiment was concerned with decision making, that details about the experiment would be provided when they arrived, and that they would be paid in cash at the experiment's conclusion.

When subjects arrived at the laboratory, they chose their own seats at computer terminals. All terminals were separated by partitions so that subjects did not have a line of sight to one another's screens. Once the requisite number of subjects appeared, they went through an instructional introduction at their own computer terminal that consisted of text, extended examples,

and a series of manipulation checks. The instructions and examples were self-paced and subjects were encouraged to ask questions prior to the beginning of the experiment.

All subjects participated in either 1) a single-shot coordination game, 2) a repeated coordination game, 3) a single-shot, discrete public good game with a provision point and binary contributions (a variant of a Voluntary Contribution Mechanism with a provisioning point. Below, we refer to these games as VCM games), or 4) a repeated public good game of the same type. Regardless of the game in which they participated, each subject was randomly assigned a letter-based identity at the outset of the initial round (which consisted of letters A, B, etc.) In single-shot games, each subject was randomly assigned a new identity at the beginning of each decision. In the repeated games subjects retained their identities throughout the course of the entire game. Table 1 gives a general overview of the experimental manipulations.

Table 1
Main Conditions for the Experiment

Type of Game	Group Size	New ID Each Period?	Group Changed Each Period?	Colors Changed Each Period?
Single-Stage Coordination	4	Yes	Yes	Yes
Repeated Coordination	4	No	No	Yes
Single-Stage VCM	8	Yes	No	Yes
Repeated VCM	8	No	No	Yes

All coordination games were played in two groups of four people. During the single-stage coordination games, individuals were randomly reshuffled into two groups of four at the outset of each decision. During the repeated coordination games, group composition was randomly determined at the beginning of the game, and then held constant for the duration of the experiment.

In the public good games, eight-person groups were used. In the single-stage public good games, subjects' letter-IDs were randomized at the beginning of each decision and group composition remained the same. For the repeated public good games, letter-IDs were held constant after an initial randomized assignment.

Across all the experiments, the general task for a subject was to select a color. Associated with each color was a row payoff indicating that the outcome was contingent on the color choice of one or more other subjects. Throughout the experiments, six different colors were used: brown, gray, orange, green, yellow and white.¹ Colors were chosen to represent the available choices because they consist of neutral labels without any apparent ordering effects (important for the coordination games) or inherent meaning (important in the public good games). In the coordination games, colors represent the arbitrary choice set of a coordination problem- what matters is not the particular color chosen, but how many subjects choose the same color. If letters or numbers were used, subjects might employ norms of selection based on ordering (e.g., A before B, 1 before 2).

In the public games, colors carry no inherent meaning or value. "Contributing" and "shirking," or investing in a "public" or "private" pool, are labels with ethical and moral overtones, and informing subjects of the true nature of their choice can affect the criteria by which they make that choice. While norms of responsibility or group obligation are very important, they are not the subject of inquiry here. In the public good games, therefore, one color was randomly assigned to represent the choice of contributing toward the public good, and the other was randomly assigned to represent shirking in each round.

At the outset of each decision, two (VCM) or three (coordination) colors were randomly chosen from the pool of six and presented to the subjects as the choices they had for that round. In addition, these colors were randomly ordered at each subject's terminal. Subjects were told that they would see the same colors, but were cautioned that the order of those colors might appear differently on the terminals of other group members. For example, while {yellow, green, brown} might be the order from top row to bottom row for one subject, another might see {green, yellow, brown}. This was done to control for positional norms that might develop in the course of the

¹¹ Initially the set of colors included red and blue. Pre-tests indicated that subjects settled around these two colors. Brown and beige were substituted. Pre-tests showed that beige probably because it was an odd color given the others, and served as a natural focal point. Subjects settled on were robust to providing a natural, uncontrolled focal point.

experiment (e.g., always pick the middle row). Subjects were told that it was the color that mattered, not position.

Subjects made their choices privately and all choices were revealed simultaneously. Once everyone in the group had made a selection, subjects were informed of what the others had chosen and what they had earned for that decision (as well as what they had accumulated to that point). Additionally, the appropriate cell in the payoff matrix was highlighted to reinforce the result. After all subjects had reviewed this information, subjects moved on to the next decision.

For each decision they made, subjects accumulated points (for the coordination games) or “experimental francs” (for the public good games). In the coordination games, points were converted into dollars using a Roth-Malouf procedure (Berg et al 1986), while in the VCM games, experimental francs were converted into dollars at a fixed rate, which was announced to subjects prior to the first round.

Experimental sessions varied considerably in length, ranging from 45-120 minutes. Some sessions consisted of only one game, whereas others consisted of two or three games (e.g., a single-stage coordination game followed by a repeated VCM game). If more than one game was presented, the order of the games was randomly determined. Average earnings for the experiment varied greatly depending on whether the experimental session was a “short” session (i.e. one game) or a “longer” session (i.e., more than one game). In general, earnings averaged around \$7-10 per hour.

The “macro” design of these studies may be thought of as a 2x2x4 incomplete factorial design (see Table 2). The first factor relates to the context and accordingly, all subjects participated in either a coordination or public good game. The second factor builds on conjectures about the importance of repetition. Subjects participated in either single-shot games or repeated games.

Table 2
Experimental Design
(Number of Trials in Each Cell)

	Coordination		VCM	
	One Shot	Repeated	One Shot	Repeated
No Leader	134	40	38	69
Good Leader p=1.0	118	146	50	161
Good Leader p=.85	--	214	--	138
Good Leader p=.50	--	130	--	115

The third factor explores the dynamic of leadership. Accordingly, all decisions were made without a leader or with a leader under one of three conditions of certainty. Under the no-leader condition, the group was made up of four, symmetric (i.e., non-hierarchical with no role or task distinctions) subjects in the coordination games and eight symmetric subjects in the public good games. Each individual made a private color choice and all choices were revealed after everyone had chosen their colors. In this setting, a subject's payoff was dependent on his or her own choice as well as the choices of others.

Under the remaining leadership conditions, one subject from the group was randomly selected to be the leader. That member was referred to as the "monitor" -- a more neutral term than "leader," and the followers were referred to as "participants." Irrespective of type, the leader began the game by sending a private signal to each follower suggesting a particular color. The signal was nothing more than a suggested color and was sent via the computer (e.g., "The monitor suggests that you choose orange"). Once followers observed the leader's signal, they made their own choice, and were always free to accept the leader's advice or ignore it. Followers were paid on the basis of their own choice as well as the choices of other followers. In the coordination setting, a leader's payoff was a function of what the followers chose as well as what she selected. In the public good setting, leader earnings only were a function of what the followers chose.

Subjects with a leader were given a specific probability of having a particular "type" of leader in the experiment. One motive-type of leader was referred to as a "Type W" monitor. Type

W monitors were “good leaders”: their interests were congruent with those of the followers (e.g., in the coordination experiments, good leaders were motivated to help the followers coordinate. In the public good experiments, good leaders were motivated to solicit an efficient number of contributions). The second motive-type of leader was called a “Type Z” monitor (i.e., “bad leaders”). The interests of Type Z monitors diverged from the followers, and Type Z monitors always had an incentive to send misleading signals to subjects (e.g., bad leaders in the coordination games received more points when the followers did poorly, while bad leaders in the public good games were motivated to solicit an inefficient number of contributions).

The experiment manipulated the probability with which subjects drew a Type W leader. We used either a *Perfect* probability of 1.0, a *High* probability of .85 or a *Low* probability of .50. Below, we refer to a leadership condition with a perfect probability of a good a leader as a “certain-leader condition”. Conditions with a leader played under high or low probabilities are referred to as “uncertain-leader conditions.” Leaders always knew their motive-type and were provided with a single payoff matrix reflecting that type. If the probability of a good leader was either .85 or .50 subjects were given two payoff matrices, one assigned to each motive-type of leader.

In the certain-leader conditions, leaders were always credible: the expected value of obeying the leader’s suggestion was greater than the expected value of ignoring it in equilibrium. In the uncertain-leader conditions, the credibility of a leader’s signal was a function of the payoffs displayed in the matrix as well as the probability of having a Type W monitor. In the single-stage experiments, a leader’s motive-type was revealed following the joint decision of all followers in the group. Subjects were then advanced to a new condition. In the repeated experiments, a leader’s motive-type was only unknown in the first round. In round one, therefore, a leader’s credibility was again a function of the payoffs and probability of having a good leader. After round one, a leader’s motive-type was perfectly revealed: good leaders were always credible, and bad leaders were never credible.

Analysis of Results

The Potential of Leadership

We begin with the impact of leadership in a single-stage coordination problem. The no-leader condition involves four subjects choosing among three colors in a series of one-shot games. The more subjects who chose the same color, the more points a subject earned.

The certain-leader condition introduces a “good” leader (i.e., Type W monitor) with certainty. In each trial, a randomly selected leader sends a private, “cheap talk” signal to each follower suggesting a color that might be selected, and no follower knows the signal received by others. The signal constitutes “cheap-talk” in that no individual is bound to follow the leader’s suggestion. After a signal is observed, followers and the leader choose a color. Because leaders’ incentives match those of followers, everyone has an incentive to pick the same color (and the leader has an incentive to suggest the same color to each follower). Leaders, therefore, are endowed with a perfect reputation for motive: followers believe that the leader is motivated to assist them.

We begin with the results from the no-leader condition. Because subjects privately and simultaneously chose colors and were not allowed to communicate with one another, we predicted that the level of trust and cooperation would be low, and as Table 3 shows, this proved to be the case. In 134 trials subjects managed to choose the same color only 19.4 percent of the time. Although subjects did slightly better than what was expected under a mixed strategy ($\chi^2 = 8.51$, $df=2$, $p<.05$), it is clear that absent any institutional corrective, subjects found cooperation to be elusive.

Table 3
Single-Stage Coordination Experiments
Level of Coordination for Four-Person Pure Coordination Game.

Number of Players in Group Choosing Same Color	Expected Percentage Under Null Ho	Observed Percentage	Observed Group Frequency
2	44.9	45.5	61
3	44.2	35.1	47
4	10.9	19.4	26
Totals	100	100	134

Given the apparent difficulty of this social dilemma, does introducing a leader, one who is only endowed with cheap-talk (if credible) signals, increase the level of cooperation? Consider first the behavior of leaders. In equilibrium, leaders should suggest the same color to each follower, and in large measure, this is exactly what the subject-monitors did. Leaders sent all three followers the same color signal 95.8 percent of the time. In response, followers wholeheartedly accept their leader, showing significant levels of trust in the exogenously imposed institution. Confident that the leader was motivated to assist them, followers heeded the leader's counsel at a rate of 92.4 percent. The overall result is that the introduction of a leader produces a dramatic improvement in the level of coordination from 19.4 to 79.7 percent (see Table 4).

Table 4
Single-Stage Coordination Experiments
Levels of Follower Coordination under Certain-Leader Condition

Number of Followers in Group Choosing Same Color	Expected Under Mixed Strategy	Observed Percentage	Observed Group Frequency
1	66.6	3.4	4
2	29.7	16.9	20
3	3.7	79.7	94
Totals	100	100	118

In the single stage coordination experiments, we see that absent an ameliorative institution such as leadership, cooperation is likely to be low. The introduction of a "good" leader produces

dramatic increases in cooperation. Where a leader has interests that match those of followers, the matter of trusting the leader's credible signal is trivial. Leadership proved effective in this context because leaders were motivated to assist followers, and equally important, followers knew that leaders were motivated to assist them. The leader's reputation for motive was perfect, and the result was a high level of cooperation.

As we will see, however, the impact of leadership changes with the underlying context.

Changing the Context to a Single-Stage Public Good Game

The no-leader condition of the single-stage public good experiments involved individuals playing a single-shot, eight-person, discrete public good game with a threshold. As in all such games, subjects have an incentive to "contribute" only if they believe that their contribution is critical to the provision of the good. The source of beliefs, whether those beliefs are accurate, and whether others share the same beliefs all serve to undermine the play of a strategy dependent on making a critical contribution. As such, another equilibrium to the game suggests that subjects should choose the "shirk" color or mix their pure strategies.

The certain-leader condition introduces a "good" leader with a perfect reputation for motive. Much like the coordination experiments, one subject was randomly selected as the monitor at the outset of the trial. Monitors received a "big" payoff if the good was provided (i.e., if five or more followers to choose to contribute) and a "small" payoff if the good was not provided (i.e., if less than five contribute). As shown in Rhodes (1998a), in equilibrium, the leader is expected to send five (the threshold value) contribute signals to the followers in each round. This enables the followers to solve the coordination problem that is inherent in reaching the threshold by communicating who is critical to the good's provision and who is not. Followers are fully informed that the leader is motivated to assist them in provisioning the good (i.e., followers can view the leader's payoffs for each outcome). The leader's signal is credible in these games in the sense that the expected value of following the leader's signal is higher than the expected value of ignoring it in equilibrium. Note, therefore, that the "institution" leadership is identical to the

single-stage coordination experiments discussed above: a known, “good” leader who is motivated to assist the followers sends cheap-talk signals. The only distinction is that the underlying social dilemma has changed from a coordination game to a public good game.

We begin again with the results of the no-leader condition. In only two trials, less than six percent of the total, did subjects manage to successfully provision the public good. On the other hand, subjects did not adhere to the pure strategy Nash equilibrium in which no one in the group contributes: only one trial produced no contributors. Instead, in most rounds, few subjects contribute, with 91 percent of trials landing somewhere between successful provision and complete shirking (see Table 5). Clearly, absent some ameliorative institutional mechanism, provision of the good is extremely unlikely.

Table 5
Single-Stage Public Good Experiments
No-Leader and Certain-Leader Conditions Compared:
Contribution Levels With and Without a Leader

Number of Contributors in Group	0	1	2	3	4	5	6	7	8
Observed # of Group Decisions - No-leader	1	11	14	6	4	2	0	0	0
(% in Parentheses)	(2.6)	(28.9)	(36.8)	(15.8)	(10.5)	(5.3)	(0)	(0)	(0)
Observed # of Group Decisions - Certain-Leader	4	11	15	11	4	5	0	0	NA
(% in Parentheses)	(8.0)	(22.0)	(30.0)	(22.0)	(8.0)	(10.0)	(0)	(0)	NA
Chi square 3.01 (d.f.=7) (p<.75)									

Unlike the single-stage coordination context, however, the introduction of a leader does little to stem the tide of shirking. Indeed, subjects provisioned the good in only 10 percent of the trials under leadership. Although the percentage of trials in which the good is provided is nearly double that produced without a leader (10.0 versus 5.3 percent), the impact of leadership in this context is much lower than its impact in the coordination experiments ($\chi^2 = 2.45$, d.f.=1, $p < .15$).

In part, the failure of leaders to increase levels of trust and cooperation among followers stems from the fact that leaders were less consistent in adopting an equilibrium signaling strategy

than were leaders in the single-stage coordination experiments. As shown in Table 6, leaders sent five contribute signals (the threshold) and two shirk signals in only 18 percent of the trials. A much more common approach was to try to convince all seven followers to contribute. Fully 44 percent of trials had leaders who tried this approach. Even more surprising is that 30 percent of the leaders chose to send too few signals.

Table 6
Single-Stage Public Good Experiments
Certain-Leader Condition
Total Number of Contribution Signals sent to Group

Number of Contribute Signals sent to Group	0	1	2	3	4	5	6	7
Observed # of Group Decisions	4	1	1	2	7	9	4	22
(% in Parentheses)	(8.0)	(2.0)	(2.0)	(4.0)	(14.0)	(18.0)	(8.0)	(44.0)

Regardless of which strategy leaders adopted, followers were much more hesitant to heed their leaders' advice than they were in the coordination experiments. As shown in Table 7, followers were somewhat wary of leaders in this context, and the rate of following the leader was quite low. The leader was perfectly obeyed in only one trial (two percent). In 68 percent of the trials, too few (i.e., less than the threshold of five) followers were willing to obey the leader to provision the good even if asked.

Table 7
Single-Stage Public Good Experiments
Certain-Leader Condition
Group Congruence with Leadership Signals

Number of Followers in Group Obeying Leader's Signal	0	1	2	3	4	5	6	7
Observed Frequencies	0	5	7	15	7	12	3	1
(% in Parentheses)	(.0)	(10.0)	(14.0)	(30.0)	(14.0)	(24.0)	(6.0)	(2.0)

It would be a mistake, however, to infer from these results that leaders were completely ineffectual, or that they were not trusted at all. Subjects proved to be sensitive to leadership signals and were significantly more likely to chip in when asked to do so by their leader. Followers who were asked to contribute did so at an average rate of .39. Although still quite low, it is a significant

improvement over the probability of chipping in when not asked to contribute: .16 ($T = -4.76$ d.f.=206.6 $p < .0001$). A leader's solicitation in these games does increase the likelihood of chipping in, but it is insufficient to lead to provision of the public good. Leaders, then, have an impact in these games, but it is too small to bring about substantial cooperation.

A comparison of the single-stage coordination and public good results reveals an important insight into the problem of trust in general, and more specifically, the prospects for any institution designed to increase it. In short, the dynamic of leadership, as an institutional corrective, is extremely dependent on the context in which it occurs. A shift in the social dilemma (here, switching from a coordination to a step-level public good problem) results in radically different levels of trust and cooperation despite the fact that the leadership mechanism remains the same. Followers accept the institution of leadership in one context, but are hesitant to offer their trust in another. In the threshold VCM game, unlike the pure coordination game, followers have powerful incentives to shirk if they cannot trust one another or if they cannot trust the leader. If followers believe that enough others will shirk so that the threshold is not reached, then there is no reason to contribute. At the same time, if followers do not trust the leader's signal, this too undermines the likelihood of contribution. Unless followers can trust others and the leader, their own incentives to shirk swamp any interest in provisioning the public good. The impact of leadership, and many other institutional fixes we suspect, is conditional on the context in which it occurs.

The Impact on a Leader's Reputation for Competence on Cooperation


Our understanding of the impact of leadership is further complicated by the reputations, motives, and behavior of leaders themselves. Ostrom (1980) has argued that institutions are simultaneously "artisanship and artifact," consisting of not only a set of rules and positions, but also the behavior, values, and history of those within them (also see Ostrom, 1986 and Ostrom, Gardner, and Walker, 1994). Our research to date supports this view, and to show why, we present some results from the repeated coordination and public good experiments.

The design of the repeated experiments is similar to the single-stage experiments discussed above, although with a number of important exceptions. First, subjects maintain the same letter-identity, group, and role (leader or follower) for the duration of the experiment. Second, rather than participate in a series of one-shot decisions, subjects engage in a multiple-round game. Public good games consisted of 23 rounds, whereas coordination games were either finitely repeated (five or 11 rounds) or indefinitely repeated (with the probability of any round being the last either .50, .20, or .05. For details, see Rhodes and Wilson forthcoming and Rhodes 1999b). Third, subjects participate in the same condition for the duration of the game (e.g., a repeated coordination game under the certain-leader condition). Finally, in the uncertain-leader conditions, followers are only uncertain of their leader's motive-type in the first round of play. After the first round, the leader's motive-type is fully revealed, and followers know with certainty if they are being led by a good or bad leader.

Repeated VCM.

We begin with the repeated public good games. In the no-leader condition, eight subjects simultaneously and privately decide to contribute to or withhold their contributions from a public good across 23 rounds. Figure 1 presents the average number of contributions by round. Overall, the average number of contributions per round in this condition is only 2.07 (out of possible total of eight), and successful provision of the good is extremely rare. In only one experiment did subjects successfully provision the good, and the only in the first round. Much like the single-stage experiments, then, groups exhibited little cooperation from the outset and contributions declined over time.

Figure 1
Repeated Public Good Experiments
No-Leader Condition
Average Number of Contributions by Round (standard errors in brackets)



Like the single-stage experiments, we are interested in the overall impact of leadership on trust and cooperation. Additionally, because these games are repeated, we can examine the impact of the leader's revealed reputation for *competence* (the degree to which followers believe that the leader is capable of assisting them). In the single-stage experiments, followers could not adjust their assessment of the trustworthiness of the leader, because the leader was selected anew for each trial. In the repeated experiments, however, followers can react to the leader's revealed reputation for competence and thereby assess her trustworthiness.


One of the key problems for followers is that they can only observe their own signal from the leader, and following the group decision, the number of followers who contributed. As noted above, competent "good" leaders should send five contribute signals in every round. If they deviate from this strategy, sending either too few or too many contribute signals, contributions from followers should decline in response, and followers can draw inferences about their leader's

competence level by observing the total number of contributions in each passing round (see Rhodes 1998a for a fuller discussion).

Empirically, we can gain insight into the role of a leader's reputation for competence by comparing the behavior of two groups. Of the seven leaders who participated in the certain-leader condition, two stood apart from the rest. In the first five rounds, these leaders displayed an unusually high level of "competence," sending five or more contribute signals (the threshold value) per round four or more times. The remaining five leaders sent five or more contribute signals per round three times or less in the first five rounds. A comparison of the signaling strategies of competent and incompetent leaders reveals that these early differences continued over the duration of the experiment.

Figure 2 presents the average number of contribute signals sent by these "competent" and "incompetent" leaders in each round. Competent leaders, defined only by their behavior in the first five rounds, sent an average of 6.41 contribute signals per round over the duration of the experiment (with a standard deviation of 1.38). Incompetent leaders sent an average of 4.45 contribute signals per round with a standard deviation of 2.31, a significant difference ($T = -6.63$, $d.f. = 135.7$, $p < .0001$).

Figure 2
Repeated Public Good Experiments
Certain-Leader Condition
Average Number of Contribute Signals by Competent and Incompetent Leaders




In response, followers recognized these differences and were quick to respond to a leader's revealed reputation for competence, trusting competent leaders at higher rates than incompetent leaders. Followers obeyed contribute signals from a competent leader with a .48 probability, while the same signal from an incompetent leader had only a .30 probability of being obeyed ($T = -.14$, d.f. = 805, $p < .0001$). Note that these differences extend over the duration of the experiment, even though competent and incompetent leaders are identified by the monitor's behavior in only the first five rounds. Subjects proved remarkably sensitive to a leader's revealed reputation for competence.

Moreover, contribute signals from leaders with a strong reputation for competence were more influential at encouraging contributions than contribute signals from incompetent leaders were. Competent leaders could increase the probability of a contribution by as much as .15 by sending the follower a contribute signal. Incompetent leaders could increase the probability of contributing by only .08.

Consequently, competent leaders were much more successful in soliciting contributions than incompetent leaders were. Separating the rate of contribution by early-leader competence, figure 3 shows that groups with competent leaders displayed an average contribution rate of 3.28 per round, whereas groups without such leaders could muster only 1.91 contributions per round ($T = -5.45$, d.f. = 98.0, $p < .0001$).

Overall, the impact of leaders as an institutional solution to the trust problem depends, in part, on the leader's revealed reputation for competence. In principle, leader signals ought to be credible. However, credibility in this complicated environment depends on followers being able to trust that the leader will send useful signals. This is learned by observing the outcome of each round, what the leader privately signals, and then drawing an inference about the competence of the leader. As a solution to social dilemmas, the introduction of any leader is not enough. Only leaders with a revealed reputation for competence can improve contribution rates, but even here, their impact is insufficient to consistently push the number of contributions over the provision point. Most groups with both types of leaders failed to consistently provision the collective good.


Figure 3
Repeated Public Good Experiments
Certain-Leader Condition
Number of Contributions by Early Leadership Competence



Repeated Coordination.

A different perspective on the leader's reputation for competence can be seen in the repeated coordination experiments. Consider first Figure 4, which presents the average number of color signals sent to the group by initial leader competence. Competent leaders, defined by their signaling strategy in the first two rounds, always suggest a common color to each follower over the entire duration of the experiment. Incompetent leaders, again defined by their behavior and their group's reaction in the first two rounds, send an average of 1.3 colors to the group over the duration of the experiment ($T = -4.06$, d.f. = 65, $p < .001$). Note also that these significant differences in behavior exist exclusively in the first five rounds. After the fifth round, initially competent and incompetent leaders converge, each sending one and only one color to the group.

Figure 4
Repeated Coordination Experiments
Certain-Leader Condition
Average Number of Colors Suggested to Group by Initial Leader Competence




Like the repeated public good experiments, followers respond differently to initially competent and incompetent leaders over the duration of the experiment. If paired with a leader who displayed an initially high level of competence, an average of 2.96 followers obeyed the leader's signal in each round, an almost perfect rate of trust. If paired with an initially incompetent leader, however, an average of only 2.36 followers obeyed a leader's signals ($T = -5.36$, d.f. = 70.2, $p < .001$). Followers, then, clearly responded to the leader's initial reputation for competence, and were less likely to trust leaders with weaker reputations.

In an interesting contrast to the repeated public good experiments, however, followers did not abandon their leaders upon the first sign of trouble. Instead they systematically responded to changes in their leader's behavior and revealed reputation. Figure 5 reports the average number of followers obeying the leader by initial leader competence. Although followers offer competent leaders almost perfect levels of trust throughout the experiment, the dynamic in groups with incompetent leaders is much different. Followers facing an incompetent leader respond by

withdrawing their trust, reaching a low of 1.45 followers obeying the leader by round 2. But as shown in figure 4, incompetent leaders eventually get back on track, and followers, never completely withdrawing their trust, responded in almost lock-step. By round 6, initially incompetent leaders are indistinguishable in behavior from their competent colleagues, and in response, followers obeyed these leaders at perfect rates in round 6 and after.

Figure 5
Repeated Coordination Experiments
Certain-Leader Condition
Average Number Obeying the Leader by Initial Leader Competence



In this context, leaders with strong reputations are more likely to be trusted, but followers are hesitant to abandon their leaders even if they begin with a poor reputation for competence. This stands in marked contrast with the repeated public good experiments, in which followers never fully accept the leader. The implications seem clear. Even where a leader's signals should be credible, trust in an exogenously imposed institution such as leadership is a complex function of the context and the revealed reputation and behavior of individual leaders. Given a favorable context, followers will trust both competent leaders, and will even patiently accept (and forgive) a leader with a revealed reputation for incompetence, provided that the leader improves. In a


favorable context, the benefits of the institution, even an imperfect one, outweigh the expected return of going it alone. The dynamic of trust is much different in an unfavorable context. In the repeated public good games, all leaders are viewed with suspicion, and revealing a proclivity toward incompetence makes a difficult situation irreparably worse.

The Impact on a Leader's Reputation for Motive on Trust

To this point we have been concerned only with "good" leaders whose signals ought to be credible. In such an instance, followers need only be concerned with whether the leader enjoys a reputation for competence. What happens, however, if followers only have probabilistic priors about their leader's motive-type? In the following experimental conditions, followers know they have a leader, but are uncertain of whether the leader is good or bad. Good leaders are as before, but bad leaders have interests that run contrary to the followers. Given our design, in repeated games, after round one, the leader's motive-type is perfectly revealed.

Beginning with the repeated coordination experiments, a leader's "initial" reputation is given as the probability of having a Type W leader (a good leader). If followers were informed that the probability was 0.85, an average of 2.44 subjects obeyed in the first round. If, however, the probability of having a good leader dropped to 0.50, only 1.94 followers obeyed ($T = -2.04$, D.F. = 48, $p < .05$). In short, followers relied on their priors concerning the initial reputation of the leader (see Figure 6). They were more likely to use a leader's signal when it was highly likely that the leader shared the same interests.

Figure 6
Repeated Coordination Experiments
Certain-Leader Condition
Average Number of Followers Coordinating by Initial Leader Competence



Followers responded more strongly once the leader's motive-type was fully revealed. After round one, an average of 2.45 followers obeyed a good leader in a given around, while only 1.76 followers obeyed bad leaders ($T = 5.28$, D.F. = 292, $p < .0001$). Moreover, this disparity widens over time. With each passing round, good leaders receive greater levels of trust, while the relationship between bad leaders and followers sours (see Figure 7). This is borne out by the extent to which groups succeed in coordinating. Figure 8 presents the average number of followers who coordinate under good and bad leaders. While 2.46 followers coordinate on average under good leaders, only 1.96 followers do so under bad leaders ($T = 5.20$, D.F. = 292, $P < .0001$). The disparity widens over time, with the effectiveness of good leaders growing and the effectiveness of bad leaders declining.

Figure 7
Repeated Coordination Experiments
Uncertain-Leader Condition
Round ≥ 2
Average Number Obeying Good and Bad Leaders





Figure 8
Repeated Coordination Experiments
Uncertain-Leader Condition
Round ≥ 2
Average Number of Matches under Good and Bad Leaders




Turning to the public good experiments, followers proved insensitive to the leader's reputation for motive in the first round. Followers who were asked to contribute did so at a rate of .86 when told the probability of a good leader was 0.85. When the probability of a good leader was .50, followers obeyed at essentially the same rate as they did when the probability of a good leader was higher. In the .50 condition, the rate of obeying was .85 ($T = -.06$, D.F. = 53, $p < .96$).

After the first round, however, the leader's motive-type is perfectly revealed to followers, and in subsequent rounds, followers were less willing to trust bad leaders. A follower with a good

leader had a .48 rate of obeying. If the contribute request came from a bad leader, however, followers were less likely to obey: .37 ($T = 3.51$, D.F. = 1,175, $p < .001$). It is clear that subjects are sensitive to the leader's (revealed) reputation for motive. They are significantly less likely to trust bad leaders.

These differences result in a significantly higher rate of contribution under good leaders than under bad leaders. As shown in Figure 9, groups with good leaders produced 2.77 contributions a round, while groups with bad leaders produced 2.16 contributions a round ($T = -2.75$, D.F. = 126.3, $p < .01$). Reputation for both motive and competence, then, affect the level of trust in leaders, irrespective of the context.

Figure 9
Repeated Coordination Experiments
Uncertain-Leader Condition (Rounds 2-23)
Average Number of Contributions by Round



Conclusion

Our findings say a great deal about the nature of trust in different contexts. Many institutional settings (e.g., legislatures, informal and voluntary groups) endow their leaders with few carrots or sticks, forcing them instead to rely on persuasion. The findings presented here suggest that for these leaders, a strong reputation for both motive and competence is perhaps a necessary (although not sufficient) condition for success. Subjects are significantly more likely to obey and contribute when the leader's reputation for both motive and competence is strong. Strong reputations, however, are no guarantee of success. Even leaders with excellent reputations for motive and competence were unable to endow their followers with the ability to consistently provision a collective good. In these experiments, highly constrained leaders proved unable to overcome a challenging social dilemma.

These findings stand in stark contrast to Wilson and Rhodes (1997), who find that leaders have a dramatic impact in a single-stage coordination problem. The fact that similarly endowed leaders are much less effective in a public good problem supports the arguments of Ostrom (1986) and Ostrom, Gardner, and Walker (1994), who argue that institutions are the end-result of a complex interaction of rules, positions, payoffs, and information. Subtle shifts in one institutional feature (here, changing from a coordination to a repeated step-level public good problem) will interact in dramatic ways with others (the level of trust between leader and follower), despite the fact that still other features (a leader who dispenses cheap-talk signals) remain unchanged. No definitive statement on the role of leaders as institutional solutions will be possible until the full complexity of these interactions is assessed and systematically disentangled.

Our findings highlight the importance of reciprocal relations between followers and highlight the importance of reputation for leaders. If leaders are going to be effective two things must happen. First, leaders must convince followers that leaders can be trusted. This means that leaders have to be credible and they have to build a strong reputation. Where a leader's reputation for trust is strong, then in most circumstances, it will be credible. However, where the leader has

little reputation, then credibility is crucial. As we have illustrated above, the leader's perceived competence is a powerful signal of credibility. New leaders cannot afford mistakes at the outset because it is difficult to get followers on an equilibrium path if they have fallen off at the beginning.

Second, leaders must convince followers to trust one another. This is done through building an environment that encourages reciprocity among followers. It is not enough to have a leader; that leader has to be active in getting followers to trust one another and soothe followers when mistakes happen or trust is misplaced. An important role for leaders involves providing a model for others. Equally important is the leader's monitoring role when providing a line of communication between followers. By opening such conduits of communication a leader can support norms of trust and reciprocity, even when followers fail to trust or reciprocate. As such a leader can help instill trust where trust is lacking and can help sustain the norm when it is threatened.

Our results also point to the importance of context. In our discussion the context manipulates the provisioning of a good and the incentives for followers in provisioning that good. In the simple coordination problem, followers have identical incentives for coordinated action. In the public goods problem, followers have individual incentives that differ sharply from group incentives. As this context changes, the ability of a leader to provide a coordinating role changes markedly.

Few decision making institutions employ leadership positions that are as anemic as those designed into these experiments. We consciously chose to limit the capacity of leadership in order to examine the ways in which leader credibility and reputation unfold. Most institutions, however, grant additional powers to a leader. Additional monitoring powers are normally granted leaders and leaders are occasionally given sanctions and rewards as a way of stimulating followership (although often these are minimal). We conjecture that sanctioning and rewarding mechanisms are valuable leadership tools. We doubt that such tools would be used all of the time. Instead they

would be used to get followers on a beneficial equilibrium path and primarily used at the outset. Moreover, sanctioning tools likely would be used sparingly as a leader builds a reputation.

Leadership, like cooperative behavior, exists in a complex nexus of reputation, trust, reciprocity and a contextual and institutional overlay. Disentangling the effects of these relationships and the incentives created by different contexts and different institutions for strategic choice is difficult.

We began by pondering the problems faced by a set of U.S. national leaders. Part of our claim was that President Clinton, ex-Speaker Gingrich and ex-Speaker-to-be Livingston all faced problems of trust with respect to their colleagues. There is no doubt that each had worked hard to build trust and reciprocity with members of Congress. However, it was clear from the events in November and December 1998 that any stockpile of trust held by those leaders had quickly dissipated. Although each held (or anticipated holding) a powerful leadership position, that position was insufficient to guarantee the backing and support of others. Leadership is more than holding a position. It requires instilling in followers a sense of trust that stems from a reputation. Reputations are fragile things and can quickly be destroyed. Perhaps this is the political lesson to be drawn from both our experiments and the malaise in leadership at the national level.

Bibliography

- Alchian, George A., and Harold Demsetz. 1972. "Production, Information Costs, and Economic Organization." *American Economic Review*. 62: 777-795.
- Alt, James E., Brian D. Humes, and Randall L. Calvert. 1988. "Reputation and Hegemonic Stability: A Game-Theoretic Analysis." *American Political Science Review*. 82: 445-466.
- Bendor, Jonathan, and Dilip Mookherjee. 1987. "Institutional Structure and the Logic of Ongoing Action." *American Political Science Review*. 81: 129-154.
- Berg, Joyce E., Lane A. Daley, John W. Dickhaut, and John R. O'Brien. 1986. "Controlling Preferences for Lotteries on Units of Experimental Exchange." *Quarterly Journal of Economics*. 101: 281-300.
- Bianco, William T., and Robert H. Bates. 1990. "Cooperation by Design: Leadership, Structure, and Collective Dilemmas." *American Political Science Review*. 84: 133-147.
- Calvert, Randall L. 1987. "Reputation and Legislative Leadership." *Public Choice*. 55: 81-119.
- Calvert, Randall L. 1992. "Leadership and its Basis in Problems of Social Coordination." *International Political Science Review* 13:7-24.
- Camerer, Colin. 1997. "Progress in Behavioral Game Theory" *Journal of Economic Perspectives* 11(4):167-88.
- Cox, Gary W., and Matthew D. McCubbins. 1993. *Legislative Leviathan: Party Government in the House*. Berkeley: University of California Press.
- Dawes, Robyn M. 1980. "Social Dilemmas." *Annual Review of Psychology*. 31: 169-193.
- Eckel, Catherine C. and Philip Grossman. 1996. "Altruism in Anonymous Dictator Games." *Games and Economic Behavior* 16:181-191.
- Eckel, Catherine C. and Philip Grossman. 1998. "Are Women Less Selfish Than Men?: Evidence from Dictator Games." *The Economic Journal*.
- Ellickson, Robert C. 1991. *Order Without Law: How Neighbors Settle Disputes*. Cambridge: Harvard University Press.
- Fehr, Ernst, Georg Kirchsteiger, and Arno Reidl. 1993. "Does Fairness Prevent Market Clearing? An Experimental Investigation." *Quarterly Journal of Economics* 108(2): 437-59.
- Forsythe, R., Horowitz, J. L., Savin, N. E. and Sefton M. (1994). 'Fairness in simple bargaining experiments.' *Games and Economic Behavior*, vol. 6, pp. 347-369.
- Frohlich, Norman., Joe A. Oppenheimer, and Oran R. Young. 1971. *Political Leadership and Collective Goods*. Princeton: Princeton University Press.
- Haraway, Donna Jeanne. 1989. *Primate visions : gender, race, and nature in the world of modern science*. New York : Routledge.
- Hardin, Russell. 1982. *Collective Action*. Baltimore: Johns Hopkins University Press.

- Hoffman, E., McCabe, K., Shachat, K. and Smith, V. (1994). "Preference, property rights and anonymity in bargaining games." *Games and Economic Behavior* 7, pp. 346-80.
- Holmstrom, Bengt. 1982. "Moral Hazard in Teams." *Bell Journal of Economics*. 13: 324-340.
- Jillson, Calvin, and Rick K. Wilson. 1994. *Congressional Dynamics: Structure, Coordination, and Choice in the First American Congress, 1774-1789*. Stanford: Stanford University Press.
- Kiewiet, D. Roderick, and Mathew D. McCubbins. 1991. *The Logic of Delegation: Congressional Parties and the Appropriations Process*. Chicago: The University of Chicago Press.
- Ledyard, John O. 1995. "Public Goods: A Survey of Experimental Research." In *The Handbook of Experimental Economics*, ed. John H. Kagel and Alvin Roth. Princeton: Princeton University Press.
- Miller, Gary J. 1992. *Managerial Dilemmas: The Political Economy of Hierarchy*. New York: Cambridge University Press.
- Oliver, Pamela E.. 1980. "Rewards and Punishments as Selective Incentives for Collective Action: Theoretical Investigations." *American Journal of Sociology*. 85: 1356-1375.
- Olson, Mancur. 1965. *The Logic of Collective Action*. Cambridge: Harvard University Press.
- Ostrom, Elinor. 1986. "An Agenda for the Study of Institutions." *Public Choice*. 48: 3-25.
- Ostrom, Elinor. 1998. "A Behavioral Approach to the Rational Choice Theory of Collective Action; Presidential Address, American Political Science Association, 1997." *American Political Science Review*. 92: 1-22.
- Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. New York: Cambridge University Press.
- Ostrom, Elinor, Roy Gardner, and James Walker. 1994. *Rules, Games, and Common-Pool Resources*. Ann Arbor: The University of Michigan Press.
- Ostrom, Vincent. 1980. "Artisanship and Artifact." *Public Administration Review*. 40 (4): 309-317.
- Rhodes, Carl M. 1999a. "Leaders as Institutional Solutions: Some Experimental Evidence." Unpublished manuscript. Franklin and Marshall College.
- Rhodes, Carl M. 1999b. "The Good, the Bad, and the Incompetent: Leaders and Followers in a Repeated Public Good Problem." Unpublished manuscript. Franklin and Marshall College.
- Rhodes, Carl M. Rick K. Wilson. forthcoming. "The Problem of Common Choice in Symmetric N-Person Coordination Games." Forthcoming. *Handbook of Experimental Economic Results*. Charles R. Plott and Vernon Smith (eds.). Berlin: Elsevier Press.
- Rohde, David W. 1991. *Parties and Leaders in the Postreform House*. Chicago: University of Chicago Press.
- Rohde, David W., and Kenneth A. Shepsle. 1987. "Leaders and Followers in the House of Representatives: Reflections on Woodrow Wilson's Congressional Government." *Congress and the Presidency*. 14: 111-34.

- Salisbury, Robert H. 1969. "An Exchange Theory of Interest Groups." *Midwest Journal of Political Science*. 13: 1-32.
- Schubert, Glendon. 1991. "Primate Politics." In Glendon Schubert and Roger D. Masters (eds.), *Primate Politics*. Carbondale, IL: Southern Illinois Press, pp. 37-56.
- Sinclair, Barbara. 1992. "The Emergence of Strong Leadership in the 1980s House of Representatives." *The Journal of Politics*. Vol. 54. No. 3. pp. 657-684.
- Sinclair, Barbara. 1995. *Legislatures, Leadership, and Lawmaking: The U.S. House of Representatives in the Postreform Era*. Baltimore: Johns Hopkins University Press.
- Van de Kragt, Alphons J. C., John M. Orbell, and Robyn M. Dawes. 1993 "The Minimum Contribution Set as a Solution to the Public Goods Problem." *American Political Science Review*. 77: 112-122.
- de Waal, Frans. 1982. *Chimpanzee Politics*. New York: Harper and Row.
- de Waal, Frans. 1996. *Good Natured: The Origins of Right and Wrong in Humans and Other Animals*. Cambridge: Harvard University Press.
- Wilson, Rick K. 1995. "'Say It Ain't So Uncle Joe': Leadership, Cooperation, and Collective Dilemmas." Presented at the Annual Meeting of the Midwest Political Science Association, Chicago, Illinois.
- Wilson, Rick K., and Carl M. Rhodes. 1997. "Leadership and Credibility in N-Person Coordination Games." *Journal of Conflict Resolution*. 41: 767-791.