JOURNAL OF THEORETICAL POLITICS

Article

Greater than the sum of their parts? When combinations of institutions improve citizens' decisions

Journal of Theoretical Politics
24(1) 90–109
©The Author(s) 2011
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI:10.1177/0951629811423234
jtp.sagepub.com
SAGE

Cheryl Boudreau
University of California, USA

Abstract
Many scholars show that institutions help citizens with their political decisions. However, real-world contexts contain multiple institutions that are imposed together. Thus, I develop a theory and experimental test of the conditions under which combinations of two institutions induce citizens to trust a speaker's statements and make better decisions than when only one institution is present. The theoretical model demonstrates that a second institution typically should not alter a speaker's propensity to make truthful statements, nor subjects' decisions to trust those statements. The experimental results reveal important departures from such rational behavior. Specifically, a second institution makes subjects more willing to trust the already mostly truthful statements they receive, which enables them to make better decisions than when only one institution is imposed upon the speaker. These findings suggest lessons about the conditions under which institutions can increase trust and improve decision making in political, legal, and economic contexts.

Keywords
decision making; experiment; institution; trust

1. Introduction

When do institutions help citizens to improve their decisions? Decades of political science research shows that American citizens are ignorant of basic facts about politics and have little interest in day-to-day political matters. Thus, when making political decisions, citizens often lack factual knowledge about candidates and policies and must base their decisions upon information that substitutes for such political knowledge. Although there are many substitutes that uninformed citizens may use (e.g. party labels, polls, candidates' appearances), a common practice is to rely on the statements of others (Boudreau,

Corresponding author:
Cheryl Boudreau, Assistant Professor, Department of Political Science, University of California, Davis, One Shields Avenue, Davis, CA 95616, USA
Email: clboudreau@ucdavis.edu

2009a,b; Downs, 1957; Druckman, 2001a,b; Lupia, 1994; Lupia and McCubbins, 1998; Popkin, 1991; Sniderman et al., 1991). Indeed, when choosing among candidates for office, uninformed citizens often rely upon the statements of politicians, endorsers, or the candidates themselves. Similarly, citizens may use the endorsements of interest groups to help them decide whether to vote for or against particular ballot propositions.

Because uninformed citizens often rely upon the statements of others (including people they do not know personally), many scholars study the conditions under which citizens trust the statements of individuals personally unknown to them. For example, several scholars suggest that a speaker's party identification conveys whether his or her interests are aligned with those of citizens and, therefore, whether citizens should trust the speaker (Druckman, 2001b; Popkin, 1991; Sniderman et al., 1991). Others note that a speaker's race and gender may indicate whether he or she shares common interests with citizens (Iyengar et al., 1997; McDermott, 1998). Still others demonstrate that when institutions are imposed upon a speaker, they can provide citizens with signals about whether they should trust a speaker's statements (Boudreau, 2009a; Lupia and McCubbins, 1998). Specifically, these scholars demonstrate that when a speaker is subject to an institution (such as a sufficiently large penalty for lying or threat of third-party verification), citizens trust the speaker's statements and improve their decisions.

In this paper, I extend existing research on the effects that institutions have on trust and citizen decision making (Boudreau, 2009a,b; Lupia and McCubbins, 1998; see also Berggren, 2001; Capra et al., 2009; Gordon and Segura, 1997; Morton and Williams, 1999; Ostrom, 1990). Specifically, I analyze whether and when different *combinations of institutions* induce citizens to trust a speaker's statements and achieve even larger improvements in their decisions. I study combinations of institutions because existing research primarily examines the effects of institutions when they are imposed one at a time. However, speakers in real-world contexts are subject to multiple institutions that are imposed in combination with one another. For example, when uninformed voters rely upon the statements of politicians, the politicians' statements can be verified by an opponent or by the media. These politicians might also suffer a penalty (e.g. a loss of reputation, monetary sanctions) if they lie. Similarly, in trial settings, jurors rely upon the statements of witnesses. These statements can be verified during cross-examination, and witnesses are also subject to penalties for lying (e.g. perjury). By analyzing the effects of one institution at a time, existing research leaves open the questions of how these institutions interact with one another, as well as whether and when these institutions can be combined to produce even larger improvements in citizens' decisions.

I take one step toward answering these questions by conducting laboratory experiments in which subjects receive the statements of a knowledgeable, but untrustworthy, speaker in various institutional contexts. Specifically, I impose *both* a penalty for lying and a threat of verification upon the speaker, thereby combining two different institutions that might induce an otherwise untrustworthy speaker to make truthful statements and might induce subjects to trust the speaker's statements. I also vary the size of the penalty for lying and the chance of verification to create several different combinations of these two institutions. I then derive predictions about the conditions under which combinations of two institutions should induce subjects to make better decisions than when only one institution is present.

The results reveal important differences between rational behavior as predicted by formal theory and subjects' actual behavior. Specifically, my theoretical model demonstrates that a second institution typically should not alter the speaker's propensity to make truthful statements, nor should it alter subjects' decisions to trust those statements, relative to when only one institution is imposed upon the speaker. Empirically, however, I find that a second institution makes subjects more willing to trust the already mostly truthful statements they receive, which induces them to make better decisions than when only one institution is present. Further, the improvements that subjects achieve with two institutions can be larger than when these institutions are imposed separately. Thus, combinations of institutions can be greater than the sum of their parts.

These results suggest a number of lessons about how institutions affect trust and decision making when citizens receive information from someone they do not know personally (as is often the case in political and legal settings). Firstly, the results imply that even relatively small penalties for lying and slim chances of verification can increase trust and improve decision making when appropriately combined with other institutions. Secondly, increasing the probability of verification (perhaps by increasing the quality of lawyers who represent criminal defendants or by increasing the number of political watchdog groups) might reduce the need for large penalties for lying. Thus, by strengthening existing institutions or appropriately combining them with others, we can substantially increase citizens' propensity to trust the statements of others and make informed decisions.

2. How do institutions affect citizens' decisions?

Although there are many substitutes for political knowledge (see, e.g., Kuklinski et al., 2001; Lupia, 1994; Popkin, 1991), scholars often examine whether and when institutions help citizens with their decisions (Berggren, 2001; Boudreau, 2009a; Capra et al., 2009; Gordon and Segura, 1997; Lupia and McCubbins, 1998; Morton and Williams, 1999). For example, in their analysis of 12 European countries, Gordon and Segura (1997) argue that particular electoral institutions and political party systems make the acquisition of political information less costly for citizens, which helps them to become more politically sophisticated. Similarly, Berggren (2001) demonstrates that certain institutional contexts can level the playing field between sophisticated and unsophisticated citizens.

In addition to studying institutions in real-world contexts, scholars use laboratory experiments to assess the effects that particular institutions have on citizens' decisions. Although laboratory experiments necessarily lack the complexities of real-world politics, they are useful for analyzing the effects of institutions. Indeed, laboratory experiments provide scholars with a controlled setting in which they can systematically manipulate the presence or absence of institutions, as well as the strength of particular institutions. Scholars can then directly observe the effects of particular institutions (and institutional changes) without the confounding events that occur in the real world. Precisely because of these advantages, scholars have used laboratory experiments to study a variety of institutions, including voting rules (Guarnaschelli et al., 2000; Morton and Williams, 1999), communication methods (Capra et al., 2009; Dickson et al., 2010; Ostrom et al., 1992), penalties for lying and threats of third-party verification (Boudreau, 2009a; Lupia and McCubbins, 1998), as well as monitoring and contract enforcement (Bohnet et al., 2001; Ostrom, 1990). Taken together, these studies demonstrate that institutions affect citizens' decisions in significant and often positive ways.

Although the experiments described above yield many key insights about the effects of institutions, they primarily study institutions when they are imposed one at a time. Therefore, they do not typically assess the effects that combinations of institutions have on citizens' decisions, nor do they systematically vary the strength of those institutions. While studying individual institutions is necessary and important, it is also important to assess how institutions of different strengths interact with one another, as well as whether and when particular institutions can be combined to produce even larger improvements in decision making. To this end, I adapt the experimental designs of Lupia and McCubbins (1998) and Boudreau (2009a,b) and assess the conditions under which combinations of two institutions (namely, a penalty for lying and a threat of third-party verification) induce citizens to trust a speaker's statements and make better decisions than they make when only one of these institutions is present. Such an assessment is important, because citizens often have opportunities to learn from others in real-world contexts, where multiple institutions are imposed upon speakers. Thus, the results may suggest lessons about how to design and combine institutions to improve the decisions that citizens make.

3. Research design

In order to analyze the conditions under which combinations of institutions help citizens to improve their decisions, I conduct laboratory experiments in which subjects are randomly assigned to treatment and control groups. In each group, subjects answer a series of binary choice math problems. Specifically, subjects choose whether answer 'a' or answer 'b' is the correct answer to a problem, or they may leave the problem blank. The math problems are drawn from a SAT math test and consist of different types of problems and levels of difficulty. I tell subjects in the treatment and control groups that they have 60 seconds to answer each math problem and that they will earn 50 cents for each problem they answer correctly, lose 50 cents for each problem they answer incorrectly,[1] and neither earn nor lose 50 cents if they leave a problem blank. In this way, the amount of money subjects earn reflects the quality of their decisions in the treatment and control groups.

I ask subjects to make choices about math problems (instead of asking them to vote for fictional candidates or policies) for several reasons. Firstly, this task provides a straightforward way of identifying correct decisions and assessing whether and when institutions induce an improvement in decision making. Stated differently, although it is often difficult to identify when citizens have chosen the 'correct' candidate or policy,[2] it is easy to tell when they have chosen the correct answer to a math problem. Secondly, as discussed by Boudreau (2009a,b), even though math problems do not *look* like political decisions on the surface, they capture key elements of the psychological processes used by voters in real-world political contexts. Thus, these experiments have a great deal of psychological realism and can tell us much about how citizens in the real world make political choices (Aronson et al., 1998).[3] Thirdly, even though math problems are more abstract than many tasks used in political science experiments, such abstraction confers a number of benefits. Specifically, abstract experimental settings often provide greater control than more contextually rich settings, where it may be difficult to isolate which of the many features of those settings influenced subjects' decisions. Further, abstract game-theoretic settings (like the one used in these experiments) enable the calculation of a rational choice baseline to which subjects' actual behavior can be compared (Dickson, 2011).[4]

The difference between the treatment and control groups is that subjects in the control group answer the problems on their own, while subjects in the treatment conditions hear the statements of a speaker in a particular institutional context before they answer the problems. Specifically, before subjects in the treatment conditions answer the math problems, the experimenter randomly selects a subject to act as 'the speaker.' The speaker's role in the experiment is different from that of the other subjects. That is, unlike the other subjects (whose role is to answer the math problems), the speaker is shown the correct answer to each math problem (that is, the speaker is given knowledge about the best choice for subjects) and then makes a statement to the other subjects about the answer to the math problem.[5] After the speaker makes her statement, the other subjects must decide within 60 seconds whether to answer the problem, and if they choose to answer, whether to pick 'a' or 'b.'

In the treatment conditions, both the speaker and the subjects know that the speaker can lie about the correct answer to the math problem or tell the truth. The speaker's ability to lie or tell the truth is designed to resemble Crawford and Sobel's (1982) and Lupia and McCubbins's (1998) models, as well as many real-world political settings. Further, in each treatment condition, the speaker and subjects have conflicting interests; that is, the speaker earns money when subjects make incorrect decisions, but subjects earn money when they make correct decisions. Thus, as in many real-world political settings, the speaker may have an incentive to misrepresent the truth. Although the speaker and subjects have conflicting interests, I vary whether the speaker is subject to a *penalty for lying* (either \$1, \$5, or \$15), a *threat of verification* (either a 30%, 70%, 90%, or 100% chance of verification), or some *combination* of a particular penalty for lying and chance of verification.[6] Thus, in some treatment conditions, there is only one institution imposed upon the speaker, while in other treatment conditions, two institutions are imposed upon the speaker. I then assess the effects that these institutions have on the quality of subjects' decisions, as indicated by the amount of money they earn.

To establish conflicting interests between the speaker and subjects and impose institutions upon the speaker, I manipulate how the speaker and subjects earn money. To establish conflicting interests between the speaker and subjects, I pay subjects 50 cents for each math problem that they answer correctly. The speaker, on the other hand, earns 50 cents for each subject who gets a math problem *wrong*, loses 50 cents for each subject who answers a problem correctly, and neither earns nor loses 50 cents for each subject who leaves a problem blank. For example, if 11 subjects answer a math problem correctly, the subjects earn 50 cents each, and the speaker loses \$5.50 (i.e. 50 cents for each of the 11 subjects who answer correctly).

Simultaneously, I also impose an institution – namely, a penalty for lying – upon the speaker. In the penalty for lying treatment conditions, the speaker and subjects have conflicting interests, but I announce to both the speaker and subjects that the speaker will incur a penalty (of either \$15, \$5, or \$1, depending on the particular treatment condition) if she lies about the correct answer to the math problem. Importantly, the procedures and payoffs used in each treatment condition are common knowledge.

In another treatment condition, I maintain conflicting interests between the speaker and subjects, but instead of imposing a penalty for lying upon the speaker, I verify the speaker's statement with some probability to make sure that it is a true statement before it is read to subjects. In the 100% chance of verification condition, if the speaker chooses to make a false statement about the correct answer to the math problem, then I silently verify the speaker's statement, charge her $2, and replace the false statement by announcing the correct answer to subjects. If the speaker chooses to make a true statement, then I simply announce the speaker's statement to subjects. In the 90% chance of verification condition, I roll a 10-sided die before the speaker makes her statement. If the die lands on 1, 2, 3, 4, 5, 6, 7, 8, or 9, then I silently verify the speaker's statement, charge her $2 if she chooses to make a false statement, and announce the correct answer to subjects.[7] If the die lands on 10, then I announce the answer that the speaker chooses to report, regardless of whether it is true or false. In this way, subjects know that there is a 90% chance that the speaker will be verified, but they do not know whether the speaker has been verified on any particular problem. The 70% and 30% chance of verification conditions proceed in a similar manner.

In order to assess whether and when combinations of institutions help citizens to improve their decisions further, I impose *both* a penalty for lying and a threat of verification upon the speaker. Specifically, I examine seven combinations of these institutions: (1) a 30% chance of verification plus a $1 penalty for lying; (2) a 30% chance of verification plus a $5 penalty for lying; (3) a 70% chance of verification plus a $1 penalty for lying; (4) a 70% chance of verification plus a $5 penalty for lying; (5) a 90% chance of verification plus a $1 penalty for lying; (6) a 100% chance of verification plus a $1 penalty for lying; and (7) a $15 penalty for lying plus a 30% chance of verification. Although there are many different combinations that I could examine, I use these particular combinations for three reasons. Firstly, by combining each chance of verification with a $1 penalty for lying, I am able to assess whether the addition of a $1 penalty has different effects at each level of verification. Secondly, by combining a $5 penalty for lying with a 30% and 70% chance of verification, I am able to examine whether and to what extent the addition of a larger penalty for lying further improves subjects' decisions. Thirdly, by combining a 100% chance of verification with a $1 penalty for lying and a $15 penalty for lying with a 30% chance of verification, I can compare whether there is a difference between combining a large penalty for lying with a small chance of verification versus combining a large chance of verification with a small penalty for lying.

4. Predictions

I now derive predictions about the conditions under which a second institution should increase the amount of money subjects earn, relative to when only one institution is imposed upon the speaker. Following Lupia and McCubbins (1998), I model the interaction between a speaker and one subject. I do this because, although there are multiple subjects in each experiment, the game is actually between one subject and one speaker.[8] Specifically, in the experiments, each subject earns or loses money only for his own decisions. Thus, an individual subject's payoff is not affected by the decisions of other subjects, and from a subject's perspective, it is as though the experiment is only between the individual subject and the speaker. Similarly, because each subject earns money in the same way and cannot interact with other subjects, from the speaker's perspective, it is also a game between an individual subject and the speaker (see Lupia and McCubbins, 1998, for further discussion). Thus, the speaker's payoffs and penalties for lying are calculated on a per subject basis in the analyses that follow.

On each problem in each treatment condition, a subject has two strategies: (1) trust the speaker's statement; or (2) not trust the speaker's statement. If a subject trusts the speaker's statement, he chooses 'a' if he receives the statement 'a' and chooses 'b' if he receives the statement 'b.' If a subject does not trust the speaker's statement, then he answers the problem on his own. Assume that if a subject answers the problem on his own, he has a 60% chance of answering correctly (thereby earning \$0.50) and a 40% chance of answering incorrectly (thereby losing \$0.50).[9] This assumption captures the fact that each subject has preexisting knowledge about how to solve math problems and, therefore, has more than a 50% chance of answering correctly on his own (i.e. does better than randomly choosing 'a' or 'b'). This assumption is also justified empirically by the results from the control group. In the control group, subjects solve the problems on their own and earn, on average, \$0.10 per problem, which is the expected payoff that results from a 60% chance of answering correctly and a 40% chance of answering incorrectly. Thus, a subject's expected payoff of not trusting the speaker's statement and answering the problem on his own is:

$$E_{\text{NotTrust}} = (60\% \times \$0.50) + (40\% \times \$-0.50) = \$0.10 \qquad (1)$$

If a subject trusts the speaker's statement, then his expected payoff depends upon the probability that the speaker's statement is verified, as well as the probability that the speaker lies. Let v equal the probability that the speaker's statement is verified and f equal the probability that the speaker lies. A subject's expected payoff of trusting the speaker's statement can then be calculated as follows in each treatment condition:

$$E_{\text{Trust}} = (v)(\$0.50) + (1-v)[(f)(\$-0.50) + (1-f)(\$0.50)] \qquad (2)$$

As for the speaker, she also has two strategies: (1) tell the truth about the correct answer; or (2) lie about the correct answer. Recall that in each treatment condition the speaker earns \$0.50 if a subject answers incorrectly and loses \$0.50 if a subject answers correctly. Thus, the speaker's expected payoff of telling the truth or lying depends on the probability that the subject trusts the statement she receives (denoted q), as well as the probability of verification (v) and the size of the penalty for lying (denoted p). If the speaker tells the truth, her expected payoff can be calculated as follows for each subject in each treatment condition:

$$E_{\text{Truth}} = (q)(\$-0.50) + (1-q)(40\% \times \$0.50 + 60\% \times \$-0.50) \qquad (3)$$

If the speaker lies, her expected payoff is calculated as follows:

$$\begin{aligned} E_{\text{Lie}} = {} & (v)[(q)(\$-0.50) + (1-q)(40\% \times \$0.50 + 60\% \times \$-0.50) - \$0.18] \\ & + (1-v)[(q)(\$0.50) + (1-q)(40\% \times \$0.50 + 60\% \times \$-0.50)] - p \end{aligned} \qquad (4)$$

I use these equations to determine the equilibrium outcome and payoffs in each treatment condition. These predictions are summarized in Table 1. The full calculations are shown

Table 1. Predicted behavior for the speaker and subjects in each treatment condition.

Treatment condition	Predicted speaker behavior	Predicted subject behavior	Predicted subject payoff
30% verification	Mix True/False	Mix Trust/Not Trust	\$0.10
30% verification + \$1 penalty	Mix True/False	Mix Trust/Not Trust	\$0.10
30% verification + \$5 penalty	Mix True/False	Mix Trust/Not Trust	\$0.10
70% verification	False statements	Trust	\$0.20
70% verification + \$1 penalty	False statements	Trust	\$0.20
70% verification + \$5 penalty	True statements	Trust	\$0.50
90% verification	True statements	Trust	\$0.50
90% verification + \$1 penalty	True statements	Trust	\$0.50
100% verification	True statements	Trust	\$0.50
100% verification + \$1 penalty	True statements	Trust	\$0.50
\$15 penalty	True statements	Trust	\$0.50
\$15 penalty + 30% verification	True statements	Trust	\$0.50

in Online Appendix 2.[10] Firstly, I do not expect to observe a difference in the amount of money subjects earn in the 30% chance of verification condition, the 30% chance of verification plus a \$1 penalty for lying condition, and the 30% chance of verification plus a \$5 penalty for lying condition. The logic behind this prediction can be seen by plugging a 30% chance of verification into Equations (1)–(4). When $v = 0.30$, whether the subject should Trust or Not Trust the speaker's statement depends upon the probability that the speaker lies, f. Likewise, whether the speaker should choose Lie or Truth depends upon the probability that the subject trusts, q. Setting Equations (1) and (2) equal to one another and solving for f reveals that the subject should Not Trust when $f > 0.57$, Trust when $f < 0.57$, and mix between Trust and Not Trust when $f = 0.57$. Setting Equations (3) and (4) equal to one another and solving for q reveals that the Speaker should choose Lie when $q > 0.07$, Truth when $q < 0.07$, and mix between Lie and Truth when $q = 0.07$. Plotting the best responses shows that the equilibrium for the 30% chance of verification treatment condition is one in which the subject Trusts with probability 0.07 and does Not Trust with probability 0.93 and the speaker Lies with probability 0.57 and tells the Truth with probability 0.43. Given these probabilities, the equilibrium payoff for the subject is \$0.10. Using Equations (1)–(4), it is also clear that while the addition of a \$1 or \$5 penalty for lying alters the value of q at which the speaker is willing to mix ($q = 0.2$ with a \$1 penalty and $q = 0.71$ with a \$5 penalty), it does not alter the equilibrium payoff for the subject, which remains approximately \$0.10.

When a 70% chance of verification is imposed upon the speaker, I do not expect to observe a difference in the amount of money subjects earn in the 70% chance of verification condition and in the 70% chance of verification plus a \$1 penalty for lying condition. I do, however, expect to observe a difference between the amount of money subjects earn in these two treatment conditions and the amount of money they earn in the 70% chance of verification plus a \$5 penalty for lying condition. The logic behind the first prediction can be seen by plugging a 70% chance of verification into Equations

(1)–(4). When $v = 0.70$, the expected payoff to the subject of Trust is strictly greater than the expected payoff of Not Trust for all values of $f \in [0, 1]$. Thus, the subject has a dominant strategy to Trust, and $q = 1$. Given that $q = 1$ and $v = 0.70$, the expected payoff to the speaker of Lie is strictly greater than the expected payoff of Truth. The speaker, therefore, has a dominant strategy to lie, and $f = 1$. Thus, the equilibrium for the 70% chance of verification treatment condition is one in which the speaker lies and the subject trusts the statement he receives (because the speaker's lies are verified 70% of the time and replaced with truthful statements). The equilibrium payoff for the subject is, thus, \$0.20. Using Equations (1)–(4), it is also clear that the addition of a \$1 penalty for lying to the 70% chance of verification does not change the speaker's or subject's dominant strategies, nor does it change the equilibrium payoff for the subject. When a \$5 penalty for lying is added to the 70% chance of verification, the subject still has a dominant strategy to Trust, but the speaker now has a dominant strategy to make truthful statements (because of the larger \$5 penalty for lying). Thus, in this treatment condition, the equilibrium payoff for the subject is \$0.50.

When a 90% chance of verification is imposed upon the speaker, I do not expect to observe a difference in the amount of money subjects earn in the 90% chance of verification condition and in the 90% chance of verification plus a \$1 penalty for lying condition. Specifically, when $v = 0.90$, the expected payoff to the subject of Trust is strictly greater than the expected payoff of Not Trust for all values of $f \in [0, 1]$. Thus, the subject has a dominant strategy to Trust, and $q = 1$. Given that $q = 1$ and $v = 0.90$, the expected payoff to the speaker of Truth is strictly greater than the expected payoff of Lie. The speaker, therefore, has a dominant strategy to tell the truth, and $f = 0$. Thus, the equilibrium in the 90% chance of verification treatment condition is one in which the speaker tells the truth and the subject trusts the statement he receives. The equilibrium payoff for the subject is, thus, \$0.50. Using Equations (1)–(4), it is also clear that the addition of a \$1 penalty for lying to the 90% chance of verification does not change the speaker's or subject's dominant strategies, nor does it change the equilibrium payoff for the subject.

When a 100% chance of verification is imposed upon the speaker, I also predict that the addition of a second institution (a \$1 penalty for lying) will not induce further increases in the amount of money the subject earns. The logic behind this prediction is straightforward. When $v = 1$, the subject is certain to receive truthful statements, regardless of whether the speaker lies or tells the truth. The subject should thus base his decision upon the statement received (see also Lupia and McCubbins, 1998). Given that the subject is certain to receive truthful statements, the subject's equilibrium payoff is \$0.50 in both the 100% chance of verification condition and in the 100% chance of verification plus a \$1 penalty for lying condition.

When a \$15 penalty for lying is imposed upon the speaker, I also expect that the addition of a second institution (in this case, a 30% chance of verification) will not induce further increases in the amount of money the subject earns. This prediction is also straightforward. Specifically, the \$15 penalty is large enough to ensure that the speaker has a dominant strategy to tell the truth and that the subject, therefore, has a dominant strategy to base his decisions upon the speaker's statement. Thus, the subject's equilibrium payoff is \$0.50 in both the \$15 penalty for lying condition and the \$15 penalty for lying plus a 30% chance of verification condition.

5. Methodology

To test the above predictions, I conducted laboratory experiments at a large public university. A total of 381 adults who were enrolled in undergraduate classes participated in one of 33 experimental sessions (four control group sessions and 29 treatment group sessions). In each experimental session, subjects were asked to answer 24 math problems. In treatment group sessions, subjects answered these problems under three or four different conditions, producing a within-subjects design for the treatment group. Most treatment group sessions included 12 subjects (11 subjects who answer the math problems plus one speaker), although some sessions contained slightly fewer subjects. In each session, subjects earned an average of $5 for the 24 math problems that they answered, in addition to a $5 show-up payment, up to $3.50 for solving practice problems correctly, and up to $6 for correctly answering quiz questions about the experimental instructions.

When analyzing the quality of subjects' decisions, I regress a dependent variable that reflects the amount of money that each subject earns on each problem on: (1) a dummy variable for each treatment condition (i.e. the *$15 penalty for lying* variable is coded 1 for the $15 penalty for lying condition and 0 otherwise); (2) a *Sophistication* variable that reflects subjects' SAT math scores and, therefore, controls for subjects' levels of sophistication at making this type of decision; and (3) a *Difficulty* variable that controls for the level of difficulty of the problems (higher values indicate harder problems). The omitted category in this model is the control group, and the unit of analysis is subject–problem observations.[9] This analysis allows me to assess whether subjects make significantly better decisions (as evidenced by larger amounts of money earned) when two institutions are imposed upon the speaker versus when each of these institutions is imposed separately. It also shows whether subjects in each treatment condition earn significantly more money than subjects in the control group, who make their decisions on their own.

Because the quality of subjects' decisions depends, in part, upon the truthfulness of the statements they receive, I also analyze the extent to which subjects receive truthful statements in each treatment condition. Specifically, I estimate a regression in which I regress a dummy variable that reflects whether subjects receive a truthful statement on each math problem (coded 1 if the statement is truthful and 0 otherwise) on a dummy variable for each treatment condition and the control variables described above. The unit of analysis in this model is speaker–problem observations.[10] The omitted category is those speakers who had conflicting interests with subjects, but who were not subject to a penalty for lying or a threat of verification.

6. Results

The results show that the addition of a second institution substantially increases the amount of money subjects earn, relative to when only one institution is imposed upon the speaker. Indeed, in nearly every treatment condition where two institutions are imposed upon the speaker, subjects earn significantly more money than when only one of these institutions is imposed. These findings are remarkable because they reveal departures from rational behavior as predicted by my theoretical model. Indeed, in most cases the second institution, theoretically, should not have induced further improvements in subjects' decisions, yet empirically, it did.

Table 2. The effect of each treatment condition on the quality of subjects' decisions.

Independent variables	Dependent variable = money earned per subject per question
$15 penalty for lying	0.247* (0.020)
$5 penalty for lying	0.032 (0.023)
$1 penalty for lying	0.022 (0.015)
100% verification	0.369* (0.027)
90% verification	0.174* (0.026)
70% verification	0.084* (0.019)
30% verification	0.023 (0.018)
30% verification + $1 penalty	0.185* (0.027)
30% verification + $5 penalty	0.197* (0.043)
70% verification + $1 penalty	0.216* (0.026)
70% verification + $5 penalty	0.217* (0.052)
90% verification + $1 penalty	0.326* (0.028)
100% verification + $1 penalty	0.403* (0.027)
$15 penalty + 30% verification	0.314* (0.026)
Difficulty	–0.005* (0.000)
Sophistication	0.000* (0.000)
Constant	–0.339* (0.048)
R-squared	0.193
N = 5046	

Standard errors in parentheses; $*p < 0.05$

Consider the results from the 30% chance of verification conditions in Tables 2 and 4. Recall that subjects' expected payoff should be $0.10 in the 30% chance of verification condition, the 30% chance of verification plus a $1 penalty for lying condition, and the 30% chance of verification plus a $5 penalty for lying condition. Thus, I did not expect to observe a difference in the amount of money subjects earn in these three treatment conditions. Empirically, I find that subjects in the 30% chance of verification condition earn approximately what I expected. That is, they earn, on average, $0.12 per problem. However, when the 30% chance of verification is combined with either a $1 or $5 penalty for lying, subjects earn significantly more money than predicted. As shown in Table 4, subjects in the 30% chance of verification plus a $1 penalty for lying condition earn, on average, $0.28 per problem, and subjects in the 30% chance of verification plus a $5 penalty for lying condition earn, on average, $0.29 per problem. Both of these amounts are significantly greater than the $0.12 per problem that subjects earn in the 30% chance of verification condition. Thus, adding a second institution (be it a $1 or $5 penalty for lying) to a 30% chance of verification significantly improves subjects' decisions even though, theoretically, it should not.

I observe a similar pattern of results in the 70% chance of verification conditions. Theoretically, subjects' expected payoff should be $0.20 in both the 70% chance of verification condition and in the 70% chance of verification plus a $1 penalty for lying condition. Subjects' expected payoff should be $0.50 in the 70% chance of verification plus a $5 penalty for lying condition. Empirically, I find that subjects in the 70% chance of verification condition earn approximately what I expected ($0.18 per problem).

Table 3. Effect of each treatment on the probability that subjects receive a true statement.

Independent variables	Dependent variable = whether subjects receive a true statement	
30% verification	1.200*	(0.300)
30% verification + $1 penalty	1.652*	(0.525)
70% verification	1.092*	(0.271)
70% verification + $1 penalty	1.240*	(0.422)
$1 penalty for lying	0.279	(0.187)
$5 penalty for lying	0.881*	(0.284)
$15 penalty for lying	1.892*	(0.438)
Difficulty	–0.012*	(0.003)
Sophistication	–0.001	(0.001)
Constant	0.341	(0.715)
Log likelihood	–222.55	
N = 459		

Standard errors in parentheses; $*p < 0.05$

Note: subjects always receive true statements in the 100% chance of verification condition, the 100% chance of verification plus a $1 penalty for lying condition, the $15 penalty for lying plus a 30% chance of verification condition, the 90% chance of verification condition, the 90% chance of verification plus a $1 penalty for lying condition, the 70% chance of verification plus a $5 penalty for lying condition, and the 30% chance of verification plus a $5 penalty for lying condition. Thus, these treatment conditions are not included in this model.

However, when the 70% chance of verification is combined with a $1 penalty for lying, subjects earn significantly more money than predicted. As shown in Table 4, subjects in the 70% chance of verification plus a $1 penalty for lying condition earn, on average, $0.31 per problem, which is significantly greater than the $0.18 that subjects earn in the 70% chance of verification condition. Interestingly, subjects earn a similar amount of money in the 70% chance of verification plus a $5 penalty for lying condition, earning $0.32 per problem on average. Thus, I again find that a second institution (namely, a $1 penalty for lying) significantly improves subjects' decisions even though, theoretically, it should not. Although adding a $5 penalty for lying to the 70% chance of verification also improves subjects' decisions, it is not nearly as effective as I expected (i.e. $0.50).

The results from the 90% chance of verification conditions also demonstrate that a second institution further increases the amount of money subjects earn. Although I expected subjects to earn $0.50 in both the 90% chance of verification condition and in the 90% chance of verification plus a $1 penalty for lying condition, I do not observe this empirically. As shown in Table 4, when a 90% chance of verification is imposed upon the speaker, subjects earn, on average, only $0.27 per problem, which is a much smaller amount of money than expected. When a 90% chance of verification is combined with a $1 penalty for lying, subjects achieve a significant increase in the amount of money they earn. Specifically, subjects in the 90% chance of verification plus a $1 penalty for lying condition earn, on average, $0.42 per problem. Although this amount is not as large as the expected $0.50, it is an improvement over the $0.27 subjects earn in the 90% chance of verification condition.

Table 4. Amount of money earned and the probability of receiving a true statement in each treatment condition and in the control group.

Experimental condition	Money earned	Probability of receiving a true statement
30% verification	$0.12 (0.088, 0.148)	0.91 (0.804, 0.973)
30% verification + $1 penalty	**$0.28** (0.231, 0.329)	0.95 (0.819, 0.998)
30% verification + $5 penalty	**$0.29** (0.210, 0.377)	**1.00**
70% verification	$0.18 (0.150, 0.208)	0.90 (0.797, 0.962)
70% verification + $1 penalty	**$0.31** (0.265, 0.354)	0.91 (0.745, 0.986)
70% verification + $5 penalty	**$0.32** (0.216, 0.417)	**1.00**
90% verification	$0.27 (0.217, 0.319)	1.00
90% verification + $1 penalty	**$0.42** (0.373, 0.473)	1.00
$15 penalty	$0.34 (0.309, 0.376)	0.97 (0.904, 0.998)
$15 penalty + 30% verification	**$0.41** (0.362, 0.458)	**1.00**
100% verification	$0.47 (0.415, 0.518)	1.00
100% verification + $1 penalty	$0.50 (0.449, 0.547)	1.00
Control	$0.10 (0.077, 0.116)	0.57 (0.493, 0.653)
$1 penalty	$0.12 (0.095, 0.140)	0.68 (0.578, 0.771)
$5 penalty	$0.13 (0.088, 0.167)	0.85 (0.712, 0.938)

Confidence intervals in parentheses. Boldface indicates that the amount of money earned or the probability of receiving a true statement in a particular 'two-institution' treatment condition is significantly different from the corresponding 'one-institution' treatment condition. Based on the models reported in Tables 2 and 3, with the results drawn from 1000 simulations performed by CLARIFY (Tomz et al., 2003). For each condition, all other treatment variables are set to zero and the control variables are held constant at their median values. The control condition for the 'probability of receiving a true statement' results is those speakers who had conflicting interests with subjects, but who were not subject to a penalty for lying or threat of verification.

The results from the $15 penalty for lying conditions are similar. Although I expected subjects to earn $0.50 in both the $15 penalty for lying condition and in the $15 penalty for lying plus a 30% chance of verification condition, I again do not observe this empirically. As shown in Table 4, when a $15 penalty for lying is imposed upon the speaker, subjects earn, on average, only $0.34 per problem, which is a smaller amount than expected. However, when a $15 penalty for lying is combined with a 30% chance of verification, subjects earn $0.41 per problem on average. Although this amount is not as large as the expected $0.50, it is a significant improvement over the $0.34 that subjects earn in the $15 penalty for lying condition. Thus, I again find that the addition of a second institution (in this case, a 30% chance of verification) significantly improves subjects' decisions, relative to when only one institution is imposed upon the speaker. Indeed, of all the treatment conditions included (and as shown in Figure 1), only the 100% chance of verification conditions yield a pattern of results that are consistent with my predictions.

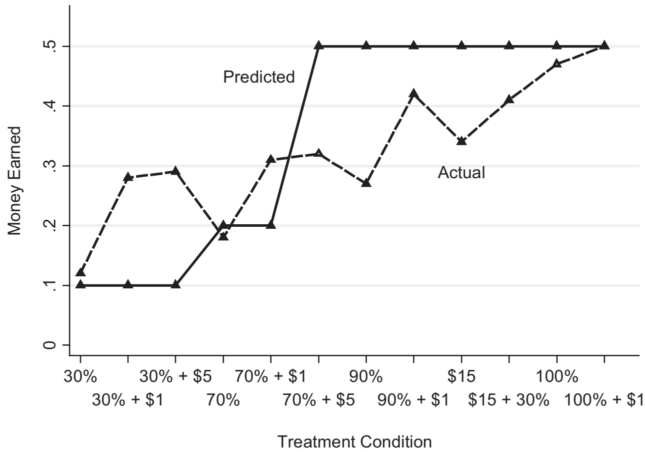

Figure 1. Predicted versus actual money earned by treatment condition.

6.1. Why combinations of institutions further improve decisions

The question the above results leave open is *why* the addition of a second institution has such a consistently positive effect on subjects' decisions, even though, theoretically, it should not. One possibility is that the addition of a second institution increases the speaker's propensity to make truthful statements. If this is the case, then I should observe an increase in the truthfulness of the statements subjects receive in the two-institution conditions, relative to the corresponding one-institution conditions. A second possibility is that the additional institution makes subjects more willing to trust the speaker's statements, even if it does not increase the truthfulness of the statements subjects receive. If this is the case, then I should observe similar amounts of truthful statements in the corresponding one-institution and two-institution conditions, but a difference in subjects' willingness to base their decisions on the statements they receive. A third possibility is that a second institution has both of these effects. If this is the case, then I should observe both an increase in the truthfulness of the statements subjects receive in the two-institution conditions and an increase in subjects' willingness to base their decisions on the statements they receive.

The results in Tables 3 and 4 rule out the first and third possibilities. That is, they show that increases in the amount of money subjects earn in the two-institution conditions are not driven by increases in the truthfulness of the statements they receive. As shown in Table 4, subjects are equally likely to receive truthful statements in the 30% chance of verification condition and in the 30% chance of verification plus a $1 penalty for lying condition.[11] They are also equally likely to receive truthful statements in the 70% chance of verification condition and in the 70% chance of verification plus a $1 penalty for lying condition. Subjects also always receive truthful statements in both the 90% chance of verification condition and in the 90% chance of verification plus a $1 penalty for lying condition and virtually always receive truthful statements in the $15 penalty for lying condition and in the $15 penalty for lying plus a 30% chance of verification condition. Although subjects are equally likely to receive truthful statements,

Table 5. Percentage of problems answered correctly, incorrectly, or left blank by treatment condition.

Treatment condition	% Correct	% Blank	% Incorrect
30% verification	45%	28%	27%
30% verification + $1 penalty	69%	14%	17%
30% verification + $5 penalty	70%	15%	15%
70% verification	58%	16%	26%
70% verification + $1 penalty	68%	22%	10%
70% verification + $5 penalty	70%	16%	14%
90% verification	70%	10%	20%
90% verification + $1 penalty	88%	8%	4%
$15 penalty	82%	13%	5%
$15 penalty + 30% verification	84%	10%	6%

they make significantly better decisions in the two-institution conditions, relative to the corresponding one-institution conditions.

Further, although there are modest increases in the truthfulness of the statements subjects receive in the 30% chance of verification plus a $5 penalty for lying condition and in the 70% chance of verification plus a $5 penalty for lying condition (subjects receive truthful statements 100% of the time in these treatment conditions, as opposed to roughly 90% of the time in the other conditions), I do not observe corresponding increases in the amount of money subjects earn in these two conditions. That is, subjects earn similar amounts of money in the 30% chance of verification plus a $1 penalty for lying condition and in the 30% chance of verification plus a $5 penalty for lying condition. They also earn similar amounts of money in the 70% chance of verification plus a $1 penalty for lying condition and in the 70% chance of verification plus a $5 penalty for lying condition. Taken together, these results indicate that something other than increases in the truthfulness of the statements received is causing subjects to earn greater amounts of money in the two-institution conditions. They also suggest that the presence of an additional institution (and not necessarily its size or strength) is what matters for increasing trust and improving decisions.

Further analysis of subjects' decisions indicates that the increases in the amount of money they earn are driven by subjects' increased propensity to trust the statements they receive in the two-institution conditions, relative to the corresponding one-institution conditions. As shown in Table 5, subjects are less likely to leave the problems blank and/or answer incorrectly in the two-institution conditions, relative to the corresponding one-institution conditions. They are also more likely to answer the problems correctly in the two-institution conditions, relative to the corresponding one-institution conditions. The decrease in subjects' propensity to leave problems blank indicates that subjects are more willing to make decisions when a second institution is imposed upon the speaker. Given that subjects receive truthful statements over 90% of the time in both the one-institution and two-institution treatment conditions, subjects' increased participation and propensity to answer correctly indicates that they are more willing to trust and base their decisions upon the mostly truthful statements they receive when a second institution is imposed.

Why are subjects more willing to trust the speaker's statements when two institutions are present? One possibility is that subjects use the number of institutions as a heuristic. That is, instead of calculating their best responses and those of the speaker, subjects may rely on a 'two are better than one' heuristic, which induces greater trust in the speaker's statements when two institutions are imposed. Of course, additional experiments are needed to determine whether subjects actually use this heuristic and, if so, when and why they use it. Such experiments will shed further light on how people make decisions in institutional contexts like the ones in these experiments.[12]

7. Conclusion

In this paper, I presented a theory and experimental test of the conditions under which combinations of institutions induce citizens to trust an unknown speaker's statements and improve their decisions. The results reveal important differences between rational behavior as predicted by formal theory and subjects' actual behavior. Specifically, my theoretical model demonstrates that a second institution typically should not alter the speaker's behavior, nor should it change subjects' behavior or payoffs. What I observe empirically, however, is that a second institution increases subjects' payoffs substantially, relative to when only one institution is imposed upon the speaker. Further, the improvements subjects achieve with two institutions can be larger than when these institutions are imposed separately; thus, combinations of institutions can be greater than the sum of their parts. Interestingly, the improvements in decision making are not driven by increases in the truthfulness of the statements subjects receive. Rather, the additional institutions make subjects more willing to trust the already mostly truthful statements they receive, and this helps them to improve their decisions.

Of course, the question of whether and when combinations of institutions improve citizens' decisions is complex, because there are countless institutions that citizens might be exposed to in the real world. That said, this study takes one step toward answering this question by examining whether and when different combinations of two particular institutions induce citizens to trust a speaker's statements and improve their decisions. Given the many political, legal, and economic contexts in which citizens rely on the statements of others, this is an important set of institutions to examine. Indeed, in trial settings, jurors rely upon the statements of witnesses, who are subject to penalties for lying and whose statements can be verified during cross-examination. Similarly, uninformed voters may rely upon the statements of politicians, who might suffer a penalty (e.g. loss of reputation) if they lie and whose statements can be verified by an opponent or by the media. What my results imply for these settings is that even relatively small penalties for lying and slim chances of verification can increase trust and improve decision making when appropriately combined with other institutions. Further, increasing the probability of verification (for example, by increasing the quality of lawyers who represent criminal defendants or by increasing the number of political watchdog groups) might reduce the need for large penalties for lying. Thus, by strengthening our existing institutions or appropriately combining them with others, we can substantially increase citizens' propensity to trust the statements of others and make informed decisions.

That said, my results also suggest important limitations on the effectiveness of institutions. Specifically, the results show that, at least for the combinations of institutions examined in this study, the *presence* of an additional institution (as opposed to its size or strength) is what increases subjects' propensity to trust the speaker's statements, which helps them to improve their decisions. Indeed, in both the 30% and 70% chance of verification conditions, I find that adding a $5 penalty for lying to the chance of verification is no more effective than adding a $1 penalty for lying. This is particularly surprising, because the $5 penalty for lying should, theoretically, improve subjects' decisions when added to a 70% chance of verification. Further, in the 30% chance of verification conditions, neither the additional $1 penalty for lying nor the additional $5 penalty for lying should, theoretically, improve subjects' decisions, yet both have similar positive effects on subjects' decisions. The lesson this suggests for institutional design is that modest increases in the strength of additional institutions may not help citizens to make better decisions, even though we might expect them to do so. Of course, the results also indicate that if we strengthen additional institutions a great deal (for example, to a $15 penalty for lying), then this may help citizens to make better decisions.

This study also suggests that when citizens have preexisting knowledge about the options from which they must choose, institutions may be less effective than expected. Figure 1 illustrates that although combinations of institutions help subjects to improve their decisions, many combinations are not nearly as effective as predicted. In fact, although I expected subjects to earn the maximum amount of money in several treatment conditions, they did not do so until a 100% chance of verification was imposed upon the speaker. Although there are a variety of possible explanations for this result, a plausible explanation (given post-experiment interviews with subjects) is that some subjects trusted their own preexisting knowledge about how to solve math problems more than the speaker's statements. Thus, when the speaker's statement conflicted with their own idea of what the correct answer was, these subjects disregarded the speaker's statement and chose the answer that they believed was correct. Unfortunately, these subjects, by and large, were incorrect in their assessments of how to solve particular problems; thus, they often disregarded truthful statements and made suboptimal decisions as a result. Given that citizens in real-world institutional contexts often have preexisting knowledge about their decisions, this study demonstrates that even the most well-designed institutions can have unintended consequences if citizens refuse to abandon their preexisting knowledge even when it is likely beneficial for them to do so.

Acknowledgements

The author wishes to thank Gary Cox, Jeff Elman, Skip Lupia, Scott MacKenzie, Mat McCubbins, Sam Popkin, David Redlawsk, Shawn Rosenberg, Lydia Tiede, and participants in the Southern California Political Psychology Workshop at UC Irvine for their helpful comments.

Funding

This work was supported by the National Science Foundation [grant number SES-0616904] and the Kavli Institute for Brain and Mind.

Notes

1. Subjects earn money before they answer the math problems (e.g. $5 dollars for showing up, money for taking quizzes on the instructions). Before subjects answer the math problems, I tell them that they will either keep the money they have earned, lose it, or have more added to it. Thus, subjects know that they can lose money in the experiment.
2. Lau and Redlawsk (1997) develop measures that assess the correctness of citizens' votes.
3. It is important to develop experimental tasks that are analogous to real-world political decisions that subjects can perform (and that scholars can observe) in laboratory settings. Many tasks that scholars have developed do not *look* like real-world political decisions, but the psychological processes that occur as subjects perform these tasks closely relate to psychological processes that occur in the real world. Examples of this type of task include: (1) using abstract questions about gains versus losses to investigate framing effects (Druckman, 2001b; Kahneman and Tversky, 1979); (2) asking subjects to predict the outcomes of coin tosses to assess how endorsements affect decisions made under uncertainty (Lupia and McCubbins, 1998); and (3) giving subjects the opportunity to pass money to another subject to assess investment and trust (Wilson and Eckel, 2006). These examples show that even when an experimental task does not look like real-world politics, it can tell us a great deal about how citizens make political choices.
4. For further discussion of the internal and external validity of these experiments, please see Online Appendix 1: http://ps.ucdavis.edu/People/faculty/clboudre
5. The speaker makes her statement by putting a checkmark beside the answer that she wishes to report. The experimenter then reads that statement aloud to the other subjects. This prevents the speaker's tone of voice from confounding the experiment. Similarly, throughout the experiment, the speaker sits behind a partition so that the speaker's gender, race, and/or age do not affect the extent to which subjects trust the speaker's statements.
6. These treatment conditions have real-world analogues, as speakers in political settings often have incentives to misrepresent the truth, but may be deterred from doing so by sufficiently large penalties (such as a loss of reputation or monetary sanctions) and/or by the chance that another individual or organization (such as political watchdog groups) will verify their statements.
7. The $2 charge in the verification treatment conditions ensures that speakers cannot lie with impunity. Without this $2 charge, a speaker would have a dominant strategy to lie on every problem except when the chance of verification is 100% (in which case the speaker would be indifferent between telling the truth and lying). Because speakers in real-world settings typically incur costs (either reputational or monetary) if they are caught lying, they must calculate whether the expected benefits of lying are worth the risk of getting caught and incurring costs when making statements. To ensure that speakers in these experiments make similar calculations, I chose to use this $2 charge in the treatment conditions with verification.
8. I include multiple subjects in each experiment because, as Lupia and McCubbins (1998) note, this enables me to collect the data at a much lower cost (and in a much more efficient manner) than running experiments that include only one subject and one speaker.
9. My predictions are robust to different assumptions about a subject's likelihood of answering correctly on his own. Specifically, my predictions regarding the speaker's and subject's behavior remain the same if I assume that a subject has only a 50% chance of answering the problems correctly on his own (thereby doing no better than random guessing). My predictions also remain the same if I assume that a subject has anything between a 50% and 70% chance of answering correctly on his own. Further, even when I assume that a subject has as high as a 70% chance of answering correctly, only two predictions change.

10. The full calculations are shown in Online Appendix 2: http://ps.ucdavis.edu/people/faculty/clboudre
11. Thus, the number of observations listed in Table 2 reflects 381 subjects' decisions about the answers to 8–24 different math problems each, for a total of 5046 choices.
12. Thus, the number of observations listed in Table 3 reflects 43 subjects acting as 'the speaker' under different conditions. These 43 subjects made statements about the answers to 6–16 different problems each, for a total of 459 statements.
13. These results are consistent with economics experiments, which show that subjects often do not like to lie (Gneezy 2005).
14. I thank an anonymous reviewer for this suggestion.

References

Aronson E, Wilson TD and Brewer MB (1998) Experimentation in social psychology. In: Gilbert DT, Fiske ST and Lindzey G (eds) *The Handbook of Social Psychology*. 4th ed. Oxford: Oxford University Press.

Berggren HM (2001) Institutional context and reduction of the resource bias in political sophistication. *Political Research Quarterly* 54: 531–552.

Bohnet I, Frey BS and Huck S (2001) More order with less law: on contract enforcement, trust, and crowding. *American Political Science Review* 95: 131–144.

Boudreau C (2009a) Making citizens smart: when do institutions improve unsophisticated citizens' decisions? *Political Behavior* 31: 287–306.

Boudreau C (2009b) Closing the gap: when do cues eliminate differences between sophisticated and unsophisticated citizens? *Journal of Politics* 71: 964–976.

Capra CM, Tanaka T, Camerer CF, Feiler L, Sovero V and Noussair CN (2009) The impact of simple institutions in experimental economies with poverty traps. *Economic Journal* 119: 977–1009.

Crawford V and Sobel J (1982) Strategic information transmission. *Econometrica* 50: 1431–1451.

Dickson ES (2011) Economics versus psychology experiments: stylization, incentives, and deception. In: Druckman JN, Green DP, Kuklinski JH and Lupia A (eds) *Cambridge Handbook of Experimental Political Science*. Cambridge: Cambridge University Press.

Dickson ES, Hafer C and Landa D (2010) Rules of debate: theory and experiment. Working paper, New York University.

Downs A (1957) *An Economic Theory of Democracy*. New York: Harper.

Druckman JN (2001a) On the limits of framing effects: who can frame? *Journal of Politics* 63: 1041–1066.

Druckman JN (2001b) Using credible advice to overcome framing effects. *Journal of Law, Economics, and Organization* 17: 62–82.

Gneezy U (2005) Deception: the role of consequences. *American Economic Review* 95: 384-394.

Gordon SB and Segura GM (1997) Cross-national variation in the political sophistication of individuals: capability or choice? *Journal of Politics* 59: 126–147.

Guarnaschelli S, McKelvey RD and Palfrey TR (2000) An experimental study of jury decision rules. *American Political Science Review* 94: 407–423.

Iyengar S, Valentino NA, Ansolabehere S and Simon AF (1997) Running as a woman: gender stereotyping in political campaigns. In: Norris P (ed.) *Women, Media, and Politics*. Oxford: Oxford University Press.

Kahneman D and Tversky A (1979) Prospect theory: an analysis of decision under risk. *Econometrica* 47: 263–292.

Kuklinski JH, Quirk PJ, Jerit J and Rich RF (2001) The political environment and citizen competence. *American Journal of Political Science* 45: 410–424.

Lau RR and Redlawsk DP (1997) Voting correctly. *American Political Science Review* 91: 585–599.

Lupia A (1994) Shortcuts versus encyclopedias: information and voting behavior in California insurance reform elections. *American Political Science Review* 88: 63–76.

Lupia A and McCubbins MD (1998) *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* Cambridge: Cambridge University Press.

McDermott ML (1998) Race and gender cues in low-information elections. *Political Research Quarterly* 51: 895–918.

Morton RB and Williams KC (1999) Information asymmetries and simultaneous versus sequential voting. *American Political Science Review* 93: 51–67.

Ostrom E (1990) *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.

Ostrom E, Walker J and Gardner R (1992) Covenants with and without a sword: self-governance is possible. *American Political Science Review* 86: 404–417.

Popkin SL (1991) *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns*. Chicago, IL: University of Chicago Press.

Sniderman PM, Brody RA and Tetlock PE (1991) *Reasoning and Choice: Explorations in Political Psychology*. Cambridge: Cambridge University Press.

Tomz M, Wittenberg J and King G (2003) *CLARIFY: Software for Interpreting and Presenting Statistical Results, Version 2.1*. Available at: http://gking.harvard.edu/.

Wilson RK and Eckel CC (2006) Judging a book by its cover: beauty and expectations in the trust game. *Political Research Quarterly* 59: 189–202.