Cursed Beliefs with Common-Value Public Goods

Caleb A. Cox

Department of Economics and Finance, Durham University Business School
Mill Hill Lane, Durham DH1 3LB, United Kingdom
Email: caleb.cox@durham.ac.uk; Phone: +44 (0) 191 33 45425; Fax: +44 (0) 191 33 45201

Abstract

I show how improper conditioning of beliefs can reduce contribution in public goods environments with interdependent values. I consider a simple model of a binary, excludable public good. In equilibrium, provision of the public good is good news about its value. Naïve players who condition expectations only on their private information contribute too little, despite the absence of free-riding incentives. In a laboratory experiment, contributions indeed fall short of the equilibrium prediction. Using modified games with different belief-conditioning effects, I verify that subjects fail to condition beliefs properly. However, improper belief conditioning cannot fully explain the results.

Keywords: Public goods, experiments, cursed equilibrium, game theory
JEL Classification: C72, C92, D03, D71, H41

1 Introduction

The provision of public goods is a central issue in economics. Research on public goods has primarily focused on incentives to free-ride and various mechanisms for overcoming these incentives. In this paper, I demonstrate another force that may impede the provision of public goods, even in the absence of free-riding. In public goods environments with common or interdependent values, individuals may fail to correctly condition their beliefs about the uncertain value of a public good. Many public goods in the real world may have substantial common-value components, such as dispersed information about uncertain quality. Real-world public goods such as pollution abatement, national defense, police protection, and flood control may be of uncertain value, and information about the value may be decentralized. Individual contributors to such public goods should condition their beliefs about value on not only their private information, but also the information implicit in the strategic contribution choices of others. Failure to do so may lead to incorrect expectations about the value of the public good.

To isolate the belief-conditioning effect of interest in the absence of free-riding incentives, I consider a simple case of a binary, excludable public good (or club good), such as a toll road or private park. To illustrate, consider the choice of whether to participate in some costly group activity. The value of this activity is unknown, and information about the value is dispersed among the potential participants. Such information might come from individual experiences and knowledge or simply from intuition. Examples of such activities might include purchasing a membership to a planned recreation facility or a home in a new gated community, joining a joint business venture or working on a co-authored research project, or registering as a student in a new course at a university.[1] In order for the group activity to be viable, some minimum threshold of participants must be reached. If the threshold is not reached, the activity is cancelled and individuals who chose to participate pay no cost. Potential participants each observe private signals correlated with the uncertain value, and then simultaneously choose whether or not to participate. Each individual should consider two possible cases: the minimum threshold of participants is either reached or it is not. If the threshold is not reached, her decision to participate is inconsequential, as she will pay no cost. Thus, she should condition her expectations on the event that the threshold is reached. It is important to note that this event contains useful information about the value of the activity, since in equilibrium it implies that other participants observed relatively favorable signals. Thus, an individual who correctly conditions her beliefs on this event should expect the value to be higher than she would conditional on her private signal alone. Failure to properly condition beliefs would reduce contribution and provision relative to equilibrium.

The ability to share private information might alleviate this problem. However, there are a number of reasons why it may be difficult to share information. Beyond simple barriers to communication (such as difficulty sharing technical knowledge or simply not knowing each other), there may be incentives not to be truthful about private information. If there is a private value component so that the total value of the public good is not purely common to everyone, then there may be an incentive to lie to influence others. Similarly, if the good is not purely excludable, or if contributions may be unequal, then some form of free-riding incentive may prevent truthful communication. If the good is congestible, again it may be in an individual's interest to misrepresent her private information. In the simple case I consider in the experiment, incentives are fully aligned

[1] The last example comes from personal experience as a student registering for new course in game theory and experimental economics at the University of North Carolina at Charlotte in 2007, which I feared might be cancelled due to low enrollment.

so that individuals would have no incentive to lie if they could communicate. However, incentives to lie may exist in more complex cases.[2]

In the theoretical portion of this paper, I develop a simple model of excludable public goods with interdependent values and compare the predictions of Bayesian Nash equilibrium with naïve strategies, formalized by the cursed equilibrium model of Eyster and Rabin (2005). In their model, agents believe that, with some probability, others ignore their private information and choose an action according to the (equilibrium) ex ante distribution of actions. For this reason, each agent's belief about the distribution of actions chosen by others is correct, but agents do not fully account for the link between others' actions and their private information. I show that cursed beliefs reduce contribution relative to Bayesian Nash equilibrium, including the possibility of zero contribution for some parameter values.

Testing these predictions in the field would be problematic, since individuals' private information is unobservable. Therefore, I design a laboratory experiment to test whether improper conditioning of beliefs reduces contribution. The main treatment (the common-value threshold game) has 5 players in a group, with a threshold of 4 contributors required for provision. I vary the cost of contribution to determine whether contribution levels conform to Bayesian Nash equilibrium or naïve strategies for high, low, and intermediate costs. Rather than closely mimicking any particular real-world application, the experiment is designed to create a stark separation between the Bayesian Nash equilibrium and fully-cursed equilibrium predictions to examine the degree to which subjects (fail to) properly condition beliefs in making contribution choices.

Improper belief conditioning has been previously observed in other contexts, most famously in the winner's curse in common-value auctions. In common-value auctions, bidders should update their belief about value downward conditional on winning, while in my context, contributors should update their belief about value upward conditional on provision of the public good. In order to compare the results of the main treatment to the more well-known winner's curse in common-value auctions, I consider an "anti-threshold" game with the same environment, except that the public good is provided to contributors if and only if *no more* than 2 players contribute. The anti-threshold game is analogous to a simple common-value, two-unit auction with restricted bids and no

[2]These barriers to communication are similar to the discussion of Fedderson and Pesendorfer (1998) about why jury members may be unable to fully share private information.

trade in the case of excess demand. This treatment allows for comparison of behavioral responses to favorable and unfavorable belief conditioning effects, as well as comparison of how subjects learn to account for these effects over several rounds of play.

Sources of error other than improper belief conditioning might drive behavior away from equilibrium. To isolate the effect of belief conditioning, I consider a treatment with uncertain private values. Each subject has an uncertain private value for the excludable public good and observes a signal correlated with this value. While there is still uncertainty in this treatment, a given subject's value is uncorrelated with other subjects' signals. Therefore no subject has information about the value of the public good to others, which is a key difference from the common-value case. Play proceeds as in the main treatment. In this case, the symmetric Bayesian Nash equilibrium strategy precisely corresponds to the naïve (or fully-cursed) strategy from the common-value threshold game. Thus, if subjects are naïve, there should be no difference in behavior between these treatments, while correct conditioning of beliefs should lead to higher contribution in the common-value setting than the uncertain private values setting.

The experimental results show that contribution falls well below the BNE benchmark in the main treatment. Despite sharp differences in the Bayesian Nash equilibria of the games with favorable, unfavorable, and no belief-conditioning effects, actual behavior is quite similar between games, and in fact indistinguishable between the main treatment and the uncertain private values treatment. Thus, the results suggest that subjects *completely* fail to condition their beliefs in the proper direction. While fully-cursed equilibrium succeeds in predicting this similarity between treatments, it does not predict contribution levels very accurately. Moreover, behavior differs substantially from equilibrium even in the uncertain private values treatment, which cannot be explained by cursedness. This result highlights the importance of including a baseline without the potential for belief conditioning rather than using only theoretical benchmarks to examine belief conditioning effects.

The paper is organized as follows. Section 2 explores the related literature. Section 3 describes the model and theoretical predictions. Section 4 details the experimental procedures. Section 5 shows the results. Section 6 concludes with a discussion of the key findings. The Appendix contains proofs of the theoretical results from Section 3. Separate Online Appendices A and B contain supplementary data analysis and experimental instructions, respectively.

2 Related Literature

Many previous experiments consider non-excludable, step-level public goods and provision points, including Van de Kragt et al. (1983), Dawes et al. (1986), Isaac et al. (1989), Marks and Croson (1999), and Croson and Marks (2000). Provision point or threshold mechanisms have been generally successful in such environments under complete information or private values. Several experiments, such as Croson et al. (2006), Kocher et al. (2005), Swope (2002), and Bchir and Willinger (2013) find that excludability tends to increase contribution in a variety of linear and step-level public goods environments, while Czap et al. (2010) find higher contribution to non-excludable projects compared to excludable projects. Gailmard and Palfrey (2005) compare alternative cost-sharing mechanisms for excludable public goods and find that a voluntary cost-sharing mechanism with proportional rebates performs best.

Several papers explore uncertain returns in public goods experiments. In a voluntary contribution, linear public goods game, Dickinson (1998) finds that uncertain provision of the public good reduces contribution relative to certain returns in early rounds of play by a small but significant amount. Gangadharan and Nemes (2009) also find reduced contribution under uncertain provision of the public good in cases of known and unknown probability of provision. In a strategy-method public goods game with heterogeneous marginal returns, Fischbacher et al. (2014) find that uncertainty about one's own marginal return slightly decreases conditional contribution, but not unconditional contribution. Levati et al. (2009) find a large negative effect of uncertain marginal benefits of contribution, while Levati and Morone (2013) find mixed results depending on game parameters. Stoddard (2014a) finds that uncertain public good provision reduces contribution only when subjects are first exposed to the certainty baseline, suggesting potential order effects. Stoddard et al. (2014) and Stoddard (2014b) and find little difference between contributions under uncertain group returns compared to certainty. Results on the effect of uncertainty *per se* are thus somewhat mixed overall.

To my knowledge, the only prior consideration of interdependent-value public goods (excludable or non-excludable) is in the literature on leading by example, beginning with Hermalin (1998), and expanded to charitable giving by Vesterlund (2003), Potters et al. (2005), Andreoni (2006), and Potters et al. (2007). Unlike my symmetric, simultaneous-move setting, this literature examines informed and uninformed players moving sequentially, which is likely to make the information content of the leader's action relatively transparent compared to simultaneous-move games. Indeed, uninformed second movers do respond to the the information contained in the contribution choices of informed first movers in this environment.

This paper contributes to the public goods literature by showing how naïve beliefs can reduce contribution in public goods environments with common or interdependent values, even when free-riding incentives are absent. This effect is conceptually related to the winner's curse in common-value auctions (Thaler, 1988; Kagel, 1995; Kagel and Levin, 2002). In these environments, the bidder with the highest value estimate tends to win the auction, but because her estimate was the highest, it tends to be higher than the true value. In Bayesian Nash equilibrium, rational agents account for this adverse selection effect and condition their value expectations on winning the auction. However, in many experiments such as Kagel and Levin (1986), Kagel et al. (1995), and Levin et al. (1996), subjects fail to properly condition beliefs, leading to overbidding and low or negative profits. In my setting, similar naïvety causes subjects to choose not to contribute, even when their signals are high enough that contributing is optimal.

This paper is also closely related to the literature on strategic voting in common-value environments. Seminal theoretical analysis of such environments by Fedderson and Pesendorfer (1996, 1997, 1998) examines the behavior of strategic voters who condition their beliefs on being pivotal. Experiments including Guarnaschelli et al. (2000), Ali et al. (2008), Battaglini et al. (2008), Battaglini et al. (2010), and Esponda and Vespa (forthcoming) find evidence that laboratory subjects sometimes behave strategically, though their behavior is not always explained well by symmetric Bayesian Nash equilibrium.

I am also concerned with comparing behavior under favorable and unfavorable belief conditioning effects. Holt and Sherman (1994) compare these effects in the context of a takeover game. They find evidence of a "loser's curse" as well as a winner's curse, with subjects behaving naïvely in both environments.

The concept of naïve behavior in common-value auctions, strategic voting, takeover games, and related environments is formalized by the cursed-equilibrium model of Eyster and Rabin (2005). I use Eyster and Rabin's cursed equilibrium model as an alternative prediction to Bayesian Nash equilibrium and discuss the extent to which this model can explain the experimental data. While cursed equilibrium correctly predicts the similarity in behavior between treatments in my experiment, it does not explain the observed contribution levels.

3 Theory

3.1 Bayesian Nash Equilibrium

I first give the basic definitions and assumptions. The set of agents is $N = \{1,...,n\}$, where $n \geq 2$. I will use i and j to denote typical agents in N. Each agent observes a private signal x_i, which is a realization of a random variable X_i. The private signals are iid with probability density function $f : [\underline{x},\overline{x}] \to \Re_+$, which is assumed to be continuous and strictly positive everywhere on the interval $[\underline{x},\overline{x}]$, where $0 \leq \underline{x} < \overline{x} < \infty$. Let $F : [\underline{x},\overline{x}] \to [0,1]$ denote the corresponding cumulative distribution function and X denote an arbitrary random variable distributed according to F.

There is a binary excludable public good, and its uncertain value to agent i is v_i, given by:

$$v_i = \alpha x_i + \tfrac{1-\alpha}{n-1} \sum_{j \neq i} x_j, \tag{1}$$

where $\alpha \in [\frac{1}{n}, 1]$. The case of $\alpha = \frac{1}{n}$ corresponds to pure common value, where the value of the public good to all agents is the arithmetic mean of the private signals. The case of $\alpha = 1$ corresponds to pure private values.

The agents observe their private signals and then simultaneously choose whether or not to contribute an exogenous amount $w \in (\underline{x},\overline{x})$ toward provision of the public good. Denote the contribution decision of agent i given the signal x_i as $c_i(x_i)$, where $c_i(x_i) = 1$ indicates contribution and $c_i(x_i) = 0$ indicates non-contribution. The public good is provided if at least $k \in \{2,...,n\}$ agents contribute, otherwise contributions are refunded and no public good is provided. Any agent who does not contribute is excluded and gets a utility of zero. Contributors to the public good get a utility of $v_i - w$ if the public good is provided, and zero otherwise. All agents are assumed to be risk neutral.

I consider symmetric Bayesian Nash equilibria (BNE), so that in equilibrium, $c_i \equiv c$ for each agent i. That is, all agents have identical contribution decision functions. Lemma 1 shows that all such BNE involve "cutoff" strategies.

Lemma 1. In any symmetric BNE, there exists $x^* \in \Re$ such that each agent $i \in N$ strictly prefers to contribute to the public good if and only if $x_i > x^*$.

All proofs are contained in Appendix A. Intuitively, Lemma 1 holds because in symmetric BNE, each agent's expected utility of contributing is non-decreasing in the private signal, and strictly increasing when others contribute with positive probability.

Lemma 2 establishes that, conditional on at least $k-1$ others contributing, agent i's expectation of the mean signal of the other $n-1$ agents is non-decreasing in the cutoff x^*.

Lemma 2. Let the function $G_i(x^*)$ be given by:

$$G_i(x^*) = E\left[\frac{1}{n-1}\sum_{j\neq i} X_j \,\middle|\, \sum_{j\neq i} c^*(X_j) \geq k-1\right], \tag{2}$$

where:

$$c^*(X_j) = \begin{cases} 1 & : X_j \geq x^* \\ 0 & : X_j < x^*. \end{cases} \tag{3}$$

Then $G_i(x^*)$ is non-decreasing in x^*.

The result in Lemma 2 simply means that the expectation of the mean signal of the agents other than i conditional on at least $k-1$ others contributing is higher than the unconditional expectation, and this conditional expectation is non-decreasing in the cutoff. This result will be useful in proving the first Proposition.

In symmetric BNE, conditional on observing a signal $x_i = x^*$, agent i must be indifferent between contributing and not contributing. Thus,

$$\sum_{l=k-1}^{n-1}\binom{n-1}{l}(1-F(x^*))^l F(x^*)^{n-1-l}\left(\alpha x^* + \frac{(1-\alpha)l}{n-1}E[X|X\geq x^*] + \frac{(1-\alpha)(n-1-l)}{n-1}E[X|X<x^*] - w\right) = 0. \tag{4}$$

Clearly, $x^* = \overline{x}$ is a solution, so non-contribution by all agents is a symmetric BNE.[3] Proposition 1 gives conditions for the existence of an interior equilibrium.

Proposition 1. There exists a symmetric BNE cutoff $x^* \in (\underline{x}, \overline{x})$ if and only if:

$$\alpha\underline{x} + (1-\alpha)E[X] < w < \left(\alpha + \frac{(1-\alpha)(k-1)}{n-1}\right)\overline{x} + \frac{(1-\alpha)(n-k)}{n-1}E[X] \tag{5}$$

Moreover, there is at most one such interior symmetric BNE cutoff.

The key to Proposition 1 is to consider agent i's expected utility of contributing, given a signal of x^* and conditional on the public good being provided, treated as a function of

[3] In some cases, this trivial equilibrium may be weakly dominated. If $w < \alpha\overline{x}$ then agent i prefers to contribute conditional on observing $x_i > w/\alpha$.

the cutoff x^*. If w is within the given bounds, this function crosses zero somewhere in the interval $(\underline{x},\overline{x})$. Lemma 2 implies that this function is also strictly increasing in the cutoff, guaranteeing uniqueness of the interior equilibrium cutoff.

Corollary 1 gives comparative static predictions for changes in the cost of contribution and the provision threshold.

Corollary 1. Any symmetric BNE cutoff $x^* \in (\underline{x},\overline{x})$ is increasing in w and decreasing in k.

Intuitively, a higher cost of contribution makes agents less willing to contribute. A higher provision threshold strengthens the favorable belief conditioning effect, increasing willingness to contribute.[4]

3.2 Cursed Equilibrium

In (symmetric) χ-cursed equilibrium, agents fail to fully account for the connection between the actions of other agents and their private information. Each agent $i \in N$ believes that with probability χ, any given other agent j contributes with ex ante equilibrium probability regardless of j's signal.

Denote the χ-cursed equilibrium cutoff by x^*_χ. Proposition 2 establishes a simple condition under which a symmetric interior χ-cursed equilibrium exists and gives a simple explicit solution for the cutoff in fully-cursed equilibrium, where $\chi = 1$.

Proposition 2. There exists a symmetric χ-cursed equilibrium cutoff $x^*_\chi \in (\underline{x},\overline{x})$ if and only if:

$$\alpha\underline{x} + (1-\alpha)E[X] < w < \left(\alpha + \tfrac{(1-\chi)(1-\alpha)(k-1)}{n-1}\right)\overline{x} + \left(\tfrac{\chi(1-\alpha)(k-1)}{n-1} + \tfrac{(1-\alpha)(n-k)}{n-1}\right)E[X] \tag{6}$$

Moreover, there is at most one such interior symmetric χ-cursed equilibrium cutoff. Finally, for $\chi = 1$, if there exists an interior symmetric fully-cursed equilibrium cutoff, denoted by x^*_1, then it is given by:

[4]This comparative static prediction is not experimentally tested here. However, it guides the experimental design, as choosing k large relative to n increases the strength of the favorable belief conditioning effect and thus separation between symmetric BNE cutoffs and cursed cutoffs.

$$x_1^* = \frac{w}{\alpha} - \frac{1-\alpha}{\alpha} E[X] \tag{7}$$

The proof of Proposition 2 is similar to the proof of Proposition 1. Straightforward manipulation of the expression for the fully-cursed equilibrium cutoff reveals the intuitive interpretation: given a signal equal to the cutoff, the cost of contributing must equal the (naïve) expected benefit (neglecting belief conditioning).

Corollary 2 establishes that in symmetric cursed equilibrium, agents contribute less than the symmetric BNE prediction, and gives comparative static predictions for the cursed equilibrium cutoff.

Corollary 2. The interior symmetric χ-cursed equilibrium cutoff x_χ^* is non-decreasing in χ, increasing in w, and decreasing in k. In particular, $x_\chi^* \in [x^*, x_1^*]$.

Neglect of the favorable belief conditioning effect causes agent i's expectation of v_i to be too low, which reduces willingness to contribute. The greater the degree of cursedness (χ), the more severe is the reduction of contribution relative to the BNE benchmark.

Finally, Corollary 3 shows that, for some parameter values, cursedness may completely eliminate contribution.

Corollary 3. If $\alpha < 1$ and

$$\begin{aligned} &\left(\alpha + \tfrac{(1-\chi)(1-\alpha)(k-1)}{n-1}\right)\overline{x} + \left(\tfrac{\chi(1-\alpha)(k-1)}{n-1} + \tfrac{(1-\alpha)(n-k)}{n-1}\right) E[X] \\ &\leq w < \left(\alpha + \tfrac{(1-\alpha)(k-1)}{n-1}\right)\overline{x} + \tfrac{(1-\alpha)(n-k)}{n-1} E[X], \end{aligned} \tag{8}$$

then there is a symmetric BNE such that each agent contributes with positive probability, but in symmetric χ-cursed equilibrium contribution occurs with probability zero.

Intuitively, symmetric BNE and fully-cursed equilibrium coincide in the case of pure private values, where other agents' information does not affect agent i's expected utility of contributing conditional on the public good being provided. However, when values are interdependent, for some range of w contribution breaks down completely in χ-cursed equilibrium because agents ignore favorable belief conditioning in forming their expectations.

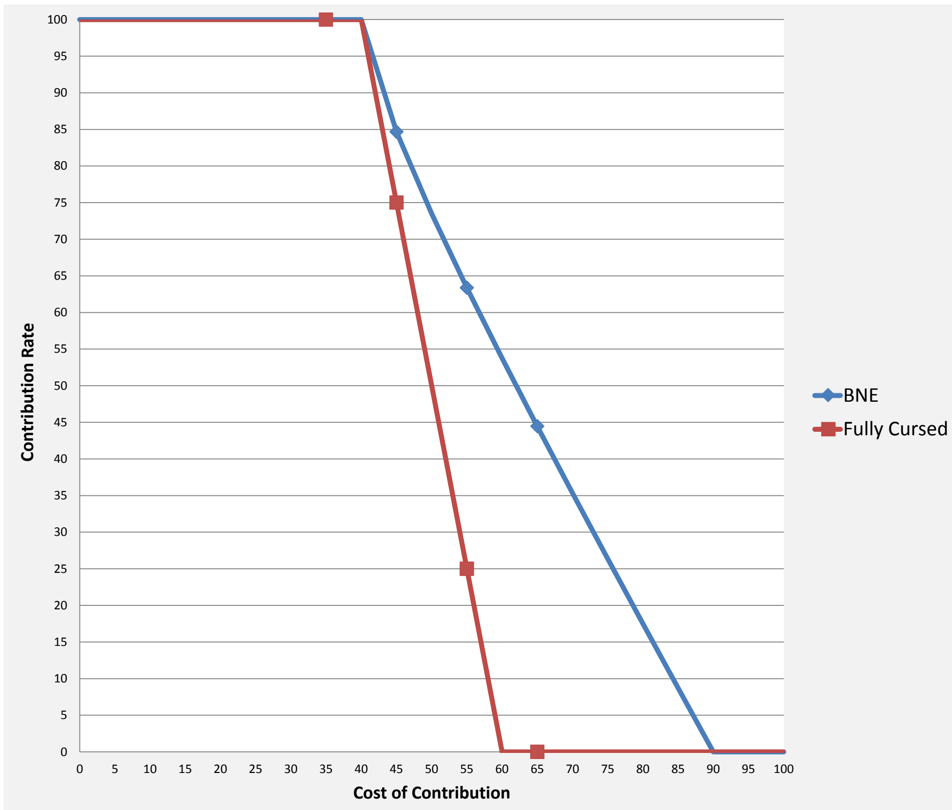

Figure 1: *Expected contribution rates in BNE and fully-cursed equilibrium*

3.3 Experimental Special Case

In the common-value threshold game (CVT) treatment, I consider a special case of the threshold game with $n = 5$, $k = 4$, $\alpha = \frac{1}{5}$, and private signals that are uniformly distributed on $[0,100]$, with w varying across rounds of play. Since $\alpha = \frac{1}{n}$, this special case is one of pure common value. The pure common value case is used in the experiment because it puts the most weight on the private signals of others and thus gives the greatest contrast between symmetric BNE and fully-cursed equilibrium. Henceforth I will omit the word "symmetric," since symmetric equilibria are the focus of the paper.

Figure 1 shows the expected contribution rates in BNE and fully-cursed equilibrium for different values of w in the interval $[0,100]$. The expected contribution rate (or percentage of agents who contribute on average) is equal to the ex ante probability that an individual agent contributes, or equivalently, 100 minus the equilibrium cutoff signal. The fully-cursed equilibrium contribution rates lie (weakly) below the BNE contribution rates for all values of w. Furthermore, as in Corollary 3, whenever $60 \leq w < 90$, contribution breaks down completely in fully-cursed equilibrium.

This special case illustrates the fact that the BNE is not fully efficient ex post. For example, for cost level 35, all players contribute for all possible signals in BNE, which may lead to inefficient provision if signals are low. The opposite problem may also occur: for example, suppose the cost level is 65 (where players contribute less than half the time) and there are three signals of 100 and two signals of 50. The common value is thus $80 > 65$, but provision will not occur in BNE.

Risk aversion might lead to reduced contribution relative to the risk-neutral BNE prediction, and thus it is important to check the robustness of the equilibrium prediction. Allowing for risk aversion makes analytical study of the model much less tractable, but approximate solutions can be found numerically. I use a constant relative risk aversion utility function of the form $u(y) = \frac{y^{1-r}}{1-r}$ and a coefficient of relative risk aversion of $r = 0.67$.[5] BNE cutoffs change very little with risk aversion, rising only by 1-2 percentage points compared to the risk-neutral prediction. Fully-cursed equilibrium cutoffs rise slightly more. Thus, cutoffs exceeding the BNE prediction by magnitudes shown in the fully-cursed equilibrium prediction could not be alternatively explained by plausible risk aversion. Furthermore, the presence of risk aversion does not affect the predicted treatment effects between CVT and the related games of interest.

3.4 Anti-Threshold Game with Unfavorable Belief Conditioning

To compare the favorable belief conditioning effects in the threshold game to similar unfavorable conditioning effects, I consider an "anti-threshold" (AT) treatment. The environment in the anti-threshold game is the same as in the common-value threshold game, except that the public good is provided if and only if *no more* than m agents contribute. If more then m contribute, the public good is not provided and contributions are refunded. The general case of the anti-threshold game is of less interest than the threshold game, so much of the theoretical analysis of the anti-threshold game is omitted. However, the equation characterizing the BNE cutoff is:

$$\sum_{l=0}^{m-1} \binom{n-1}{l} (1-F(x^*))^l F(x^*)^{n-1-l} \left(\alpha x^* + \frac{(1-\alpha)l}{n-1} E[X|X \geq x^*] + \frac{(1-\alpha)(n-1-l)}{n-1} E[X|X < x^*] - w \right) = 0 \qquad (9)$$

[5]This level of risk aversion has been found to be in the upper range of parameters typical of laboratory subjects by Holt and Laury (2002).

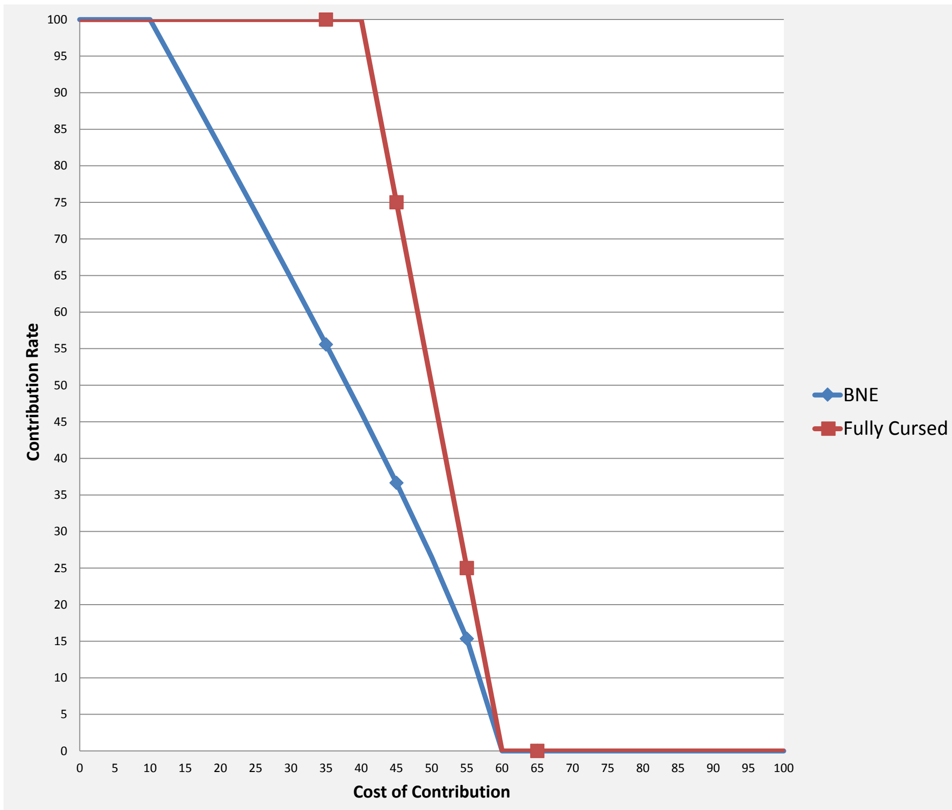

Figure 2: *Anti-threshold game expected contribution rates in BNE and fully-cursed equilibrium*

Notice that $x^* = 0$ (full contribution) is always a BNE. Under conditions similar to those in the previous section, interior BNE also exist. The key difference from the threshold game is that in the anti-threshold game, the public good being provided is *bad news* about its value, while in the threshold game it is good news.

In the AT treatment, I consider the special case of $n = 5$, $m = 2$, $\alpha = \frac{1}{5}$, and private signals uniformly distributed on $[0, 100]$, with w varying across rounds of play. Figure 2 shows the expected contribution rates in the anti-threshold game in BNE and fully-cursed equilibrium for varying w. As in the case of CVT, the expected contribution rate equals 100 minus the equilibrium cutoff signal. Notice that contribution rates in fully-cursed equilibrium are exactly the same as those from CVT. However, in AT, fully-cursed agents contribute more than the BNE prediction.

There is a simple symmetry between the AT and CVT. Fixing $\delta \in [-50, 50]$, the absolute difference between the BNE and fully-cursed equilibrium contribution rates in CVT with $w = 50 + \delta$ is equal to the absolute difference between BNE and fully-cursed equilibrium contribution rates in the anti-threshold game with $w = 50 - \delta$. Thus, the belief conditioning effects in CVT and AT are in this sense comparable in magnitude, but opposite in direction.

3.5 Private-Value Threshold Game with No Belief Conditioning Effect

In the private-value threshold game (PVT) treatment I consider a game similar to the one in CVT, except that each agent's value for the public good is the mean of five agent-specific *iid* random draws. One of the five is observed by the agent, while the other four are not observed by anyone. Thus, the ex ante marginal distribution of each agent's value is the same as in CVT, but there is no conditioning effect. In fact, the symmetric fully-cursed equilibrium in CVT is identical to the symmetric Bayesian Nash equilibrium (and any cursed equilibrium) in PVT. Thus, by comparing contributions between PVT and CVT, the effect of favorable belief conditioning in CVT can be observed.

I also study an individual-choice version of the PVT game, in which the threshold k is equal to 1 rather than 4. This treatment is denoted PVTk1. Changing k does not alter the optimal strategies in this case: the predicted cutoffs are the same as the BNE cutoffs in the PVT game and the fully-cursed cutoffs in the CVT game. However, turning the PVT game into an individual choice problem steepens incentives, since every player is always pivotal in this case.[6] This treatment was added to explore some anomalous behavior in PVT, discussed in the Section 5.

4 Experimental Procedures

To avoid negative payoffs, the cost of contributing is implicit, so that each participant is faced with a choice between a certain payoff of w and an uncertain payoff of v.[7] The conversion rate is \$0.20 for each experimental currency unit (or "token"), so that the maximum possible earnings are \$20 per person. Participants also received a \$5 show-up fee.

[6]In addition to steepening incentives, changing from a game to an individual choice problem could also alter pro-social motives for contribution. To minimize this potential confound, the framing of the decision is kept as close as possible to the original PVT treatment. Subjects are matched into groups of five, and the risky alternative is again referred to as a "group project" despite the contribution threshold of only one. Subjects are given feedback on other group members choices, which also allows for possible imitation learning effects which might be present in PVT.

[7]Framing in terms of explicit rather than implicit costs might affect behavior and learning (Lind and Plott, 1991). However, in the treatment of primary interest (CVT), the experience from which subjects are expected to learn is the failure to realize profitable public goods, which is inherently implicit.

There are two treatment variables. The first, varied between subjects, is the game: CVT, AT, PVT, and the supplementary treatment PVTk1. Only one of these games appeared in any given session. The second treatment variable, varied within subject, is the cost of contributing: 35, 45, 55, and 65 experimental currency units, with each value repeated five times in randomized order. Each session had twenty rounds, one of which was selected randomly for payment. Subjects in each session were randomly assigned to groups of five at the start of each round (stranger matching).

In each round, each participant observed the cost of contributing and her own private signal. Contribution choices were then made simultaneously. At the end of each round, all participants observed the signals and choices of the other four group members (ordered from highest to lowest), the value of the public good and whether it was provided, and their own earnings in tokens for the round.[8]

The experiment was programmed and conducted using z-Tree (Fischbacher, 2007). All sessions were run in the experimental economics laboratory at The Ohio State University. Nine sessions were run (3 CVT, 2 AT, 2 PVT, and 2 PVTk1), each with 20 subjects, except session of PVTk1, with only 15 subjects due to absences.[9] Participants earned approximately $15.50 on average, and each session lasted about 45 minutes.

5 Results

5.1 Hypotheses

To organize the results, I first summarize the key hypotheses to be tested. The main hypotheses come from the predictions of cursed equilibrium compared to Bayesian Nash equilibrium.

Hypothesis 1 (Contribution within Games)**.** Under full or partial cursedness, subjects will choose to allocate tokens to the group project too little in CVT and too much in AT, relative to BNE.

Hypothesis 2 (Contribution between Games)**.** Under full cursedness, subjects will choose to allocate tokens to the group project with the same frequency in CVT, PVT, and AT.

[8] The signals and choices of other group members were displayed in decreasing order by signal to make it easier to notice any correlation between signals and choices.

[9] Due to a recruitment system error, two subjects were mistakenly allowed to participate a second time. The choices made by each of these subjects in their second session of participation have been excluded from the analysis.

Hypothesis 1 comes from the neglect of belief conditioning in CVT and AT under full or partial cursedness. Hypothesis 2 comes from the fact that the fully-cursed equilibrium cutoffs in CVT, PVT, and AT are identical.

I will also test whether steeper incentives in PVTk1 relative to PVT reduces non-equilibrium behavior. Additional questions of secondary interest are concerned with learning over multiple rounds of play and individual heterogeneity. Subjects might learn to play closer to BNE, and individuals with greater cognitive or quantitative ability might play strategies closer to BNE. However, results on learning and ability have proven negative, and so much of this analysis is omitted from the main body of the paper and instead discussed in Online Appendix A.

5.2 Contribution Rates

Figure 3 shows aggregate contribution rates (the average percentage of subjects choosing to contribute) in CVT, PVT, and AT. Graphical presentation of PVTk1 is left to Online Appendix A. Logit regressions reported in Table 1 are used to formally test Hypotheses 1 and 2. The dependent variable is an indicator for contribution to the group project, and the omitted category for the game indicators is PVT. Standard errors are clustered by session.[10] Notice that in Figure 3, contribution rates in CVT and PVT appear virtually the same for all cost levels, and indeed the regression results show no significant difference between contributions in CVT and PVT, consistent with Hypothesis 2. Even where the difference is greatest (the higher cost levels) it is in the opposite direction predicted by BNE, with slightly lower contribution in CVT than in PVT. Contributions are significantly lower in AT than in PVT for higher cost levels, but still appear much closer to one another than BNE predicts.

The logit regression results in Table 1 can also be used to compare the contribution

[10]Fréchette (2012) suggests clustering by session as a robust approach to account for possible correlation between subjects interacting with one another in the same session. To adjust for potential bias in hypothesis tests due to the small number of clusters, I use the t-distribution with degrees of freedom equal to the number of clusters minus the number of regressors rather than using the standard normal reference distribution (Donald and Lang, 2007; Cameron and Miller, forthcoming). An alternative approach to session-level clustering also suggested by Fréchette (2012) is controlling for feedback effects but assuming independence between subjects conditional on this feedback. A variety of specifications of this kind show similar results. The results also remain similar simply clustering at the subject level without such feedback controls. I have also used non-parametric chi-squared proportion tests (using the clustering adjustment of Rao and Scott 1981, 1984; see also Sribney, 1998). These non-parametric tests show similar results comparing PVT with CVT and AT, but find no significant difference between PVT and PVTk1.

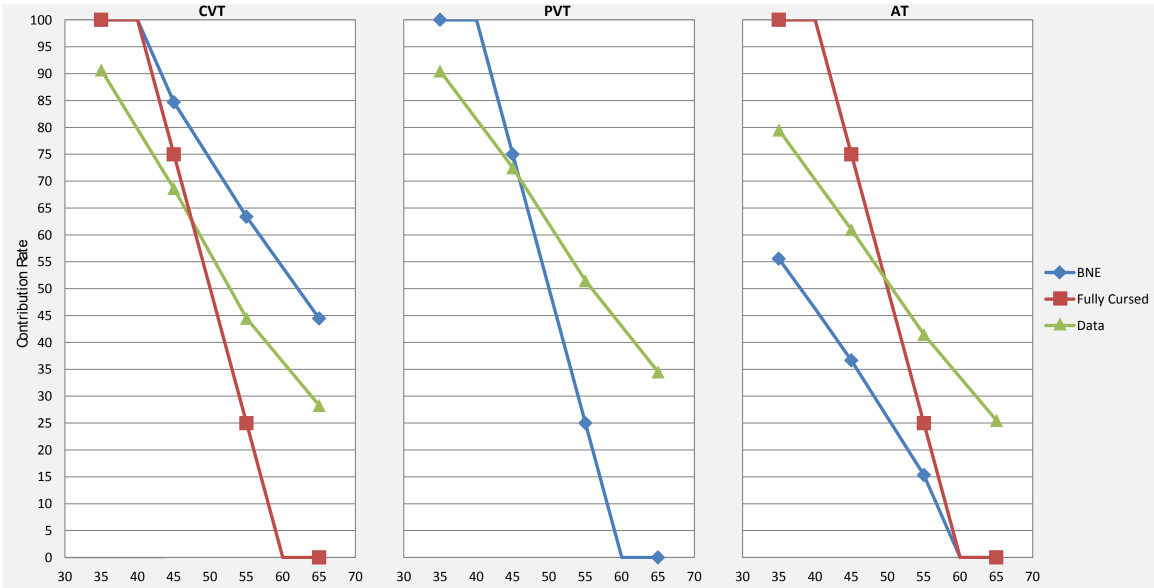

Figure 3: *Aggregate rates of contribution.*

rates in each treatment with theoretical benchmarks.[11] Contribution rates in CVT are significantly lower than the BNE prediction at every cost level. Moreover, contribution rates in AT are significantly higher than the BNE prediction at every cost level. These results are consistent with Hypothesis 1. However, cursedness cannot explain contributing too infrequently in CVT at cost level 35, nor can it explain over-contributing in AT at cost level 65. In both of these cases, BNE and cursed equilibrium are identical. There are also statistically significant differences between the contribution rates in the data and the fully-cursed equilibrium benchmark predictions. Some of these differences cannot be explained by partial cursedness either, such as above-fully-cursed contributions in AT at higher cost levels.

Cursed equilibrium also cannot explain contribution rates in PVT, which differ substantially from the (identical) predictions of BNE and cursed equilibrium. Such deviation from BNE even without the potential for belief conditioning underscores the importance of studying belief conditioning by comparing the CVT and PVT treatments rather than only comparing the data to theoretical benchmarks within one treatment. Clearly, forces other than cursedness must drive behavior away from the BNE prediction.

[11] I derive such comparisons by computing margins (contribution rates) and cluster-robust standard errors for each treatment and testing for equality with theoretical benchmarks, using the t-distribution with degrees of freedom equal to the number of clusters minus the number of regressors.

	w=65		w=55		w=45		w=35	
Variable	(i)	(ii)	(i)	(ii)	(i)	(ii)	(i)	(ii)
signal	-	1.025**	-	1.033***	-	1.028**	-	1.018**
		(0.006)		(0.004)		(0.005)		(0.006)
round	-	0.963	-	0.992	-	0.986	-	1.054*
		(0.017)		(0.013)		(0.007)		(0.019)
CVT	0.748	0.794	0.755	0.662	0.829	0.831	1.023	1.026
	(0.141)	(0.169)	(0.156)	(0.129)	(0.177)	(0.225)	(0.499)	(0.433)
AT	0.650***	0.659**	0.668***	0.582**	0.593*	0.551	0.407	0.329
	(0.066)	(0.068)	(0.071)	(0.084)	(0.144)	(0.167)	(0.213)	(0.161)
PVTk1	0.527*	0.511*	0.812	0.777	0.948	1.074	0.814	0.845
	(0.158)	(0.124)	(0.107)	(0.090)	(0.221)	(0.295)	(0.502)	(0.424)

Table 1: *Logit regression results. The dependent variable is an indicator for contribution. Odds ratios are reported with cluster-robust standard errors in parentheses. Each regression includes 865 observations and 9 session-level clusters. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.*

In particular, there is a puzzling tendency for subjects to over-contribute in PVT (and AT) with $w = 65$. Even with a signal of 100, the expected value of the group project is no greater than 60 tokens, and thus subjects should never contribute in PVT when the cost is 65 tokens. As discussed in more detail in Online Appendix A, this result does not seem to be driven by mathematical errors or risk-seeking preferences, nor is it driven by only a few subjects. By comparing PVT with PVTk1, I am able to test whether over-contribution in PVT at the highest cost level might be due to weak incentives caused by the low probability of provision of the group project. Changing the threshold k does not alter the optimal strategies in this case: the predicted cutoffs are the same as the BNE cutoffs in the PVT game and the fully-cursed cutoffs in the CVT and AT games. However, incentives are steeper in PVTk1, since every player is always pivotal.

For the highest cost level, contributions are lower in PVTk1 compared to PVT. Though the difference is only marginally significant, this result is consistent with the supplementary hypothesis that over-contribution in PVT at cost level 65 is driven by weak incentives due to the low probability of provision. No significant differences are apparent for lower cost levels, where provision in PVT is more frequent.[12] More detailed comparison of PVT and PVTk1 is left to Online Appendix A.

Note that round of play is included in the regressions in Table 1 as a control, showing a significant effect only for the lowest cost level. However, learning effects may go in different directions for different treatments, as subjects may, for example, intially contribute too infrequently in CVT and too frequently in AT. Since the number of session-level clusters is small, I omit round-treatment interactions for the sake of parsimony. In Online Appendix A, I examine potential learning effects in more detail (and under stronger assumptions about the nature of intra-session correlation), finding little evidence of consistent learning patterns.

The following main results summarize the findings on Hypotheses 1 and 2.

Result 1 (Contribution within Games). Relative to BNE, subjects contribute to the group project too infrequently in CVT and too frequently in AT. However, some deviations from BNE cannot be explained by cursedness. In particular, subjects contribute too often in PVT with cost levels of 55 and 65, which is not predicted by cursedness.

Result 2 (Contribution between Games). Contribution choices in CVT and PVT are indistinguishable. Contribution choices in AT differ from those in the other games, but this difference is smaller than predicted by BNE.

5.3 Cutoff Strategies

While the rates of non-contribution (choosing the private account) may provide a rough estimate of the average cutoff subjects use, perhaps a more appealing method is maximum likelihood, similar to the method of El-Gamal and Grether (1995). Under the assumption that all subjects use the same cutoff (but may make errors), I estimate the cutoff for each game and cost level by checking all possible cutoffs and finding the one that explains the most data, or equivalently, minimizes the number of errors. I assume that with probability $1-\epsilon$, an agent makes a contribution choice consistent with the hypothesized cutoff, and with probability ϵ (the error rate), she makes the opposite choice. The maximum likelihood cutoff is the cutoff that minimizes the observed error rate.

[12] The similarity of contribution rates between PVT and PVTk1 at cost levels 55, 45, and 35 suggests that pro-social motives do not drive the difference between PVT and PVTk1 at cost level 65. If pro-sociality were the driving force, contributions would likely be lower in PVTk1 at all cost levels.

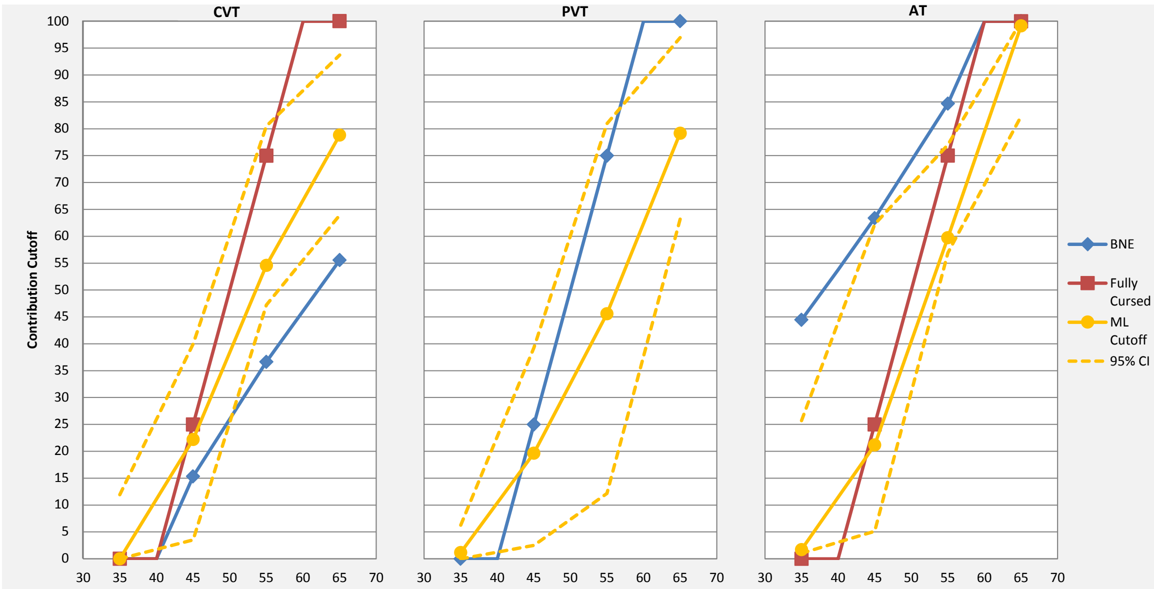

Figure 4: *Maximum Likelihood Cutoffs*

Figure 4 shows maximum-likelihood cutoffs with 95% bootstrap confidence intervals.[13] Once again, the estimated cutoffs suggest below-BNE contribution in CVT and above-BNE contribution in AT. It is also clear that estimated cutoffs are very similar between CVT and PVT. I find no signficant differences between the CVT and PVT cutoffs for any cost level using bootstrap hypothesis tests.[14] This similarity further suggests that subjects treat the CVT and PVT games as equivalent, despite the substantial differences between their BNE.

While behavior does not appear to be consistent with BNE, it is also of interest how closely behavior approximates an empirical best response. Figure 5 compares maximum likelihood cutoffs with empirical best response cutoffs. Empirical best response cutoffs can be easily computed from equation 1 using the empirical probability of contribution to the group project, average signal conditional on contribution, and average signal conditional on non-contribution. Maximum likelihood cutoffs are generally not very close to empirical best response cutoffs where the cutoffs are interior, with the exception of

[13]To adjust for potential correlation within sessions, I use a two-stage resampling procedure (Davison and Hinkley, 1997; McCullagh, 2000). Sessions are randomly selected with replacement, and then subjects within the selected sessions are randomly selected with replacement. This procedure yields wider confidence intervals compared to a naïve bootstrap, but there is little substantive difference in results.

[14]These tests are not independent of the previous comparisons of contribution rates, but together they give a clearer description of the similarity between the CVT and PVT data.

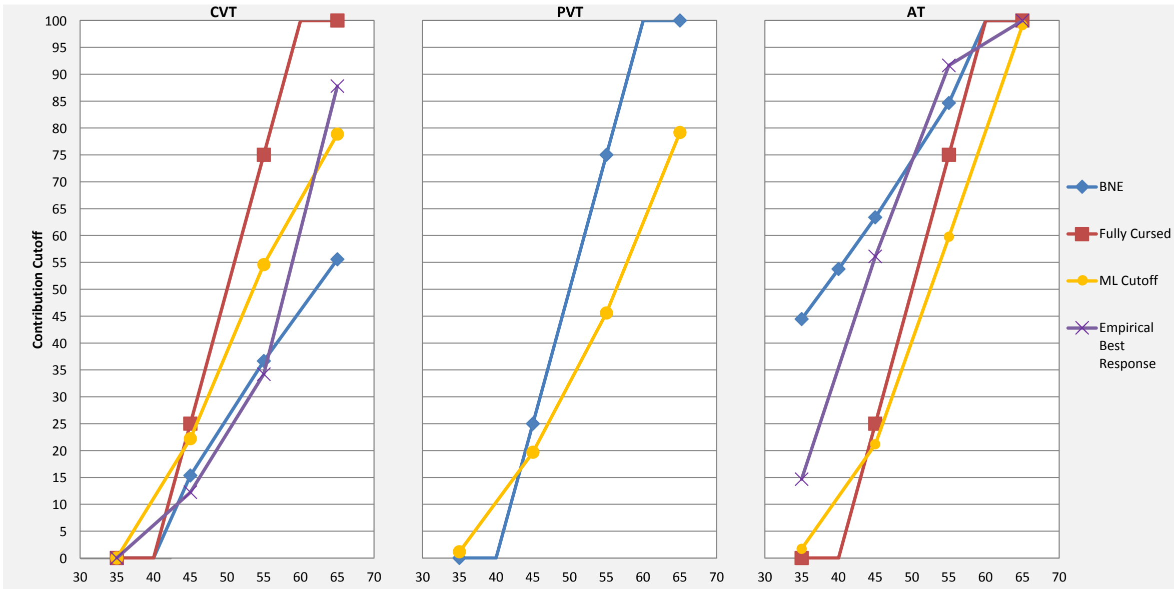

Figure 5: *Empirical Best Responses*

CVT with $w = 65$.[15] Thus, neither BNE, nor fully-cursed equilibrium, nor empirical best response appears to explain the aggregate data well.

Assuming some partially-cursed equilibrium holds across all rounds and all cost levels, the cursedness parameter χ can be estimated by maximum likelihood for CVT and AT.[16] For CVT, the maximum likelihood estimate of the cursedness parameter is 0.60, with a 95% bootstrap confidence interval of [0.54, 0.92]. In fact, the 0.60-cursed equilibrium cutoffs are quite close to the previous (unrestricted) maximum likelihood cutoffs estimates for CVT.

For AT, the maximum likelihood estimate of the cursedness parameter is 0.91 with a 95% bootstrap confidence interval of [0.48, 1.00]. The 0.91-cursed equilibrium cutoffs are not particularly close to the unrestricted maximum likelihood cutoff estimates, since for cost levels 45 and 55 the unrestricted estimates do not fall between the BNE and fully-cursed equilibrium cutoffs. The data in AT are somewhat noisier than in the other

[15]Empirical best response cutoffs may be either higher or lower than BNE cutoffs, depending on behavior. Using a cutoff above the BNE cutoff tends to drive the empirical best response cutoff downward in CVT, since the favorable conditioning effect is strengthened. However, the opposite type of "mistake" (contributing when the signal is too low) has the opposite effect on the empirical best response cutoff.

[16]This estimation follows the same approach of selecting cutoffs to minimize errors as previously discussed. However, I add the restriction that cutoffs for each of the four cost levels within a game (CVT or AT) must be consistent with some partially-cursed equilibrium cutoff.

games, perhaps due to the less intuitive nature of the game.

Result 3 (Cutoff Strategies)**.** Estimated cutoff signals are above BNE in CVT and below BNE in AT (consistent with Hypothesis 1), and again show strong similarity between CVT and PVT (consistent with Hypothesis 2). While partially-cursed equilibrium fits the estimated CVT cutoffs well, cursedness cannot explain the estimated cutoffs in PVT and AT, particularly for the highest cost level.

In this subsection, I have assumed that subjects in a particular game use the same cutoff strategy, as in the BNE and cursed-equilibrium benchmarks. However, this assumption may not hold. In Online Appendix A, I show graphically the estimated probability of contribution as a function of the signal in CVT, PVT, and AT at each cost level. If all subjects in a particular game use the same cutoff strategy (but may make errors), these graphs should resemble step functions. However, the estimated probabilty of contribution appears to increase smoothly with the signal. It is possible that subjects may employ probabilistic strategies. Another possibility, explored next in Subsection 5.4, is that subjects choose heterogeneous cutoff strategies.

5.4 Individual Heterogeneity

I next specify several candidate strategies to which individual choices are compared. The candidate strategies include BNE and fully-cursed equilibrium, as well as several other possible heuristics. One possibility is that a subject may misperceive the correlation between her own signal and the value of the group project. In the extreme case of perfect correlation, she would contribute if and only if the signal exceeds w. I denote this heuristic strategy "Signal Bias." In the opposite extreme case of perceiving no correlation between the signal and the value of the group project, she would simply contribute if and only if $w < 50$, that is, when the cost of contributing is less than the prior expected value of the group project. I denote this heuristic strategy "Prior." I also consider the simple strategies of always contributing and never contributing.

I use a Bayesian approach to estimate the proportion of subjects in each candidate strategy. I assume that each individual is playing one of the candidate strategies, but may make errors. In any individual game I assume that with probability $1-\epsilon_i$, player i follows her chosen strategy, and with probability ϵ_i she deviates. First, an individual subject's choices over all twenty rounds are compared to the predictions of each candidate strategy. The error rate ϵ_i for player i is estimated as her smallest observed fre-

	CVT		AT	
Strategy	Proportion	Std. Error	Proportion	Std. Error
Fully-Cursed	0.312	0.047	0.359	0.054
BNE	0.233	0.046	0.130	0.035
Prior	0.219	0.043	0.154	0.035
Signal Bias	0.120	0.034	0.187	0.044
Always	0.108	0.036	0.097	0.036
Never	0.008	0.003	0.073	0.036

Table 2: *Estimated strategy proportions in CVT and AT*

quency of deviations over all candidate strategies.[17] So for example, if player i's choices were consistent with the fully-cursed strategy 95% of the time and less frequently consistent with any other strategy, her estimated error rate would be 0.05. Next, I set a uniform prior over all candidate strategies and update for each observation according to Bayes' rule to arrive at a posterior over the candidate strategies for each individual subject. For each candidate strategy, the posterior probability is averaged across subjects to estimate the overall proportion of subjects playing that strategy.[18]

Table 2 shows the estimated proportion of subjects playing each candidate strategy in CVT and AT. In both games, the fully-cursed strategy is modal, consistent with the aggregate results showing neglect of belief conditioning. The BNE strategy is second-most prevalent in CVT with a proportion of nearly one quarter, though the vast majority appear to play some boundedly-rational strategy. The BNE strategy is less prevalent in AT than CVT, which might suggest that AT is a more difficult game.

To check for correlations between consistency with the BNE strategy and cognitive/quantitative ability, I have run a number of regressions, reported in more detail in Online Appendix A. However, the results have been negative, suggesting that some subjects' behavior may simply appear to closely match BNE by chance rather than strategic sophistication.[19] Furthermore, estimating strategy proportions in the PVT data using

[17] Very few subjects are always consistent with a single candidate strategy, but 21.7% of subjects in CVT are at least 95% consistent. Just over half of such subjects closely match the fully-cursed strategy, while about one quarter closely match BNE. Only 7.5% of subjects in AT are at least 95% consistent.

[18] The results are reasonably robust to alternative error structures and non-uniform priors. The MLE partially-cursed strategy is not included as a candidate strategy since it is a free parameter estimated from the data rather than being specified a priori. However, if it is included, it becomes modal in CVT and second-most prevalent after the fully-cursed strategy in AT, and the prevalence of BNE falls substantially.

[19] This finding is similar to Georganas et al. (2013), who found very little correlation between measures of cognitive ability and playing more sophisticated strategies in undercutting and guessing games.

the CVT strategies also yields an estimate of approximately one quarter of subjects playing the BNE for the CVT game. However, the BNE strategy from CVT does not have any particular justification or heuristic intuition in the PVT game. Recall that the actual BNE strategy in PVT is identical to the fully-cursed equilibrium strategy from CVT. Thus, there is no reason to expect subjects in PVT to play the BNE strategy from CVT, except perhaps by chance. The similarity of estimated proportions in CVT and PVT playing the BNE strategy from CVT further suggests that strategic sophistication does not drive consistency with the BNE strategy in CVT. Therefore, it appears that very few, if any, subjects are able to properly condition beliefs in this setting.

Result 4 (Individual Heterogeneity)**.** Estimated proportions of strategic types show that fully-cursed behavior is modal, and that the great majority of subjects play some boundedly-rational or heuristic strategy.

6 Discussion

In this paper, I have demonstrated that a severe neglect of belief conditioning can impede the provision of common-value excludable public goods. In the CVT game, a favorable conditioning effect arises in Bayesian Nash equilibrium because the expected value of the public good conditional on sufficiently many others contributing is higher than this expectation conditional on the private signal alone. However, subjects fail to account for this effect, consistent with cursed equilibrium. Furthermore, behavior in this game is indistinguishable from behavior in the closely-related PVT game, in which conditioning effects are absent. There is also a surprising similarity in behavior between the CVT game (with a favorable conditioning effect) and the AT game (with an unfavorable conditioning effect). The fully-cursed equilibria of CVT, PVT, and AT are identical, while there are sharp differences in their Bayesian Nash equilibria. Thus, the similarity in behavior between games is consistent with cursedness.

However, cursed equilibrium cannot explain the observed contribution rates. In particular, the level of contribution in PVT with the highest contribution cost is unexplained by cursed equilibrium or Bayesian Nash equilibrium. The decrease in contribution in PVTk1 (the individual-choice version of PVT) suggests that weak incentives partially drive contribution in this case, since the probability of provision is low. Such weakness of incentives is also present in CVT with the highest contribution cost, and may have also driven some contributions in this case.

This paper contributes to the the literature on public goods by identifying a novel impediment to contribution that is distinct from free-riding. While this experiment is designed to provide a clear separation between equilibrium and naïve contribution choices, the behavioral phenomenon found here may also be important in field environments. In a number of applications within public economics and industrial organization, such as the provision of gated communities and the formation of joint ventures, naïve contribution choices may cause a failure to coordinate on efficiency-enhancing outcomes. Future research might examine the design of optimal mechanisms for information aggregation in such environments. In the simple case that I consider, the incentives of individual agents are aligned under pure common value, so that agents would truthfully reveal their signals if they could. However, incentives to lie may exist in closely-related cases in which some form a free-riding is possible. Private value components, unequal contributions, or lack of excludability all lead to the possibility of free-riding in some form, which may give individual players an incentive to misrepresent their private information.

This paper also contributes to the literature on cursedness in related contexts such as common-value auctions and voting games by examining cursed equilibrium in a novel game and showing a potentially important consequence of this type of bounded rationality. Importantly, my experimental design demonstrates the failure to properly condition beliefs by the comparison of the CVT and PVT treatments. While comparing behavior to theoretical benchmarks within a treatment is also useful, the treatment comparison controls for other potential sources of decision error while varying only the presence of belief conditioning effects. The treatment comparison suggests that subjects not only fail to fully condition beliefs, but actually fail to condition *at all*.

I have focused on the case of excludable public goods (such as gated communities and private parks) to isolate the neglect of belief conditioning in the absence of free-riding incentives. Future research might explore the idea of pure public goods with interdependent values.[20] Examples include pollution abatement and flood control, for which values are likely to be strongly correlated, but uncertain. This study provides a first step toward a promising line of inquiry on coordination and information aggregation in environments with common-value public goods.

[20]The neglect of belief conditioning in a pure public goods context might be called a "Free-Rider's Curse," though in the current excludable context, there is no free-riding.

Acknowledgements

I thank Paul J. Healy, Matthew Jones, Yaron Azrieli, John Kagel, Dan Levin, James Peck, Lucas Coffman, Katie Baldiga Coffman, Daeho Kim, David Blau, Semin Kim, Hong Il Yoo, Brock Stoddard, Xi Qu, Greg Howard, Kerry Tan, Alan Horn, Michael Caldara, and Dimitry Mezhvinsky for helpful comments and suggestions. I am also very grateful for helpful feedback from Co-Editor Tim Cason and two anonymous referees. Any remaining errors are my own responsibility. This research was funded in part by National Science Foundation grant #SES-0847406 (Paul J. Healy, P.I.) and in part by the JMCB Grants for Graduate Student Research Program.

Ali, S. N., Goeree, J. K., Kartik, N., Palfrey., T. R., 2008. Information aggregation in standing and ad hoc committees. *American Economic Review* 98, 181–186.

Andreoni, J., 2006. Leadership giving in charitable fund-raising. *Journal of Public Economic Theory* 8, 1–22.

Battaglini, M., Morton, R. B., Palfrey, T. R., 2008. Information aggregation and strategic abstention in large laboratory elections. *American Economic Review* 98, 194–200.

Battaglini, M., Morton, R. B., Palfrey, T. R., 2010. The swing voter's curse in the laboratory. *Review of Economic Studies* 77, 61–89.

Bchir, M. A., Willinger, M., 2013. Does a membership fee foster successful public good provision? An experimental investigation of the provision of a step-level collective good. *Public Choice* 157, 25–39.

Cameron, A. C., Miller, D. L., forthcoming. A practitioner's guide to cluster-robust inference, *Journal of Human Resources*.

Croson, R., Fatás, E., Neugebauer, T., 2006. Excludability and contribution: A laboratory study in team production, Working paper.

Croson, R. T. A., Marks, M. B., 2000. Returns in threshold public goods: a meta- and experimental analysis. *Experimental Economics* 2, 239–259.

Czap, H. J., Czap, N. V., Bonakdarian, E., 2010. Walk the talk? The effect of voting and excludability in public goods experiments. *Economics Research International* 2010, 15 pages, article ID 768546, doi:10.1155/2010/768546.

Davison, A., Hinkley, D., 1997. Bootstrap Methods and Their Application. Cambridge: Cambridge University Press.

Dawes, R. M., Orbell, J. M., Simmons, R. T., Van de Kragt, A. J. C., 1986. Organizing groups for collective action. *American Political Science Review* 80, 1171–1185.

Dickinson, D., 1998. The voluntary contributions mechanism with uncertain group payoffs. *Journal of Economic Behavior and Organization* 35, 517–533.

Donald, S. G., Lang, K., 2007. Inference with difference-in-differences and other panel data. *Review of Economics and Statistics* 89, 221–233.

El-Gamal, M. A., Grether, D. M., 1995. Are people Bayesian? Uncovering behavioral strategies. *Journal of the American Statistical Association* 90 (432), 1137–1145.

Esponda, I., Vespa, E., forthcoming. Hypothetical thinking and information extraction: Strategic voting in the laboratory, *American Economic Journal: Microeconomics*.

Eyster, E., Rabin, M., 2005. Cursed equilibrium. *Econometrica* 73, 1623–1672.

Fedderson, T., Pesendorfer, W., 1996. The swing voter's curse. *American Economic Review* 86, 408–424.

Fedderson, T., Pesendorfer, W., 1997. Voting behavior and information aggregation in elections with private information. *Econometrica* 65, 1029–1058.

Fedderson, T., Pesendorfer, W., 1998. Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. *American Political Science Review* 92, 23–35.

Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10, 171–178.

Fischbacher, U., Schudy, S., Teyssier, S., 2014. Heterogeneous reactions to heterogeneity in returns from public goods. *Social Choice and Welfare* 43, 195–217.

Fréchette, G. R., 2012. Session-effects in the laboratory. *Experimental Economics* 15, 485–498.

Gailmard, S., Palfrey, T. R., 2005. An experimental comparison of collective choice procedures for excludable public goods. *Journal of Public Economics* 89, 1361–1398.

Gangadharan, L., Nemes, V., 2009. Experimental analysis of risk and uncertainty in provisioning private and public goods. *Economic Inquiry* 47, 146–164.

Georganas, S., Healy, P. J., Weber, R. A., 2013. On the persistence of strategic sophistication, Working Paper.

Guarnaschelli, S., McKelvey, R. D., Palfrey, T. R., 2000. An experimental study of jury decision rules. *American Political Science Review* 94, 407–423.

Hermalin, B. E., 1998. Toward an economic theory of leadership: Leading by example. *American Economic Review* 88, 1188–1206.

Holt, C. A., Laury, S. K., 2002. Risk aversion and incentive effects. *American Economic Review* 92, 1644–1655.

Holt, C. A., Sherman, R., 1994. The loser's curse. *American Economic Review* 84, 642–652.

Isaac, M. R., Schmidtz, D., Walker, J. M., 1989. The assurance problem in the laboratory. *Public Choice* 62, 217–236.

Kagel, J. H., 1995. Auctions: a survey of experimental research. In: Kagel, J. H., Roth, A. E. (Eds.), The Handbook of Experimental Economics. Princeton, NJ: Princeton University Press, pp. 501–585.

Kagel, J. H., Levin, D., 1986. The winner's curse and public information in common value auctions. *American Economic Review* 76, 894–920.

Kagel, J. H., Levin, D., 2002. Common value auctions and the winner's curse. Princeton, NJ: Princeton University Press.

Kagel, J. H., Levin, D., Harstad, R. M., 1995. Comparative static effects of number of bidders and public information on behavior in second-price common value auctions. *International Journal of Game Theory* 24, 293–319.

Kocher, M., Sutter, M., Waldner, V., 2005. Exclusion from public goods as an incentive system – an experimental examination of different institutions, Working paper.

Levati, M. V., Morone, A., 2013. Voluntary contributions with risky and uncertain marginal returns: the importance of the parameter values. *Journal of Public Economic Theory* 15, 736–744.

Levati, M. V., Morone, A., Fiore, A., 2009. Voluntary contributions with imperfect information: An experimental study. *Public Choice* 138, 199–216.

Levin, D., Kagel, J. H., Richard, J. F., 1996. Revenue effects and information processing in English common value auctions. *American Economic Review* 86, 442–460.

Lind, B., Plott, C. R., 1991. The winner's curse: Experiments with buyers and with sellers. *American Economic Review* 81, 335–346.

Marks, M. B., Croson, R. T. A., 1999. The effect of incomplete information in a threshold public goods experiment. *Public Choice* 99, 103–118.

McCullagh, P., 2000. Resampling and exchangeable arrays. *Bernoulli* 6, 285–301.

Potters, J., Sefton, M., Vesterlund, L., 2005. After you — endogenous sequencing in voluntary contribution games. *Journal of Public Economics* 89, 1399–1419.

Potters, J., Sefton, M., Vesterlund, L., 2007. Leading-by-example and signaling in voluntary contribution games: an experimental study. *Economic Theory* 33, 169–182.

Rao, J. N. K., Scott, A. J., 1981. The analysis of categorical data from complex sample surveys: Chi-squared tests for goodness of fit and independence in two-way tables. *Journal of the American Statistical Association* 76, 221–230.

Rao, J. N. K., Scott, A. J., 1984. On chi-squared tests for multiway contingency tables with cell proportions estimated from survey data. *The Annals of Statistics* 12, 46–60.

Sribney, W. M., 1998. Two-way contingency tables for survey or clustered data. *Stata Technical Bulletin* 45, 33–49.

Stoddard, B., 2014a. Probabilistic production of a public good. *Economics Bulletin* 34, A194.

Stoddard, B., 2014b. Uncertainty in payoff-equivalent appropriation and provision games, Working Paper.

Stoddard, B., Walker, J., Williams, A., 2014. Allocating a voluntarily provided common-property resource: An experimental examination. *Journal of Economic Behavior and Organization* 101, 141–155.

Swope, K., 2002. An experimental investigation of excludable public goods. *Experimental Economics* 5, 209–222.

Thaler, R., 1988. Anomalies: the winner's curse. *Journal of Economic Perspectives* 2, 191–202.

Van de Kragt, A. J. C., Orbell, J. M., Dawes, R. M., 1983. The minimal contributing set as a solution to public goods problems. *American Political Science Review* 77, 112–122.

Vesterlund, L., 2003. The informational value of sequential fundraising. *Journal of Public Economics* 87, 627–657.

Appendix: Proofs

Proof of Lemma 1. In symmetric BNE, agents contribute with equal probability $p = \Pr(c(X) = 1)$. Given a signal x_i, agent i's expected payoff of contributing is given by:

$$U_i(x_i) = \sum_{l=k-1}^{n-1} \binom{n-1}{l} p^l (1-p)^{n-1-l} \left(\alpha x_i + \tfrac{(1-\alpha)l}{n-1} E[X|c(X)=1] + \tfrac{(1-\alpha)(n-1-l)}{n-1} E[X|c(X)=0] - w\right) \tag{A.1}$$

Differentiating with respect to the x_i, it is straightforward to verify that agent i's expected payoff is non-decreasing in x_i, and is strictly increasing whenever $p > 0$.

Suppose in symmetric BNE, $\{x_i \in [\underline{x}, \overline{x}] \mid U_i(x_i) > 0\} = \emptyset$. Then simply let $x_i^* = \overline{x} + 1$. Next, suppose in symmetric BNE, $\{x_i \in [\underline{x}, \overline{x}] \mid U_i(x_i) > 0\} = [\underline{x}, \overline{x}]$. Then let $x_i^* = \underline{x} - 1$. Now suppose in symmetric BNE, $\{x_i \in [\underline{x}, \overline{x}] \mid U_i(x_i) > 0\}$ is a non-empty, proper subset of $[\underline{x}, \overline{x}]$. Take $x_i \in \{x_i \in [\underline{x}, \overline{x}] \mid U_i(x_i) > 0\}$. Then since the expected payoff of contributing is non-decreasing in the signal, whenever $x_i' > x_i$, it must be that $U_i(x_i') \geq U_i(x_i) > 0$. Since $\{x_i \in [\underline{x}, \overline{x}] \mid U_i(x_i) > 0\}$ is bounded below by $\underline{x}$, it has an infimum. Let $x_i^* = \inf \{x_i \in [\underline{x}, \overline{x}] \mid U_i(x_i) > 0\}$. By continuity, $U_i(x_i^*) = 0$. By symmetry $x_i^* = x^*$ for each agent $i \in N$.

Given a signal $x_i'' \leq x^*$, it must be that $U_i(x_i'') \leq 0$, by definition of x^* and continuity of U_i.

□

Proof of Lemma 2. Take $r \neq i$. Then:

$$\begin{aligned}
&E\left[X_r \,\middle|\, \sum_{j\neq i} c^*(X_j) \geq k-1\right] \\
&= \Pr\left(\sum_{j\neq i} c^*(X_j) - c^*(X_r) < k-1 \,\middle|\, \sum_{j\neq i} c^*(X_j) \geq k-1\right) E[X_r | X_r \geq x^*] \qquad \text{(A.2)} \\
&+ \Pr\left(\sum_{j\neq i} c^*(X_j) - c^*(X_r) \geq k-1 \,\middle|\, \sum_{j\neq i} c^*(X_j) \geq k-1\right) E[X_r]
\end{aligned}$$

Intuitively, it is possible to partition the event where at least $k-1$ signals other than x_i exceed x^* into two cases. In the first, exactly $k-1$ signals other than x_i exceed x^*, one of which is x_r. In the second case, at least $k-1$ signals other than x_i and x_r exceed x^*, in which case the expectation of X_r is simply the prior expectation.

As x^* increases, $E[X_r|X_r \geq x^*]$ weakly increases by first-order stochastic dominance. In particular, $E[X_r|X_r \geq x^*] \geq E[X_r]$. Since the expression in equation A.2 is a convex combination of these two expectations, it suffices to show that the first probability in equation A.2 is also non-decreasing in x^*.

$$\begin{aligned}
&\Pr\left(\sum_{j\neq i} c^*(X_j) - c^*(X_r) < k-1 \,\middle|\, \sum_{j\neq i} c^*(X_j) \geq k-1\right) \\
&= \frac{\binom{n-2}{k-2}(1-F(x^*))^{k-1}F(x^*)^{n-k}}{\sum\limits_{l=k-1}^{n-1}\binom{n-1}{l}(1-F(x^*))^{l}F(x^*)^{n-1-l}} = \frac{\binom{n-2}{k-2}}{\sum\limits_{l=k-1}^{n-1}\binom{n-1}{l}\left(\frac{1-F(x^*)}{F(x^*)}\right)^{l-k+1}}
\end{aligned} \tag{A.3}$$

It is straightforward to verify that $\frac{1-F(x^*)}{F(x^*)}$ is non-increasing in x^*, which implies that the probability in equation A.3 is non-decreasing in x^*. Thus the expectation in equation A.2 is non-decreasing in x^*, which implies the result.

□

Proof of Proposition 1. Let w belong to the given interval. First, this interval is non-empty since:

$$\begin{aligned}
&\left(\alpha + \frac{(1-\alpha)(k-1)}{n-1}\right)\overline{x} + \frac{(1-\alpha)(n-k)}{n-1}E[X] - \alpha\underline{x} - (1-\alpha)E[X] \\
&= \alpha(\overline{x} - \underline{x}) + \frac{(1-\alpha)(k-1)}{n-1}(\overline{x} - E[X]) > 0
\end{aligned} \tag{A.4}$$

Morever, the interval is contained in $[\underline{x}, \overline{x}]$ since the lower bound is a convex combination of $\underline{x}$ and $E[X]$ while the upper bound is a convex combination of $\overline{x}$ and $E[X]$.

Restricting to $x^* < \overline{x}$, equation 4 can be rewritten as $H(x^*) = 0$ where:

$$\begin{aligned}
H(x^*) &= \frac{\sum\limits_{l=k-1}^{n-1}\binom{n-1}{l}(1-F(x^*))^{l}F(x^*)^{n-1-l}\left(\alpha x^* + \frac{(1-\alpha)l}{n-1}E[X|X\geq x^*] + \frac{(1-\alpha)(n-1-l)}{n-1}E[X|X<x^*] - w\right)}{\sum\limits_{l=k-1}^{n-1}\binom{n-1}{l}(1-F(x^*))^{l}F(x^*)^{n-1-l}} \\
&= \alpha x^* - w + \frac{\sum\limits_{l=k-1}^{n-1}\binom{n-1}{l}(1-F(x^*))^{l}F(x^*)^{n-1-l}\left(\frac{(1-\alpha)l}{n-1}E[X|X\geq x^*] + \frac{(1-\alpha)(n-1-l)}{n-1}E[X|X<x^*]\right)}{\sum\limits_{l=k-1}^{n-1}\binom{n-1}{l}(1-F(x^*))^{l}F(x^*)^{n-1-l}} \\
&= \alpha x^* - w + (1-\alpha)G_i(x^*)
\end{aligned} \tag{A.5}$$

Intuitively, $H(x^*)$ is the expected utility of contributing given a signal of x^* conditional on the public good being provided, where each other agents' strategy is contribution whenever their signal is at least x^*.

Notice that $H(x^*)$ is strictly increasing in x^* because the final term is non-decreasing in x^* by 2. Thus, $H(x^*)$ can cross zero at most once, guaranteeing that there is at most one interior equilibrium cutoff.

Consider the behavior of $H(x^*)$ as $x^* \to \overline{x}$. The key term of interest is the probability that exactly $k-1$ agents other than i contributed given that at least $k-1$ contributed.

$$\lim_{x^* \to \overline{x}} \Pr\left(\sum_{j \neq i} c(X_j) = k-1 \,\middle|\, \sum_{j \neq i} c(X_j) \geq k-1 \right) = \lim_{x^* \to \overline{x}} \frac{\binom{n-1}{k-1}(1-F(x^*))^{k-1}F(x^*)^{n-k}}{\sum_{l=k-1}^{n-1} \binom{n-1}{l}(1-F(x^*))^{l}F(x^*)^{n-1-l}} \tag{A.6}$$

$$= \lim_{x^* \to \overline{x}} \frac{1}{1+\sum_{l=k}^{n-1} \frac{(k-1)!(n-k)!}{l!(n-1-l)!}\left(\frac{1-F(x^*)}{F(x^*)}\right)^{l-k+1}} = 1$$

Since $\lim_{x^* \to \overline{x}} \frac{1-F(x^*)}{F(x^*)} = 0$. Therefore, taking the limit of $H(x^*)$ as $x^* \to \overline{x}$ yields the following:

$$\lim_{x^* \to \overline{x}} H(x^*) = \alpha \overline{x} - w + \frac{(1-\alpha)(k-1)}{n-1}\overline{x} + \frac{(1-\alpha)(n-k)}{n-1}E[X] \tag{A.7}$$

Which is positive if and only if $w < \left(\alpha + \frac{(1-\alpha)(k-1)}{n-1}\right)\overline{x} + \frac{(1-\alpha)(n-k)}{n-1}E[X]$. Now consider $x^* = \underline{x}$.

$$H(\underline{x}) = \alpha \underline{x} - w + (1-\alpha)E[X] \tag{A.8}$$

The right-hand side is negative if and only if $w > \alpha \underline{x} + (1-\alpha)E[X]$. Thus, since $H(x^*)$ is clearly continuous, there exists $x^* \in (\underline{x}, \overline{x})$ such that $H(x^*) = 0$, thus satisfying equation 4.

Now suppose w is not within the specified bounds. Since $H(x^*)$ is strictly increasing in x^*, it is either positive everywhere or negative everywhere. Thus, no symmetric BNE cutoff exists.

□

Proof of Corollary 1. $H(x^*)$ in equation A.5 is clearly decreasing in w. Therefore, when w increases, H becomes negative at the previous value of x^*. Since H is increasing in

x^*, the value of x^* at which this function is zero must be higher. Similarly, $G_i(x^*)$ in equation 2 is clearly increasing in k, which implies that $H(x^*)$ is also increasing in k. Therefore, when k increases, the value of x^* at which $H(x^*) = 0$ must decrease.

□

Proof of Proposition 2. As in Lemma 1, in symmetric χ-cursed equilibrium, the expected utility of contributing is non-decreasing in the signal. Thus by the same argument as in Lemma 1, attention can be restricted to cutoff equilibria. The equilibrium condition in equation 4 becomes:

$$\begin{aligned}&\sum_{l=k-1}^{n-1}\tbinom{n-1}{l}(1-F(x_\chi^*))^l F(x_\chi^*)^{n-1-l}\left[\alpha x_\chi^* + \tfrac{(1-\alpha)l}{n-1}\left(\chi E[X] + (1-\chi)E[X|X \geq x_\chi^*]\right)\right.\\ &\left.+\tfrac{(1-\alpha)(n-1-l)}{n-1}\left(\chi E[X] + (1-\chi)E[X|X < x_\chi^*]\right) - w\right] = 0\end{aligned} \tag{A.9}$$

As in the proof of Proposition 1, define a function $H_\chi(x_\chi^*)$ as follows:

$$\begin{aligned}H_\chi(x_\chi^*) &= \frac{\sum_{l=k-1}^{n-1}\binom{n-1}{l}(1-F(x_\chi^*))^l F(x_\chi^*)^{n-1-l}\left[\alpha x_\chi^* + \frac{(1-\alpha)l}{n-1}\left(\chi E[X]+(1-\chi)E[X|X\geq x_\chi^*]\right) + \frac{(1-\alpha)(n-1-l)}{n-1}\left(\chi E[X]+(1-\chi)E[X|X<x_\chi^*]\right) - w\right]}{\sum_{l=k-1}^{n-1}\binom{n-1}{l}(1-F(x_\chi^*))^l F(x_\chi^*)^{n-1-l}}\\ &= \alpha x_\chi^* - w + (1-\alpha)\chi E[X] + (1-\alpha)(1-\chi)G_i(x_\chi^*)\end{aligned} \tag{A.10}$$

As in Proposition 1, a zero of this function in $(\underline{x}, \overline{x})$ corresponds to an interior symmetric equilibrium cutoff. As in Proposition 1, it can be shown that $\lim_{x_\chi^* \to \overline{x}} H_\chi(x_\chi^*) > 0$ and $H_\chi(\underline{x}) < 0$ given the bounds on w. Thus, by continuity, an interior zero exists. By Lemma 2, $H_\chi(x_\chi^*)$ is strictly increasing in x_χ^*, and so it has at most one zero. Furthermore, as in Proposition 1, if w lies outside the given interval, $H_\chi(x_\chi^*)$ has no interior zero.

For $\chi = 1$ (fully-cursed equilibrium), equation A.9 becomes:

$$\left(\alpha x_1^* + (1-\alpha)E[X] - w\right)\sum_{l=k-1}^{n-1}\binom{n-1}{l}(1-F(x_1^*))^l F(x_1^*)^{n-1-l} = 0 \tag{A.11}$$

If $x_1^* < \overline{x}$, then solving for the cutoff yields equation 7.

□

Proof of Corollary 2. From equation A.10 it is clear that H_χ is non-increasing in χ, since by Lemma 2, $G_i(x_\chi^*) \geq G_i(\underline{x}) = E[X]$. Thus, as χ increases, the zero of H_χ (weakly) increases. The proof of the comparative statics with respect to w and k is the same as the proof of Corollary 1

□

Proof of Corollary 3. The bounds on w in Propositions 1 and 2 imply the result, as long as:

$$\begin{aligned}
&\left(\alpha+\tfrac{(1-\alpha)(k-1)}{n-1}\right)\overline{x}+\tfrac{(1-\alpha)(n-k)}{n-1}E[X]-\left(\alpha+\tfrac{(1-\chi)(1-\alpha)(k-1)}{n-1}\right)\overline{x}-\left(\tfrac{\chi(1-\alpha)(k-1)}{n-1}+\tfrac{(1-\alpha)(n-k)}{n-1}\right)E[X] \\
&=\tfrac{\chi(1-\alpha)(k-1)}{n-1}(\overline{x}-E[X])>0
\end{aligned} \tag{A.12}$$

Which holds if $\alpha<1$.

□

Online Appendix to Accompany: Cursed Beliefs with Common-Value Public Goods

Caleb A. Cox

Online Appendix A Supplementary Results

A.1 Repeated Trials and Learning

Figure OA.1 shows contribution rates over repeated trials of each of the four cost levels and for each of CVT, PVT, and AT. Recall that each of the four cost levels is encountered five times in each session, with the order randomized for each session. The BNE and fully-cursed benchmarks represent the expected contribution rates under each equilibrium concept, given the signals realized in the experiment. Few clear trends are apparent. Contributions do appear to decline in PVT with $w = 65$, which is somewhat reassuring given that no contribution should occur in that case. Only in CVT and PVT with $w = 35$ (where everyone should always contribute) does behavior seem to approach BNE. Overall, only in these few simple cases, where subjects should always or never contribute, do the data seem to suggest learning patterns.

Figure OA.2 shows the proportion of choices consistent with several equilibrium concepts over repeated trials. The particular partially-cursed equilibrium used here is for the maximum likelihood values of χ for CVT and AT (0.6 and 0.91 respectively). Again, there is little evidence of significant learning or convergence toward BNE, except in the simpler cases where subjects should always or never contribute. The clearest differences in consistency of contribution choices with the equilibrium concepts are in the cases of greatest contrast between BNE and cursed cutoffs (CVT with $w = 65$ and AT with $w = 35$). However, less differentiation is apparent for cost levels where there is less contrast between cutoffs under each equilibrium concept.[1]

I explore potential learning effects more formally in the logit regressions in Table

[1]Relatively few observations in these cases fall in the range where the equilibrium concepts make different predictions.

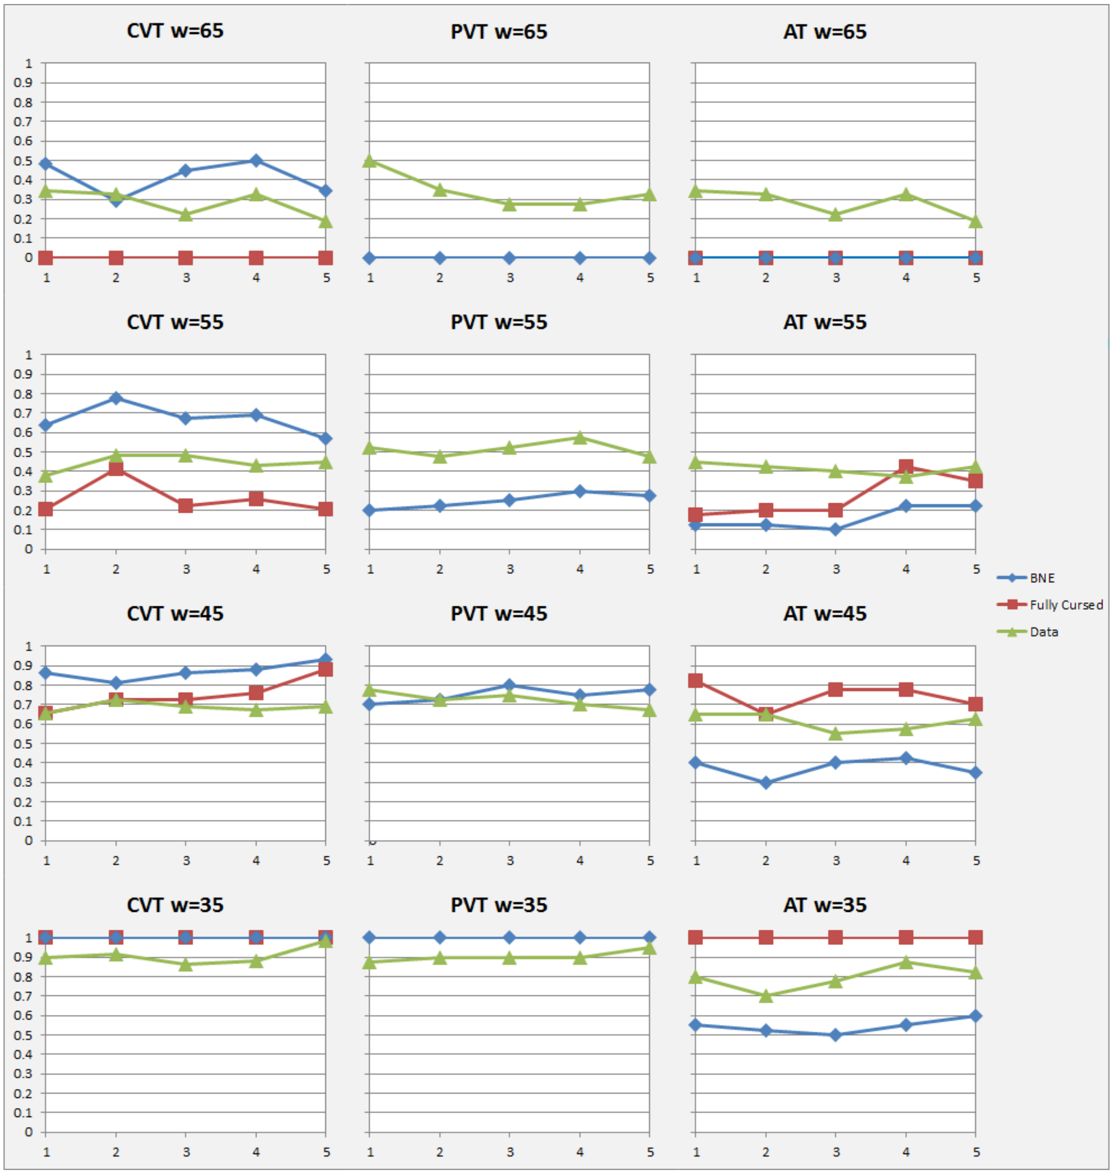

Figure OA.1: *Contribution Rates Over Repeated Trials*

OA.1. To study behavior within each treatment in more detail, I impose stronger assumptions about the nature of intra-session correlation than in logit regressions in the main body of the paper. I model session effects by including indicators for sessions and a (lagged) feedback variable equal to the contribution rate of other players in a given subject's group in the previous trial at the same cost level. I assume that conditional on session indicators and lagged feedback, observations of different subjects in the same session are independent. Standard errors are clustered by subject.

Figure OA.2: *Equilibrium Match Over Repeated Trials*

Consistent with the previous graphs, the results in Table OA.1 show little evidence of consistent learning patterns. In CVT at the lowest cost level, the is an upward trend toward the equilibrium prediction of full contribution. However, there is also a (marginally significant) upward trend in contribution in AT at the lowest cost level, moving farther away from the BNE prediction. No other significant trends are apparent. Thus it appears that very little learning takes place, or if it does, that it is very slow.

CVT

Variable	w=65	w=55	w=45	w=35
signal	1.035***	1.033***	1.031***	1.008
	(0.009)	(0.007)	(0.008)	(0.012)
round	0.968	1.018	0.991	1.087**
	(0.031)	(0.033)	(0.035)	(0.042)
feedback	0.782	0.897	0.448	0.476
	(0.523)	(0.473)	(0.316)	(0.629)

PVT

Variable	w=65	w=55	w=45	w=35
signal	1.034***	1.039***	1.029***	1.009
	(0.011)	(0.010)	(0.011)	(0.011)
round	1.010	1.020	0.946	1.036
	(0.050)	(0.049)	(0.049)	(0.049)
feedback	2.433	0.301*	1.454	2.877
	(2.356)	(0.205)	(1.349)	(4.376)

AT

Variable	w=65	w=55	w=45	w=35
signal	1.017*	1.041***	1.026***	1.023**
	(0.010)	(0.008)	(0.007)	(0.009)
round	1.149**	0.954	0.998	1.092*
	(0.069)	(0.041)	(0.034)	(0.055)
feedback	1.479	0.296	1.470	2.632
	(1.179)	(0.241)	(0.848)	(2.819)

Table OA.1: *Logit regression results for CVT, PVT, and AT and each cost level. The dependent variable is an indicator for contribution. Session indicators included but not reported. Odds ratios are reported with standard errors in parentheses, clustered at the subject level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.*

A.2 Over-Contribution in PVT

As shown in the main body of the paper, there is a puzzling tendency for subjects to over-contribute in PVT (and AT) with $w = 65$. Even with a signal of 100, the expected value of the group project is no greater than 60 tokens, and thus subjects should never contribute in PVT when the cost is 65 tokens. Interestingly, subjects do not frequently make similar mistakes in PVT with $w = 35$, where contribution is always optimal for any signal. Thus, there is an asymmetry in behavior between cases in which subjects should always or never contribute.

This finding is not driven by a small subset of subjects. Rather, a majority of subjects contributed at least once in this case. This behavior is clearly not driven by cursedness, because in PVT full or partial cursedness yields the same prediction as BNE. Rationalizing this behavior through risk preferences would require many subjects to be implausibly risk-loving, with coefficients of relative risk aversion less than -5. This behavior might represent some form of altruism, as subjects may simply view contributing as a pro-social act. However, such motivations would be misguided, since the group project is a bad bet for other players as well.

Another possibility is that subjects are simply bad at calculating expected values. To investigate this possibility, a surprise bonus question was added at the end of the second session of PVT and the second session of AT. In this question, subjects were asked to calculate the expected value of the group project given a signal of 100. Answers within plus or minus 5 of the correct answer (60) were rewarded with a $1 bonus payment on top of any earnings from the main part of the experiment. If subjects can correctly perform this calculation, they should see that contributing at a cost of 65 is never optimal. Of the forty subjects in these two sessions, 45% got the answer exactly right (which was also the modal response), and 65% answered within plus or minus 5. I ran logistic regressions similar those in Table OA.1 at cost level $w = 65$ in these two sessions, using an indicator for an exactly correct answer as an explanatory variable. As shown in Table OA.2, I did not find any significant correlation between a over-contribution and the answer given on the bonus question. This negative result is robust to alternative specifications such as using an indicator for an answer with plus or minus 5 or using the

Variable	**PVT w=65**	**AT w=65**
signal	1.086***	1.014
	(0.025)	(0.012)
round	0.991	1.134*
	(0.083)	(0.076)
feedback	4.833	2.615
	(7.574)	(2.710)
Correct	0.683	1.291
	(0.546)	(1.192)

Table OA.2: *Logit regression results for PVT and AT at the highest cost level. The dependent variable is an indicator for contribution. Odds ratios are reported with standard errors in parentheses, clustered at the subject level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.*

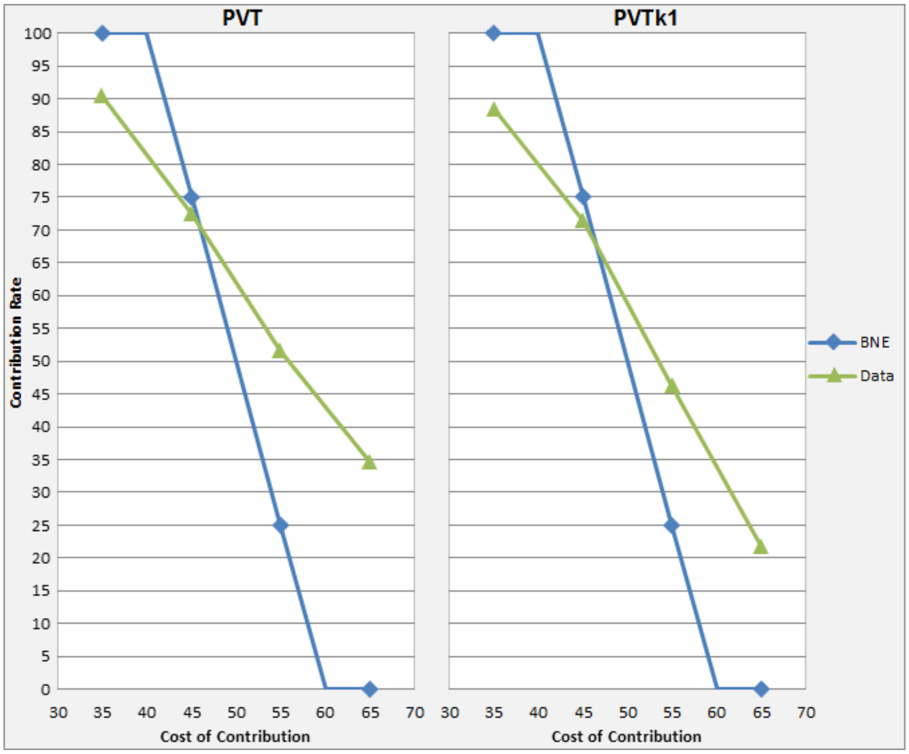

Figure OA.3: *Contribution in PVT v. PVTk1*

actual reported expected value. Thus, errors in expected value calculation do not appear to be an important reason for the observed over-contribution.

As discussed in the main body of the paper, I added the PVTk1 treatment, altering PVT into an individual choice problem (while keeping the framing as close as possible to the original PVT treatment). This treatment is designed to test whether weak incentives due to infrequent provision of the group project might drive over-contribution in this case. As shown in Table 1 in the main body of the paper, there is some evidence in support of this explanation, with lower contribution in PVTk1 relative to PVT at cost level 65. Aggregate contribution rates in PVT and PVTk1 are displayed graphically in Figure OA.3, showing visual evidence consistent with the regression result in the main body of the paper. Figure OA.4 compares contribution in PVT and PVTk1 at each cost level over repeated trials. Again, the two treatments appear somewhat different at cost level 65, but similar for lower cost levels.

Beyond differences in the aggregate rates of contribution in PVT and PVTk1 at the highest cost level, there may also differences in learning. Table OA.3 shows regres-

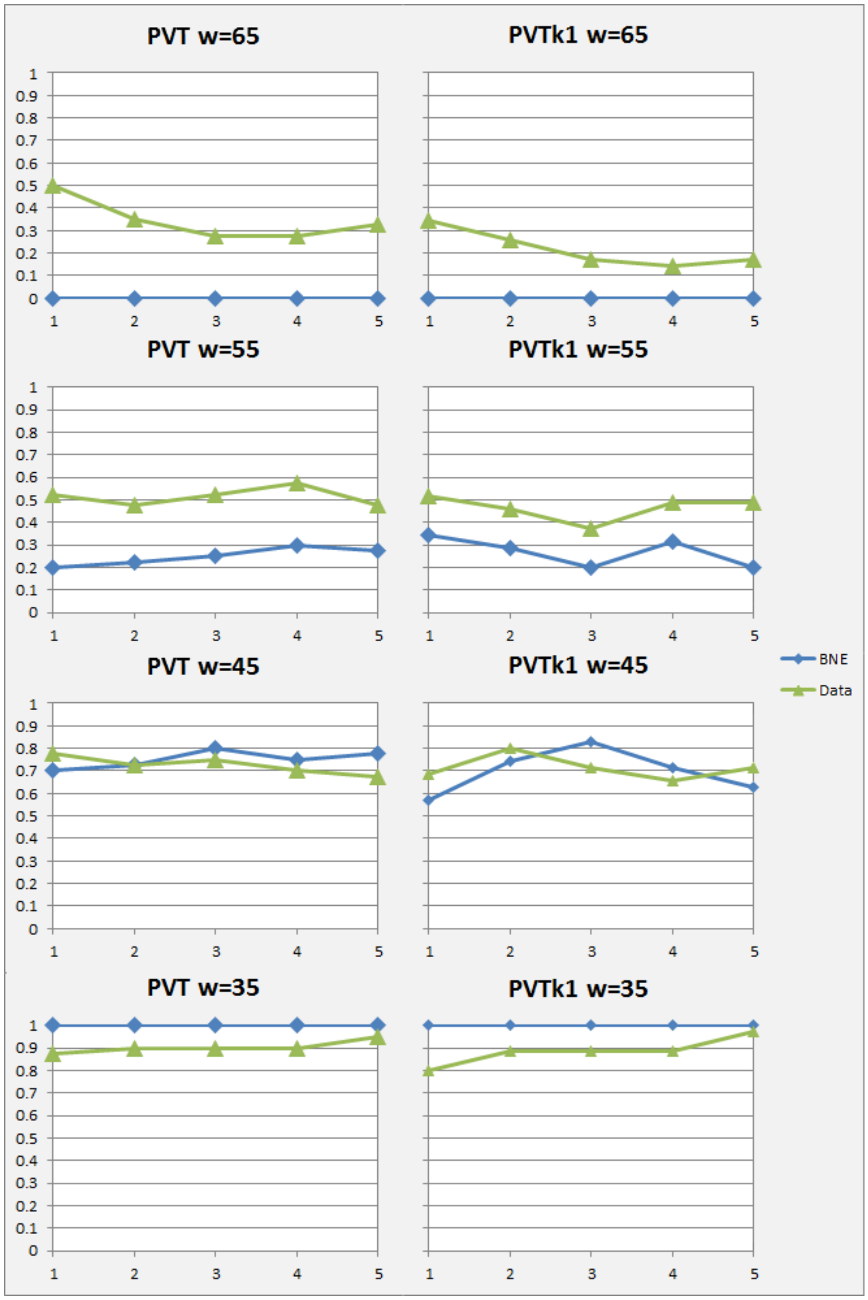

Figure OA.4: *Contribution in PVT v. PVTk1 over repeated trials.*

sions similar to those in Table OA.1, examining trends over several rounds of play in PVTk1. For cost level 65, there is a significant reduction in contributions for PVTk1. No such trend is apparent in PVT, as shown in Table OA.1. Comparing the coefficients on

PVTk1				
Variable	**w=65**	**w=55**	**w=45**	**w=35**
signal	1.046***	1.036***	1.041***	1.021
	(0.013)	(0.010)	(0.010)	(0.017)
round	0.851**	1.004	0.964	1.134*
	(0.059)	(0.049)	(0.036)	(0.077)
feedback	0.136	0.919	0.061**	0.747
	(0.181)	(0.732)	(0.077)	(1.347)

Table OA.3: *Logit regression results for PVTk1 and each cost level. The dependent variable is an indicator for contribution. Session indicators included but not reported. Odds ratios are reported with standard errors in parentheses, clustered at the subject level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.*

round of play in PVT and PVTk1, the difference is significiant (p-value=0.030).[2] Thus, there is some evidence that steepening incentives may help subjects learn not to over-contribute in this setting. However, even in the individual choice problem of PVTk1, subjects still choose the group project too often for $w = 65$. It is possible that some subjects simply enjoy gambling in small amounts, or use idiosyncratic heuristics leading to over-contribution.

A.3 Predicted Probability of Contribution as a Function of the Signal

Figure OA.5 shows predicted probabilities of contribution for CVT, PVT, and AT at each cost level, including 95% confidence intervals. These predicted probabilities are derived from the logit regression estimates in Table OA.1. Explanatory variables other than signal are fixed at their mean values. The graphs suggest that the probability of contribution increases smoothly with the signal rather than following a step function as suggested by the theoretical benchmarks. It is possible that subjects use probabilistic strategies, with contribution more likely as the signal increases. Another possibility, as explored in the main body of the paper, is that subjects' cutoff strategies are heterogeneous. Such behavior could lead to aggregate contribution that increases gradually in the signal rather than jumping sharply at a common cutoff.

[2]For this comparison, I nest the two models, interact all explanatory variables with the indicator for PVTk1, and test the significance of the interaction between the treatment indicator and round of play.

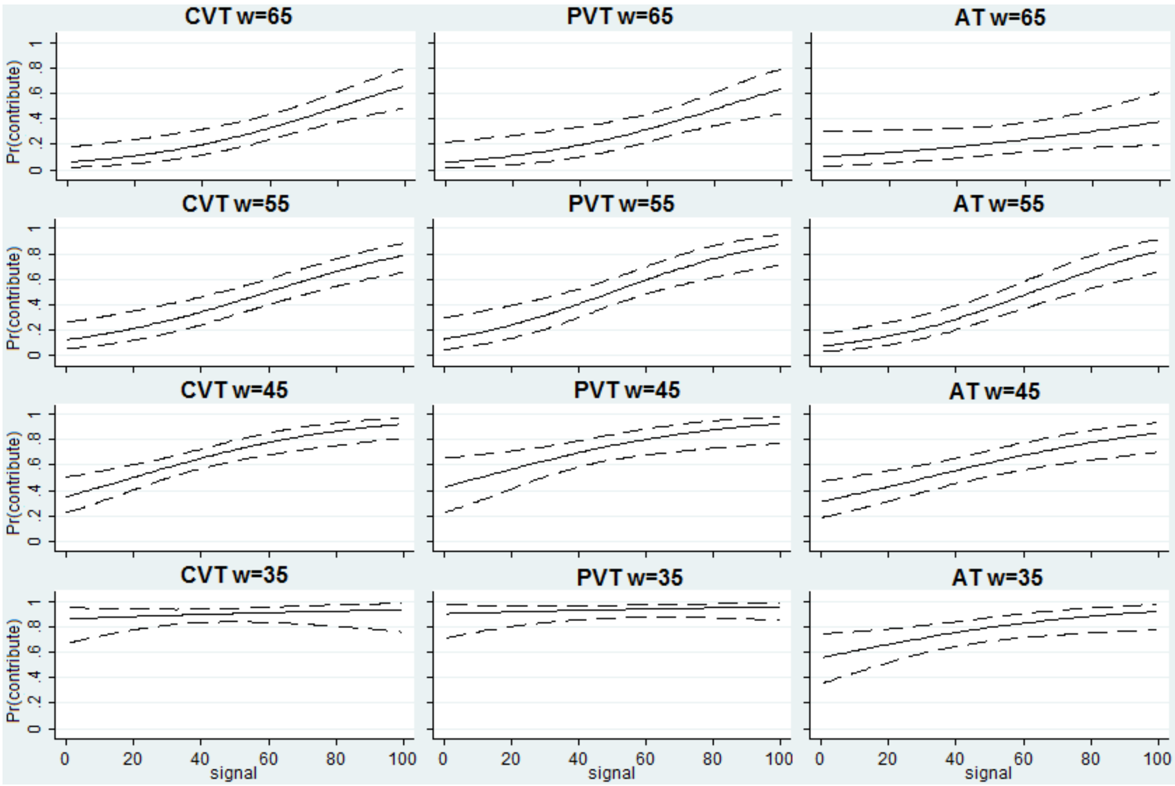

Figure OA.5: *Predicted probability of contribution as a function of the signal for CVT, PVT, and AT at each cost level, from the logit regression estimates in Table OA.1. Shown with 95% confidence intervals.*

A.4 Cognitive / Quantitative Ability

In Table 2 in the main body of the paper, I report estimated strategy proportions for CVT and AT. An additional question of interest is whether cognitive and quantitative ability are correlated with proper belief conditioning. Subjects gave consent to access academic records including GPA, ACT/SAT scores, and academic major. Academic major is used as a proxy for quantitative ability: I classify mathematics, statistics, engineering, natural and physical sciences, computer science, economics, finance, and accounting as quantitative majors. I have run a number of regressions checking whether consistency with BNE correlates with ACT/SAT percentile, GPA, and quantitative major. However, as shown in Table OA.4 the results have been negative. Perhaps surprisingly, there is little correlation between the measures of cognitive and quantitative ability, and thus multicollinearity does not seem to be a major issue in this case.

Variable	CVT	AT
ACT/SAT Percentile	0.0017	0.0010
	(0.0014)	(0.0016)
GPA	0.0328	0.0513
	(0.0834)	(0.1194)
Quant	-0.0177	0.0600
	(0.0726)	(0.1014)

Table OA.4: *Linear regressions results. The dependent variable is the proportion of choices consistent with BNE. Robust standard errors shown in parentheses. Session indicators are included but not reported.*

A.5 Efficiency

In addition to contribution decisions, efficiency is of some interest. Figures OA.6 and OA.7 show the average per person net gains in CVT and PVT at each cost of contribution for the signals realized in the experiment. As there is no public good of any kind in AT, efficiency in this case is of less interest and is omitted. The first-best efficiency benchmark shows the net gain if provision occurs if and only if provision is efficient (with the most efficient contributing set in PVT). The second-best benchmark shows the net gain if a benevolent social planner were to enforce a symmetric contribution cutoff to maximize the expected total surplus. While efficiency under BNE is somewhat lower than second-best in CVT, it is quite close. Efficiency in the data falls below the BNE benchmark in CVT, particularly for cost levels 45 and 55. In PVT, efficiency in the data is actually higher than the BNE benchmark at cost level 45, but falls short otherwise, including small negative net returns for cost levels 55 and 65.

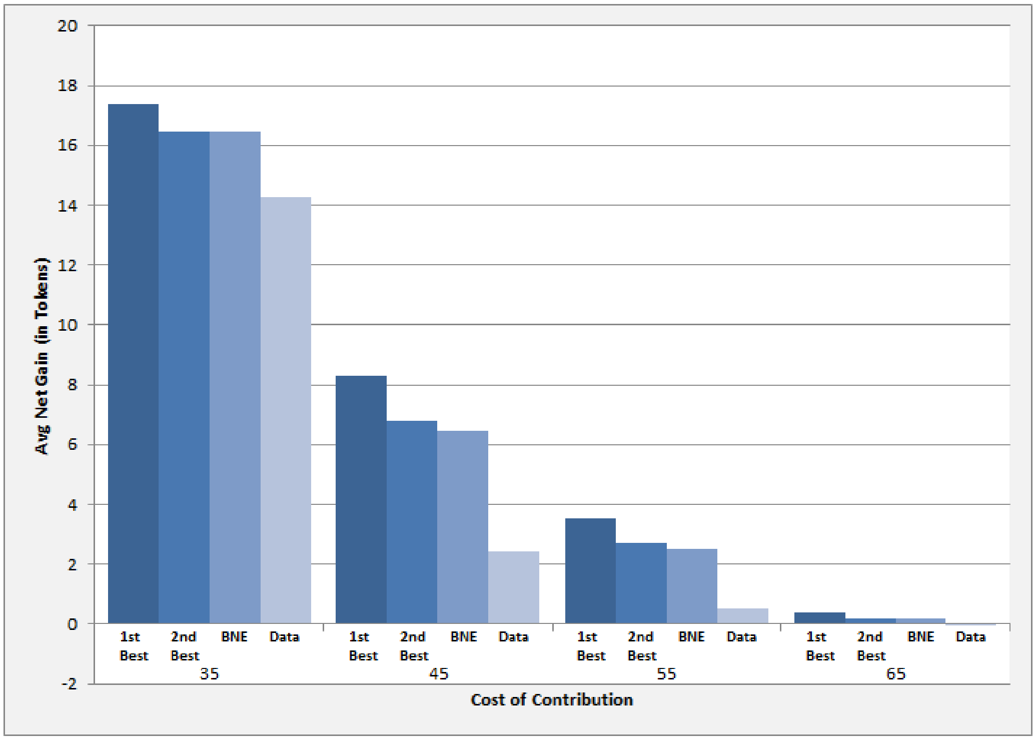

Figure OA.6: *Efficiency in CVT*

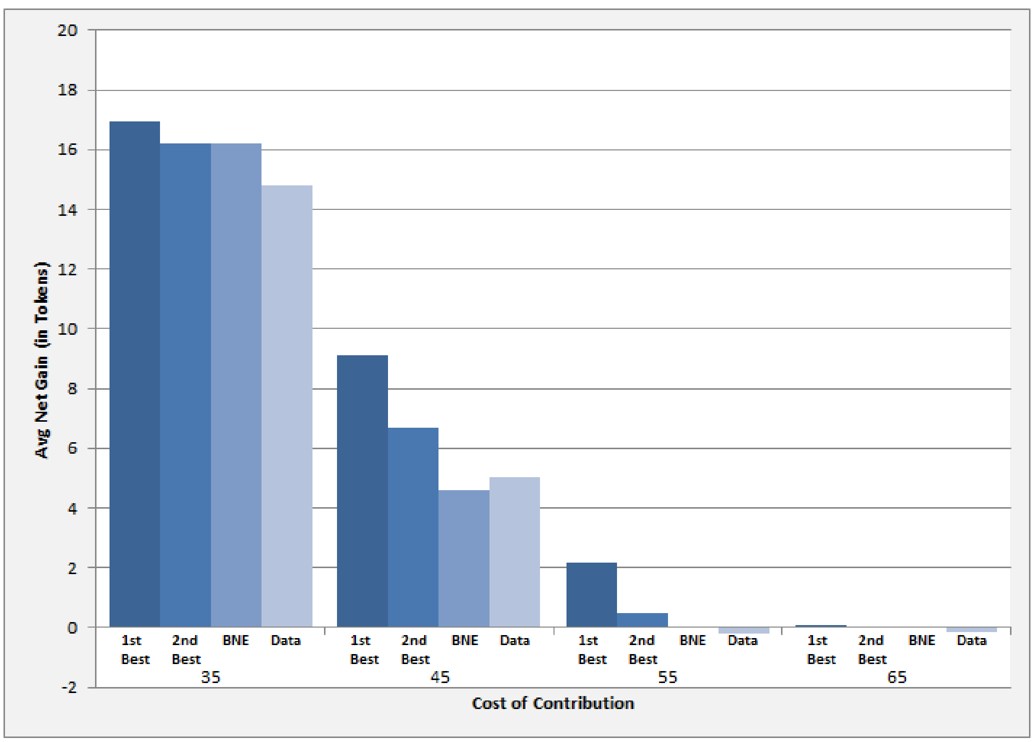

Figure OA.7: *Efficiency in PVT*

Online Appendix B Experimental Instructions

Experiment Overview

This is an experiment in the economics of decision-making. In this experiment you will make a series of choices, each of which may earn you money. The amount of money you earn will depend on the decisions you make and on the decisions of others. If you listen carefully and make good decisions, you could earn a considerable amount of money that will be paid to you in cash at the end of the experiment.

Ground Rules

Please make all decisions independently; do not communicate with others (in the room or outside the room) in any way during the experiment. This means no talking, no cell phone usage, no texting, no internet chatting, etc. Please do not attempt to use any software other than the experiment software provided. Failure to comply with these rules will lead to dismissal from the experiment.

Instructions (CVT)

During the experiment, participants earn tokens. All participants will be paid based on the number of tokens they earn. Each token is worth 20 cents, or $1 for every 5 tokens.

The experiment consists of twenty rounds. At the start of the first round, you will be randomly and anonymously matched into groups of five. At the start of each later round, you will be randomly and anonymously re-matched into new groups of five, so that your group changes every round and you never learn the identities of the other group members in any round.

In each round, you will choose how to allocate some number of tokens, which we will call T. Everyone in your group in any given round will allocate the same number of tokens, so that T is the same for everyone in your group and is known to everyone in your group. However, T may change from round to round.

The T tokens may be allocated to a private account or to a group project. If you choose to allocate T tokens to a private account, then you get T tokens for the round. The details of the group project are as follows.

In each round, a random number will be selected by the computer from a uniform distribution between 0 and 100. We will call this random number your signal. Each other member of your group will also get a signal randomly selected by the computer from the same distribution. We will call the signals of the five group members S1, S2, S3, S4, and S5. All signals are drawn independently. During each round, you will not observe the signals of the other members. Similarly, other members of the group will not observe any signal other than their own.

If you choose to allocate T tokens to the group project, and if at least three other members of your group also choose to allocate T tokens to the group project, then you get a number of tokens equal to the average of the signals of all five members of your group for the round. We will call the average of the signals of the five members of your group the value of the group project, or V, which is given by:

$$V = \frac{S1 + S2 + S3 + S4 + S5}{5}$$

So, for example, if your signal is 50 and the other members of your group get signals of 25, 40, 62, and 86, then the average of all five signals is:

$$V = \frac{50 + 25 + 40 + 62 + 86}{5} = 52.6$$

Thus, in this case, if you chose to allocate T tokens to the group project and at least three other members of your group also chose to allocate T tokens to the group project, then you would get 52.6 tokens for that round.

If you choose to allocate T tokens to the group project, but *less* than three other members of your group also choose to allocate T tokens to the group project, then you get T tokens for that round. In other words, if less than four of the five members of your group (including yourself) choose to allocate T tokens to the group project, then all tokens are automatically reallocated to private accounts and everyone in your group gets T tokens.

After all members of your group have made their choice, you will learn the value of the group project (V) and your earnings in tokens for the round. You will also learn the signals observed by the other members of your group (listed from highest to lowest) and how they allocated T tokens (to the group project or private account).

The following summarizes your choice in each round. If you choose to allocate T

tokens to a private account, then you get T tokens. If you choose to allocate T tokens to the group project, and if at least three other members of your group also choose to allocate T tokens to the group project, then you get V tokens, where V is the average of the signals of the five members of your group. If you choose to allocate T tokens to the group project, but *less* than three other members of your group also choose to allocate T tokens to the group project, then you get T tokens.

Remember that you will be randomly and anonymously re-matched into new groups of five at the start of each round. Also remember that T is the same for every member of your group. However, signals are randomly and independently drawn for each member of your group.

Of the twenty rounds, one will be randomly selected for payment. All participants will be paid their earnings in dollars for the randomly selected round, plus a $5.00 show-up payment. You will not find out which round you will be paid for until the end of the experiment, so you should treat each round as something for which you might get paid. You will not be paid for the rounds that are not randomly selected for payment.

Are there any questions before we begin the experiment?

Instructions (AT)

During the experiment, participants earn tokens. All participants will be paid based on the number of tokens they earn. Each token is worth 20 cents, or $1 for every 5 tokens.

The experiment consists of twenty rounds. At the start of the first round, you will be randomly and anonymously matched into groups of five. At the start of each later round, you will be randomly and anonymously re-matched into new groups of five, so that your group changes every round and you never learn the identities of the other group members in any round.

In each round, you will choose how to allocate some number of tokens, which we will call T. Everyone in your group in any given round will allocate the same number of tokens, so that T is the same for everyone in your group and is known to everyone in your group. However, T may change from round to round.

The T tokens may be allocated to a private account or to a group project. If you choose to allocate T tokens to a private account, then you get T tokens for the round. The details of the group project are as follows.

In each round, a random number will be selected by the computer from a uniform distribution between 0 and 100. We will call this random number your signal. Each other member of your group will also get a signal randomly selected by the computer from the same distribution. We will call the signals of the five group members S1, S2, S3, S4, and S5. All signals are drawn independently. During each round, you will not observe the signals of the other members. Similarly, other members of the group will not observe any signal other than their own.

If you choose to allocate T tokens to the group project, and if no more than one other member of your group also chooses to allocate T tokens to the group project, then you get a number of tokens equal to the average of the signals of all five members of your group for the round. We will call the average of the signals of the five members of your group the value of the group project, or V, which is given by:

$$V = \frac{S1+S2+S3+S4+S5}{5}$$

So, for example, if your signal is 50 and the other members of your group get signals of 25, 40, 62, and 86, then the average of all five signals is:

$$V = \frac{50+25+40+62+86}{5} = 52.6$$

Thus, in this case, if you chose to allocate T tokens to the group project and if no more than one other member of your group also chose to allocate T tokens to the group project, then you would get 52.6 tokens for that round.

If you choose to allocate T tokens to the group project, but *more* than one other member of your group also chooses to allocate T tokens to the group project, then you get T tokens for that round. In other words, if more than two of the five members of your group (including yourself) choose to allocate T tokens to the group project, then all tokens are automatically reallocated to private accounts and everyone in your group gets T tokens.

After all members of your group have made their choice, you will learn the value of the group project (V) and your earnings in tokens for the round. You will also learn the signals observed by the other members of your group (listed from highest to lowest) and how they allocated T tokens (to the group project or private account).

The following summarizes your choice in each round. If you choose to allocate T

tokens to a private account, then you get T tokens. If you choose to allocate T tokens to the group project, and if no more than one other member of your group also chooses to allocate T tokens to the group project, then you get V tokens, where V is the average of the signals of the five members of your group. If you choose to allocate T tokens to the group project, but *more* than one other member of your group also chooses to allocate T tokens to the group project, then you get T tokens.

Remember that you will be randomly and anonymously re-matched into new groups of five at the start of each round. Also remember that T is the same for every member of your group. However, signals are randomly and independently drawn for each member of your group.

Of the twenty rounds, one will be randomly selected for payment. All participants will be paid their earnings in dollars for the randomly selected round, plus a \$5.00 show-up payment. You will not find out which round you will be paid for until the end of the experiment, so you should treat each round as something for which you might get paid. You will not be paid for the rounds that are not randomly selected for payment.

Are there any questions before we begin the experiment?

Instructions (PVT)

During the experiment, participants earn tokens. All participants will be paid based on the number of tokens they earn. Each token is worth 20 cents, or \$1 for every 5 tokens.

The experiment consists of twenty rounds. At the start of the first round, you will be randomly and anonymously matched into groups of five. At the start of each later round, you will be randomly and anonymously re-matched into new groups of five, so that your group changes every round and you never learn the identities of the other group members in any round.

In each round, you will choose how to allocate some number of tokens, which we will call T. Everyone in your group in any given round will allocate the same number of tokens, so that T is the same for everyone in your group and is known to everyone in your group. However, T may change from round to round.

The T tokens may be allocated to a private account or to a group project. If you choose to allocate T tokens to a private account, then you get T tokens for the round. The details of the group project are as follows.

In each round, a random number will be selected by the computer from a uniform distribution between 0 and 100. We will call this random number your signal and label it S. Each other member of your group will also get a signal randomly selected by the computer from the same distribution. All signals are drawn independently. During each round, you will not observe the signals of the other members. Similarly, other members of the group will not observe any signal other than their own.

Furthermore, in each round, four unobserved random number will be selected for you by the computer from a uniform distribution between 0 and 100. We will label these four random numbers R1, R2, R3, and R4. For each other member of your group, there will also be four unobserved numbers randomly selected by the computer from the same distribution. All of these random numbers are drawn independently of each other and independently of your signal and the signals of others in your group. You will not observe any of these random numbers, and neither will any other member of your group.

If you choose to allocate T tokens to the group project, and if at least three other members of your group also choose to allocate T tokens to the group project, then you get a number of tokens equal to the average of your signal and the four unobserved random numbers selected for you by the computer for the round. We will call this average your value for the group project, or V, which is given by:

$$V = \frac{S + R1 + R2 + R3 + R4}{5}$$

So, for example, if your signal is 50 and the four unobserved random numbers are 25, 40, 62, and 86, then the average is:

$$V = \frac{50 + 25 + 40 + 62 + 86}{5} = 52.6$$

Thus, in this case, if you chose to allocate T tokens to the group project and at least three other members of your group also chose to allocate T tokens to the group project, then you would get 52.6 tokens for that round.

If you choose to allocate T tokens to the group project, but *less* than three other members of your group also choose to allocate T tokens to the group project, then you get T tokens for that round. In other words, if less than four of the five members of your group (including yourself) choose to allocate T tokens to the group project, then all tokens are automatically reallocated to private accounts and everyone in your group gets T tokens.

After all members of your group have made their choice, you will learn your value for the group project (V) and your earnings in tokens for the round. You will also learn the signals observed by the other members of your group (listed from highest to lowest) and how they allocated T tokens (to the group project or private account).

The following summarizes your choice in each round. If you choose to allocate T tokens to a private account, then you get T tokens. If you choose to allocate T tokens to the group project, and if at least three other members of your group also choose to allocate T tokens to the group project, then you get V tokens, where V is the average of your signal, R1, R2, R3, and R4. If you choose to allocate T tokens to the group project, but *less* than three other members of your group also choose to allocate T tokens to the group project, then you get T tokens.

Remember that you will be randomly and anonymously re-matched into new groups of five at the start of each round. Also remember that T is the same for every member of your group. However, signals and unobserved random numbers are randomly and independently drawn for each member of your group.

Of the twenty rounds, one will be randomly selected for payment. All participants will be paid their earnings in dollars for the randomly selected round, plus a $5.00 show-up payment. You will not find out which round you will be paid for until the end of the experiment, so you should treat each round as something for which you might get paid. You will not be paid for the rounds that are not randomly selected for payment.

Are there any questions before we begin the experiment?

Instructions (PVTk1)

During the experiment, participants earn tokens. All participants will be paid based on the number of tokens they earn. Each token is worth 20 cents, or $1 for every 5 tokens.

The experiment consists of twenty rounds. At the start of the first round, you will be randomly and anonymously matched into groups of five. At the start of each later round, you will be randomly and anonymously re-matched into new groups of five, so that your group changes every round and you never learn the identities of the other group members in any round.

In each round, you will choose how to allocate some number of tokens, which we will call T. Everyone in your group in any given round will allocate the same number of tokens, so that T is the same for everyone in your group and is known to everyone in your group. However, T may change from round to round.

The T tokens may be allocated to a private account or to a group project. If you choose to allocate T tokens to a private account, then you get T tokens for the round. The details of the group project are as follows.

In each round, a random number will be selected by the computer from a uniform distribution between 0 and 100. We will call this random number your signal and label it S. Each other member of your group will also get a signal randomly selected by the computer from the same distribution. All signals are drawn independently. During each round, you will not observe the signals of the other members. Similarly, other members of the group will not observe any signal other than their own.

Furthermore, in each round, four unobserved random numbers will be selected for you by the computer from a uniform distribution between 0 and 100. We will label these four random numbers R1, R2, R3, and R4. For each other member of your group, there will also be four unobserved numbers randomly selected by the computer from the same distribution. All of these random numbers are drawn independently of each other and independently of your signal and the signals of others in your group. You will not observe any of these random numbers, and neither will any other member of your group.

If you choose to allocate T tokens to the group project, then you get a number of tokens equal to the average of your signal and the four unobserved random numbers selected for you by the computer for the round. We will call this average your value for the group project, or V, which is given by:

$$V = \frac{S + R1 + R2 + R3 + R4}{5}$$

So, for example, if your signal is 50 and the four unobserved random numbers are 25, 40, 62, and 86, then the average is:

$$V = \frac{50 + 25 + 40 + 62 + 86}{5} = 52.6$$

Thus, in this case, if you chose to allocate T tokens to the group project, then you would get 52.6 tokens for that round.

After all members of your group have made their choice, you will learn your value for the group project (V) and your earnings in tokens for the round. You will also learn the signals observed by the other members of your group (listed from highest to lowest) and how they allocated T tokens (to the group project or private account).

The following summarizes your choice in each round. If you choose to allocate T tokens to a private account, then you get T tokens. If you choose to allocate T tokens to the group project, then you get V tokens, where V is the average of your signal, R1, R2, R3, and R4.

Remember that you will be randomly and anonymously re-matched into new groups of five at the start of each round. Also remember that T is the same for every member of your group. However, signals and unobserved random numbers are randomly and independently drawn for each member of your group.

Of the twenty rounds, one will be randomly selected for payment. All participants will be paid their earnings in dollars for the randomly selected round, plus a $5.00 show-up payment. You will not find out which round you will be paid for until the end of the experiment, so you should treat each round as something for which you might get paid. You will not be paid for the rounds that are not randomly selected for payment.

Are there any questions before we begin the experiment?