Submission Number: PET11-11-00299

Information revelation in a public goods game: an experimental study of the role of charity ratings organizations

Daniel Jones
University of Pittsburgh

Abstract

It is becoming increasingly easy for donors to obtain information about the quality of charitable organizations, primarily through third-party organizations that evaluate nonprofits and publish ratings on the Internet. What is the impact of these rating organizations? I explore the role of third-party information providers in public goods games. In an experiment, participants are either assigned the role of "donor" or "informer." Donors choose whether or not to contribute to a public good but initially face uncertainty about the returns to contributing. The informer decides how much information should be available to donors. I test the theoretical prediction that the informer chooses to limit information to an indication of whether or not the public good is socially efficient rather than revealing the precise quality, and that contributions increase as a result.

Submitted: March 15, 2011.

Information revelation in a public goods game: An experimental study of the role of charity ratings organizations

Daniel Jones

March 15, 2011

Abstract. It is becoming increasingly easy for donors to obtain information about the quality of charitable organizations, primarily through third-party organizations that evaluate nonprofits and publish ratings on the Internet. What is the impact of these rating organizations? I explore the role of third-party information providers in public goods games. In an experiment, participants are either assigned the role of "donor" or "informer." Donors choose whether or not to contribute to a public good but initially face uncertainty about the returns to contributing. The informer decides how much information should be available to donors. I test the theoretical prediction that the informer chooses to limit information to an indication of whether or not the public good is socially efficient rather than revealing the precise quality, and that contributions increase as a result.

1 Introduction

Following the Asian tsunami of 2004, there was an outpouring of donations from the United States. At the same time, there was a dramatic increase in traffic to the websites of "charity watchdog" groups (Charity Navigator and the Better Business Bureau's Wise Giving Alliance are two prominent examples of such groups) – organizations that evaluate charities and make their findings available to donors. This demonstrates both the uncertainty that donors initially face when making charitable contribution decisions and the growing influence of the charity watchdog groups. However, while it is clear that watchdog groups' ratings influence giving (Gordon, Knock, and Neely, 2009; Sloan, 2009; Grant, 2010), it is less clear how they impact society more generally. They are frequently criticized from within the philanthropy community for excessively simplifying the rather complex matter of the quality of a charity (Lowell, Trelstad, and Meehan, 2005), providing donors information in the form of a "pass/fail" grade or a rating of one to four stars, for instance.

This leads one to ask: when voluntary contributions must be relied upon to provide a public good, what is the "right" level of information to provide? Can society achieve this level of information, perhaps through institutions like charity watchdog groups? These are the questions I address in this paper. I focus in particular on the role of a third-party information provider in a public goods game and argue that (1) social welfare can be improved as a result of the existence of information providers and (2) the mechanism through which this works may be the information provider's ability to *restrict* information. The reason for this is that, by limiting information, donors may have enough information to observe when a public good is socially efficient, but not enough information to incentivize free-riding. Therefore, partial information may lead donors out of this classic social dilemma and into the most efficient outcome. Thus, in this paper I suggest that, due to the unique incentives present in public goods situations, it may be precisely because charity watchdog groups simplify or limit the information they provide that they are beneficial.

It should of course be noted that restricting information increases welfare most effectively in a world wherein free-riding is a serious problem. As has been shown repeatedly, both empirically and in laboratory experiments, individuals are significantly more likely to contribute to public goods than naive theory might predict. Thus, although the theoretical possibility that less information is preferable is clear, the empirical reality may not be. It is with this in mind that I conduct a laboratory experiment to explore these ideas.

Participants play a simple public goods game as often seen in the literature, but with the addition of a third-party "informer" who, in advance of any contribution decisions, decides how much "donors" know about the quality of the public good. For the reasons discussed above, "partial" information is preferable in theory, even as compared to a world in which all donors have complete information. Therefore, I test which level of information maximizes social welfare and then ask whether informers in fact choose to provide this level of information. The results suggest that, although partial information is not always better (and in some cases is strictly worse), there are clear circumstances where it should be preferred. These circumstances are discussed in detail below. Importantly, with regards to the second question, there is very clear evidence that informers do in fact learn to provide the level of information that is welfare maximizing.

The remainder of the paper proceeds as follows: In the next section, I review some relevant literature. Afterwards, I describe the game that the experiment is based upon and derive theoretical predictions. Next, I describe the experimental design and procedure, and then present the results. Finally, I conclude with a discussion of the results.

2 Literature review

At the broadest level, this paper contributes to a literature which addresses the effects of uncertainty in public goods games. There is much less on this topic than one might expect given the prevalence of uncertainty in the real world provision of public goods. Many theoretical models and many experiments begin with the assumption that the involved parties have complete information about the quality of the public good (see Ledyard (1995) for a thorough review of experimental studies of public goods.) A recent exception is the work of Levati, Morone, and Fiore (2009), who address this issue directly. The authors find that uncertainty about the quality of a public good does indeed decrease contributions, with the important qualification that in their experiment a "rational agent" would not contribute in either their complete or incomplete information treatments. That is to say, they find that uncertainty serves to dampen the unexpected generosity of participants in a situation in which they might be expected to free ride. However, the present paper, and several previous papers, argues that uncertainty under different circumstances can *increase* contributions and social welfare.

Several authors have provided theoretical justification for the potential benefits of uncertainty in public goods games. For instance, Gradstein, Nitzan, and Slutsky (1994) analyze public goods games under different types of uncertainty and assumptions about preferences, attempting to determine the circumstances under which increasing the level of information available to contributors is welfare-enhancing. In an example which is particularly relevant to the present paper, they suppose that government has access to information about the magnitude of spillover that can be expected from a privately provided public good. They demonstrate that it may be preferable for government to withhold this information from the public for reasons very similar to the intuition described above – too much information can make clear to potential contributors that free-riding is the optimal strategy. In a very recent paper, Boucher and Bramoull (2010) demonstrate the potential benefit of uncertainty about public goods in the context of global public goods. Each of these papers demonstrate that, although the present paper employs a very simple example of a public goods game wherein limiting information can be beneficial, this result can hold in much richer and more realistic environments.

Of course, the appropriate amount of uncertainty must be achieved in order for it to be beneficial. Too much uncertainty, or uncertainty under which a rational agent would still be expected to free ride, can be harmful, as demonstrated by Levati et al. Thus, just as important as the question of *what the best level of information is* is the question of *whether society can achieve that level of information*. This question has been considered in a different context in the literature on "leading by example" in public goods games (and similar situations). Building on theoretical research (Vesterlund, 2003), both Potters, Sefton, and Vesterlund (2005, 2007) (PSV) and Komai, Grossman, and Deters (2007) (KGD) conduct experiments in which some participants are better informed of the quality of a group project which requires cooperation to achieve the socially best outcome[1]. They show that, when the uninformed participants are allowed to observe the contribution decision of the informed participants before making their own decisions, free-riding is reduced and social welfare is increased. In both cases, this works because the informed participant realizes that his or her contribution decision becomes a signal of the quality of the group project, thereby

[1]The differences between these papers are mostly in motivation. Both sets of authors come to roughly the same conclusions, but KGD is slightly more relevant to the present research as they set out to demonstrate the welfare gains result from incomplete information. PSV, on the other hand, are primarily interested in explaining why fundraisers would announce contributions and therefore focus on a comparison of sequential vs. simultaneous decision making in public goods games. It should also be noted that there are serveral theoretical papers by Komai and a variety of coauthors which explore the same ideas as KGD, but in much richer environments (Komai, Stegeman, and Hermalin, 2007; Komai and Stegeman, 2010).

communicating enough information to followers to incentivize contributing but not enough information to allow followers to determine whether they would rather free-ride. Each of these papers demonstrates that restricting information about the quality of a public good can lead to socially preferable outcomes, yet they rely on indirect means of communication – signaling quality through action – to achieve the desired level of information. Information about the quality of a public good is frequently communicated much more directly and often by an informed third-party agent, as is the case, for instance, with charity watchdog groups. Thus, the contribution of the present paper is a consideration of whether the benefit of partial information can be realized when the information is communicated directly by a third-party.

3 The game

The game played by participants is an extension of a simple linear public goods game. The game involves three players of two types – one "informer" and two "donors." The informer moves first, making a decision about the amount of information donors will receive. The donors then play the public goods game.

The public goods game played by donors builds upon a simple voluntary contributions mechanism framework in which the two players simultaneously decide whether to contribute an endowment w to a public good. Donors' payoffs are given by:

$$\pi_i^D = w - x_i + mX$$

where $x_i \in \{0, w\}$ is player i's contribution, $X = \sum_i x_i$ is the total supply of the public good (which, in this simple game, is equivalent to the sum of contributions), and m is the marginal per capita return to contributing to the public good (which I refer to interchangeably as the "quality" of the public good.) Here, m takes one of three equally likely values: 0.1, 0.7, or 1.5.

Donors are initially uncertain of the realized value of m. Thus, in the first stage of the game (that is, before the public goods (sub)game is played) the informer chooses whether to reveal "full" or "partial" information about m to the donors. If full information is chosen, donors know the value of m when making their contribution decisions. If partial information is chosen, donors instead observe an indication of whether m is "High" or "Low." Specifically, if $m \in \{0.7, 1.5\}$, donors observe that the m is "High;" if the true value of m is 0.1, then donors observe that the m is "Low." Upon observing the signal, donors make their contribution decisions. Thus, providing "partial" information is akin to providing donors a simple recommendation rather than fully revealing the quality of the public good. It is important to note that the informer does not know the value of m before making his or her decision.[2] I assume that informers are motivated to maximize social

[2]This is partially for experimental design reasons. If informers knew the value of m some might be tempted to only provide partial information when $m = 0.7$. This would then affect the donors' beliefs about what can be inferred from partial information when paired with different informers in future rounds,

welfare and therefore receive a payoff which is equal to the average of donors' payoffs:

$$\begin{aligned}\pi^I &= 0.5 \sum_i \pi_i^D \\ &= w - 0.5X + mX\end{aligned}$$

The game has a unique equilibrium: the informer chooses partial information; donors contribute when they observe that m is 1.5 or "High" and do not contribute otherwise. I elaborate below.

First consider donors' strategies. Note that contributing is socially optimal whenever $m \geq 0.5$. Also note that (risk neutral) donors will only contribute when $E(m) > 1$. Thus, under full information (where $E(m) = m$) donors will contribute if $m = 1.5$ and will not contribute otherwise. Under partial information, $E(m|low) = 0.1 < 1$ while $E(m|high) = (0.5 \times 0.7) + (0.5 \times 1.5) > 1$, so donors will contribute upon observing "High" and will not contribute upon observing "Low."

This makes clear the social benefits of partial information in this game – donors contribute when they would otherwise free-ride ($m = 0.7$) and as a result the socially best outcome is achieved. This leads directly to the informer's equilibrium strategy – because the informer's payoff is an average of the donors' payoffs, he or she will play the strategy which maximizes *their* payoffs, which is partial information. This is made clear by the table below, which compares each players' payoff under full and partial information for each possible value of m, assuming that donors behave according to the equilibrium strategy outlined above. (The table lists only one payoff per cell because, with donors behaving identically, the average of their payoffs is equal to each of their payoffs and so all players receive the same payoff.)

Table 1: Payoffs for all players given optimal donor decisions

$m =$	0.1	0.7	1.5
Full	w	w	3w
Partial	w	1.4w	3w

To both summarize the above and lay the groundwork for what I take as my experimental hypotheses, below are the resulting theoretical predictions:

Theoretical prediction 1. Partial information eliminates free-riding when $m = 0.7$ and therefore leads to the socially best outcome.

Theoretical prediction 2. As a result, informers always play partial information.

which introduces an undesirable dependance across rounds. Furthermore, this design feature is defensible on practical grounds – charity ratings organizations, for instance, first decide the format of the information they will provide (star ratings, pass/fail, etc.), and then "fill in" the content in their chosen format. Finally, partial information is a weakly dominant strategy whether the informer is aware of m or not, so this does not fundamentally alter the strategic considerations of the game.

4 Experimental design

The experimental design consists of three phases, the second of which is of primary interest. It is in the second phase that participants play the game that is described in the previous section. In order to obtain data regarding donor behavior under both informational environments (regardless of the choice of the first-mover) the strategy method is employed. Thus, informers and donors make their decisions simultaneously. Each round, informers choose full or partial information, while donors indicate whether they would or would not contribute their endowment to the public good conditional on the five possible messages that they might receive about its quality – namely, that m is 0.1, 0.7, 1.5, "High", or "Low"[3]. Thus, payoffs for each round depend on donors choices in the situation that is actually realized, conditional on (a) the true value of m and (b) the informers choice of full or partial information. After these decisions have been made, participants learn their payoff for the round, the true value of m, and the choices of the donors in their group for the realized situation.

Additionally, beliefs are elicited at the end of the second phase. After all participants indicate their choices in the final round but before the outcome is revealed, participants are asked to guess how many donors (or, in the case of donors, how many *other* donors) chose to contribute when $m = 0.7$ and how many chose to contribute when m is "High" in the final round. It is beliefs about these two situations that are most critical to the success or failure of partial information. Participants receive \$0.50 for each correct guess.

The first phase of the experiment is similar to the second phase, except that all participants are donors and, in the absence of designated informers, the level of information is randomly determined each round. This phase is intended to provide all participants an opportunity to learn the structure of the game and gain exposure to the incentives that donors face. Unlike the second phase, the strategy method is *not* employed. Thus, each round, given full or partial information about m, participants decide whether or not to contribute their endowment to the public good. As in the second phase, the outcome is revealed after each round.

Finally, the third phase of the experiment replicates the low-payoff version of the Holt-Laury risk elicitation procedure (Holt and Laury, 2002) and is included simply to obtain some indication of participants' risk preferences.

This design is intended to test two main hypotheses based on the theoretical predictions from above: (1) partial information is welfare maximizing (as measured by a comparison of full and partial information contribution decisions during the second phase) and (2) informers choose to limit the information they provide as a result. Risk preference and belief elicitations are included to rule out alternative explanations.

[3]In the experimental sessions, "informers" and "donors" are referred to as "first-movers" and "second-movers" respectively. The contribution decision is described as a choice between "A" and "B."

4.1 Procedures

Participants were recruited through the Pittsburgh Experimental Economics Laboratory (PEEL) recruiting database. The experiment was conducted in four sessions, each of which followed the design described above. There were 18 participants per session, with each session lasting roughly an hour. All participants received the same instructions. These instructions were read aloud to ensure that all details of the experiment were common knowledge. Participants were aware that there would be three phases but only received the instructions for each of the three phases as they occurred[4]. Participants were not aware that there would be belief elicitation questions until they occurred. Participants interacted and made all of their decisions anonymously through a computer terminal.

There were 12 rounds in the first phase and 20 rounds in the second phase, with endowments of \$4 each round. Participants were randomly paired in each round in the first phase. At the beginning of the second phase, 6 participants were randomly assigned the role of "informer" with the remaining 12 being assigned the role of "donor." These roles were maintained throughout the 20 rounds. In each round of the second phase, participants were randomly assigned to groups with one informer and two donors per group. In both the first and second phases, random assignment to groups was done in a way such that no two participants were ever in a group together in two consecutive rounds. The true value of m for each round of the first and second phases was randomly determined prior to the experiment.

Participants received \$5 for showing up in addition to their earnings during the experiment; participants were paid for one randomly selected round from the first phase, one randomly selected round from the second phase (plus any earnings from the belief elicitation), and the outcome of the lottery in the third phase. Participants were not aware of which rounds from the first and second phases had been selected for payment until the end of the experiment. Average earnings were \$15.95.

5 Results

I ask two main questions in this paper: What level of information about the quality of public good maximizes social welfare? When there exists a (socially interested) third-party with access to this information, do they provide the level of information that is welfare maximizing? The second question is of primary interest, but given the dependence of the second on the first, I address each of these in turn.

[4]Instructions from all three phases are included as an appendix.

5.1 Which level of information maximizes social welfare?

A natural initial response to this question is a simple comparison of average payoffs when informers choose full information to payoffs when informers choose partial information. This comparison is provided in Table 2. The table demonstrates that, on average, the choice of partial information leads to higher payoffs then full information.

Table 2: Mean payoffs under full and partial information

Information	**Mean payoff**	**(Std. Err.)**
Partial	6.555	(0.118)
Full	6.265	(0.133)

However, the situation is much more complicated than it would appear based on the above comparison. If we examine the relationship presented in Table 2 more carefully, then the apparent benefits of partial information disappear. Table 3 presents the results of a random-effects regression of payoffs on information (an indicator set to 1 if partial information is chosen) and a set of round fixed effects.

Table 3: RE regression - dep. var.: payoff

Variable	**Coefficient**	**(Std. Err.)**
Partial information	-0.048	(0.092)
Intercept	3.814**	(0.096)
(includes round FE's, not displayed)		
N	1440	
$\chi^2_{(20)}$	6847.044	
Significance levels : † : 10% ∗ : 5% ∗∗ : 1%		

The results of this estimation suggest that the level of information has very little impact on welfare. This too is misleading. It is instead true that partial information has very strong impacts in each of the four sessions, but that its effect is negative in two sessions and positive in two sessions. Thus, a more appropriate interpretation is that whether full or partial information is preferable depends largely on the nature of the population of donors receiving the information. To understand why this is, consider the following – there are two conditions that must be met in order for the benefits of partial information to be realized: (1) as discussed previously, it must be the case that *some* donors would free-ride under full information and (2) it must be that the fraction of donors who are willing to contribtute upon observing 1.5 but *not* "High" is small.

Condition One is entirely plausible (as has been demonstrated in countless previous public goods experiments), but it is also not damning if it is not satisfied. In particular, if every donor *were* to contribute upon observing 0.7, then partial information is *just as good* as full information. Condition Two is of greater concern. A failure of Condition Two – meaning that donors are only willing to contribute when they are certain that the public good is of the highest quality – would imply that partial information in fact reduces social welfare.

More precisely, let $E[X|s]$ be the expected sum of contributions to the public good when donors observe the message s. If we make the simplifying assumptions that (1) nobody gives upon observing that the return is 0.1 or "low" and (2) everybody gives upon observing $m = 1.5$,[5] then it is easily shown that partial information maximizes social welfare under the following condition:

$$0.2[E[X|high] - E[X|0.7]] > E[X|1.5] - E[X|high]$$

The difference on the left hand side represents the potential benefits of partial information – increased giving in the $m = 0.7$ case. The right hand side demonstrates the potentially harmful impact of partial information – namely, lost contributions when $m = 1.5$. Thus, the expected gain in total contributions when $m = 0.7$ must be five times larger than the losses when $m = 1.5$ in order for partial information to be beneficial – this simply emphasizes the fact that a lost contribution to the highest quality public good is significantly more costly to society than a lost contribution to the middle quality good.

Recall that $X = x_1 + x_2$ where $x_i \in \{0, 4\}$ is the contribution of group member i. Let p_s represent the probability that any donor i chooses $x_i = 4$ upon observing signal s. Thus, assuming independence across donors, $E[X|s] = 8p_s$. With this, it can be shown that the above inequality simplifies to the following condition: partial information maximizes social welfare if

$$6p_{high} > 5 + p_{0.7}$$

Note that even in a population of donors who *always* free-ride when $m = 0.7$ (i.e., $p_{0.7} = 0$), in order for partial information to be preferable, it must be that $p_{high} > \frac{5}{6}$. This establishes a lower-bound for the acceptable level of non-giving in the "high" case.

Returning to the data, in two of the experimental sessions there is some fraction of participants who fail to contribute upon observing "high" in every round – there are three of these individuals in one session and two in another. Thus, in any given round of each of these sessions, the proportion of donors contributing upon observing "high" can be no higher than $\frac{3}{4}$ in the session with three never-givers and $\frac{5}{6}$ in the session with two. These sessions clearly fail to meet the criteria above. Thus, even if every other donor follows the payoff-maximizing equilibrium strategy (which of course is not true in any of the sessions), in aggregate, partial information would appear to fail to maximize social welfare.

With this in mind, it is reasonable to separately explore the impact of partial information in the two sessions wherein this threshold of "high"-non-givers is violated and in the two sessions where

[5]These assumptions are approximations of what is observed in the data – both of these assumptions are true in roughly 95% of observations. They are made simply for the sake of expositional clarity. The argument presented in this section would hold in the absence of these assumptions.

it is not.[6] As an intitial demonstration of the sharp contrast across these sessions, Table 4 repeats the analysis presented in Table 3, but conducted separately for the sessions that have not crossed the threshold of "high"-non-givers and again for those that have (for the sake of brevity, these are referred to from here on as "below-threshold" and "above-threshold" sessions respectively). The top panel demonstrates a clear positive and significant impact of partial information in below-threshold sessions, whereas the bottom panel shows an equally clear and significant negative impact of partial information. Notably, as one would expect from the discussion above, the magnitude of the negative impact is larger than that of the positive impact.

Table 4: RE regression - dep. var.: payoff

Variable	Coefficient	(Std. Err.)
Below-threshold sessions	N=480	
Partial information	0.322**	(0.107)
(includes round FE's, not displayed)		
Intercept	3.786**	(0.170)
Above-threshold sessions	N=480	
Partial information	-0.564**	(0.164)
(includes round FE's, not displayed)		
Intercept	3.921**	(0.175)

Significance levels : † : 10% ∗ : 5% ∗∗ : 1%

Of course, the level of information is endogenous so comparing outcomes of full and partial information based on the outcomes that are actually realized is not ideal. Table 4 has simply been presented to demonstrate the stark contrast across sessions above and below the threshold of "high"-non-givers. From here on, I take advantage of the fact that donors' full strategy is elicited each round. Thus, I can compute expected outcomes under full and partial information independent of the level of information that is chosen.

I use two primary sets of measures. The first calculates the payoffs that an individual would have received under both full and partial information. Of course, depending on the choice of the informer one of these is the actual payoff, while the other is a counterfactual. Thus, the appealing aspect of this set of measures is that they are closely tied to the actual outcomes and therefore provide a straightforward interpretation – namely, an indication of what would have happened had the informer made a different information choice. The less appealing aspect of this set of measures is their dependence on the realized sequence of randomly determined m's. Thus, a wealth of data is ignored under this approach.

The second set of measures then captures a similar idea, but takes fuller advantage the dataset. Unlike the first set, this set of measures is *not* conditioned on m and instead provides a calculation

[6] To be clear – even in the two sessions where this threshold is not crossed, it is certainly still the case that donors *sometimes* fail to give upon observing "high." The important point is that there is no fixed fraction of donors for whom this is always true, as there is in the population of donors in the other two treatments, and therefore it is still mathematically possible that partial information, on average, would be preferable.

of the ex ante expected payoffs under full and partial information. Table 5 summarizes these measures, split by above- and below-threshold sessions. To test the significance of the difference of each full and partial measure, I conduct a Wilcoxon sign-rank test, taking the mean of each measure from each round for a particular session as the unit of observation.[7] Table 6 summarizes the same measures for the full dataset.

Table 5: Comparison of full and partial information - split by session category

	Mean	**(Std. Err.)**
Below-threshold sessions		
actual payoff if partial	6.877	(0.150)
actual payoff if full	6.530	(0.168)
difference	0.347*	(0.063)
expected payoff if partial	7.040	(0.028)
expected payoff if full	6.747	(0.022)
difference	0.293**	(0.033)
Above-threshold sessions		
actual payoff if partial	5.823	(0.131)
actual payoff if full	6.320	(0.165)
difference	-0.497	(0.112)
expected payoff if partial	5.991	(0.067)
expected payoff if full	6.520	(0.033)
difference	-0.529**	(0.061)

Significance levels : † : 10% $*$: 5% $**$: 1%

Table 6: Comparison of full and partial information - all sessions

	Mean	**(Std. Err.)**
All sessions		
actual payoff if partial	6.350	(0.101)
actual payoff if full	6.425	(0.118)
difference	-0.075	(0.065)
expected payoff if partial	6.516	(0.040)
expected payoff if full	6.633	(0.020)
difference	-0.118	(0.037)

Thus, although Table 6 once again demonstrates that, as a whole, the effect of partial information is somewhat ambiguous[8], when we look within sessions – as in Table 5 – there is a very clear

[7]The round-means are paired with the corresponding round-means for the opposite measure from the other session within the same threshold category. For example, sessions 1 and 3 are "below-threshold sessions." So to determine the significance of the the first difference in Table 5, for each round I pair the session 1 "actual payoff if partial" measure with the session 3 "actual payoff if full" measure. Also note that, though the means are the same, the standard errors in Table 5 are from the full dataset, rather than the round-mean collapsed dataset.

[8]Although, if there were any effect in the full sample, it would appear to be negative.

impact of partial information. In particular, as suggested by the discussion above, full information is unambiguously preferable in sessions where a minimum of $\frac{1}{6}$ of the donor population is guaranteed to not contribute upon observing "high" every round. However, when this is *not* the case, as in the below-threshold sessions, then the theoretical predictions regarding partial information are clearly fulfilled.

5.1.1 Why do some donors never contribute upon observing "high"?

In the above analysis, I have argued that the critical determinant in the impact of partial information is the variation in the proportion of the donor population that never gives upon observing "high" across sessions. An important question is why this variation exists – that is, why would a donor *not* give upon observing "high"? I briefly discuss several possible explanations before moving on.

One possibility (particularly amongst donors that consistently fail to give in the "high" case) is that donors are attempting to affect informers' beliefs about the way that informers will react to partial information. After all, though partial information may be socially preferable, an individual donor is always better off with complete information in that she can more accurately decide when to free-ride. Thus, sacrificing current gains to modify informers' behavior may be profitable in the long-run. The experiment is designed to minimize such behavior by randomly rematching participants each round, but with only 6 informers and 20 rounds, donors are almost sure to meet each informer multiple times. However, the donors who fail to contribute in the "high" case in the second phase of the experiment *also* tend to fail to do so in the first phase. Recall that all participants were completely unaware of the concept of "informers" during the first phase. Thus, it is very unlikely that donors are motivated to impact informers' beliefs.

Another possible explanation is that these subjects are very risk averse. To address this, risk preferences were elicited using the low-payoff version of the Holt and Laury (2002) procedure. However, these measures fail to significantly predict the behavior of "high" non-givers.[9] Thus, either the Holt-Laury risk measures fail to generalize to the game that this experiment is based upon or some donors simply have not realized that, in expectation, contributing is the best response to the "high" signal.

[9]This is based on several possible approaches – from simply incorporating the risk measures into probit estimations to estimating the CRRA utility of contributing and not contributing using beliefs about other-informer behavior constructed from previous observations. The results of these approaches are roughly consistent with the possibility that risk aversion is driving this behavior, but none of these results are significant.

5.2 Do informers choose the level of information that maximizes social welfare?

Overall, 480 information decisions are made (6 informers across 20 rounds in 4 sessions). Mirroring the ambiguity of the initial result from the previous section, informers choose full information in 49.4% of these decisions. However, given the stark contrast across sessions in society's preference for full or partial information, if informers are indeed choosing the level of information the maximizes social welfare then we should see a similar divergence across sessions in information choice. Table 7, which lists the frequencies of full information across above- and below-threshold sessions, demonstrates that this is the case.

Table 7: Frequency of full information by session category

	Freq. of full info.	**(Std. Err.)**
Above-threshold sessions	0.613	(0.018)
Below-threshold sessions	0.375	(0.018)

Notably, we see that in above-threshold sessions, where full information is preferable, informers do in fact choose full information which much greater frequency than in sessions where the possibility for partial information to be beneficial still exists. Figure 1 demonstrates the evolution of information choices across these session categories. In below-threshold sessions, we observe a very clear convergence to partial information. Interestingly, above-threshold informers start at roughly the same point as below-threshold informers, but in the absence of evidence that partial information is beneficial, they remain at a relatively constant frequency of full information.

Next I more directly address the question of whether informers are in fact choosing the level of information that maximizes social welfare. That is, I have shown thus far in this section that informers choose partial relatively often and that the frequency thereof roughly coincides with the sessions that we would expect it to. Here, I ask whether, on a period-by-period basis across all sessions, informers are choosing the level of information that is in fact expected to maximize social welfare in that particular period. To do so, using a random-effects probit model, I estimate the probability that an informer chooses partial information conditional on the *within-session* and *within-round* average of the expected gains from choosing partial information across all informers (in the same session and round). Specifically, recall the "expected payoff from partial" and "expected payoff from full" measures from Table 5 in the previous section. I average the difference between these two measures for all informers within a session for each round and take that as the independent variable of primary interest.

The results are presented in Table 8. There is a very clear positive relationship between the average expected welfare gains (or losses) that may result from choosing partial and the likelihood of doing so. This result increases our confidence that, not only does choosing partial information coincide with potential gains from doing so, but that informers are choosing partial *because* there

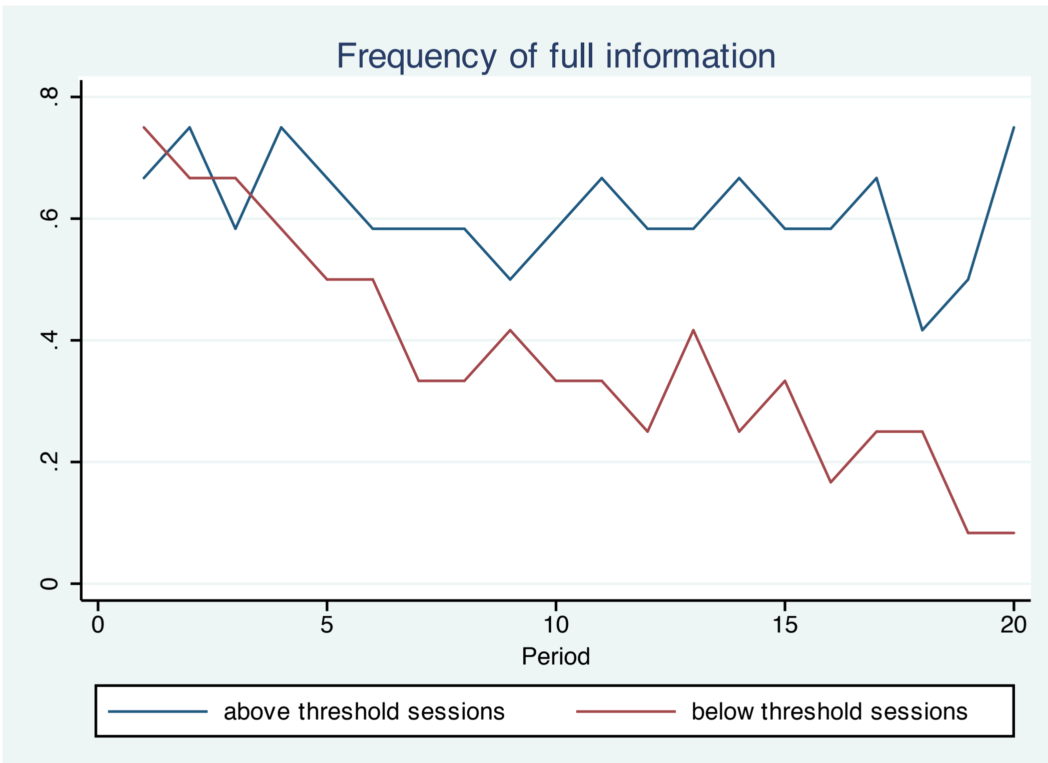

Figure 1: Frequency of full information for above- and below-threshold sessions

are gains.

Table 8: RE Probit – Do informers maximize social welfare?

Variable	**Coefficient**	**(Std. Err.)**
Dep. var.=1 if partial info. is chosen		
Average expected partial info gain	0.891*	(0.391)
(includes round FE's, not displayed)		
Intercept	-0.976*	(0.459)
N	480	
Log-likelihood	-210.256	
$\chi^2_{(20)}$	43.496	

Significance levels : † : 10% $*$: 5% $**$: 1%

Finally, the previous section emphasized the importance of the probability of donors contributing in the "high" case in determining whether full or partial information is preferable. A simple model of informer beliefs provides interesting insights into the information choice and confirms the importance of this particular probability. Specifically, for every informer in every round, I construct beliefs about the probability that donors will contribute in each of the five cases (0.1, 0.7, 1.5, low, and high). These beliefs are constructed as the simple empirical frequency with which they have observed donors contribute in each of these five situations for all rounds up to the present, including observations from the first phase of the experiment. Thus all informers have observed donor behavior in – and therefore have beliefs about – each of these situations before the second phase begins. These beliefs are continuously updated with new observations throughout the second phase.

In Table 9, I present a random-effects probit estimation of the likelihood of choosing partial information conditional on these beliefs (and a set of round fixed effects). In the table, these beliefs follow the notation of the corresponding probabilities from the previous section. (That is, p_s is informer i's belief in a given round that a donor will contribute upon observing signal s.) The previous section argued that p_{high} is critical to determining whether full or partial information is preferable. Table 9 presents evidence that this probability is also *the* critical factor when an informer decides whether to choose full or partial information, providing further evidence that informers choose the level of information to maximize social welfare – even if this means limiting the amount of information that is available.

6 Discussion & conclusion

In this paper, I study the effects of a third-party information provider in a public goods game wherein potential contributors are uncertain of the quality of the public good. This environment is not of purely theoretical interest; with the emerging influence of charity ratings organizations, this is an increasingly accurate description of the situation faced by those considering making charitable

Table 9: RE Probit – Impact of beliefs on info. choice

Variable	Coefficient	(Std. Err.)
Dep. var.=1 if partial info. is chosen		
$p_{0.1}$	-3.656	(2.343)
$p_{0.7}$	0.099	(0.737)
$p_{1.5}$	-0.280	(1.452)
p_{low}	-0.358	(1.343)
p_{high}	3.723^{**}	(0.952)
(includes round FE's, not displayed)		
N	480	
Log-likelihood	-202.531	
$\chi^2_{(24)}$	52.803	

Significance levels : † : 10% ∗ : 5% ∗∗ : 1%

contributions. Thus, the question that arises is: how should such an organization communicate the information they hold to donors, assuming their goal is to achieve the best social outcome? In this paper, I suggest that the way these organizations benefit society is not just in their ability to provide information, but also in their ability to *limit* the amount of information they provide (or, more precisely, to pool their signals.)

To explore these ideas, I conduct a laboratory experiment that extends a simple public goods game often found in the experimental literature by incorporating a third-party information provider. Within the context of the experiment, I ask (1) what the welfare maximizing level of information is and (2) whether informers choose to provide this level of information. The simple theoretical prediction is that limiting information can overcome the classic free-rider problem and therefore informers should always choose to provide partial information. The results are not as clear cut as the theoretical prediction. Namely, with regards to the first question, it is not the case that partial information is always preferable. However, the results *do* demonstrate that there are circumstances under which partial information *is* unambiguously preferable and that these circumstances are clearly defined (namely, by the threshold proportion of donors that fail to give upon observing "high"). Given this result and turning to the second question, there *is* fairly clear evidence that informers choose the level of information that maximizes social welfare.

Thus, with regards to the motivating example of charity ratings groups, I have argued that, in theory, they may be justified in limiting the information they provide to simple star or pass/fail ratings. While the results of the experiment are not conclusive evidence of this suggestion, the important point is that partial information *can* be preferable contingent on the characteristics of the donors to whom the information is directed. Namely, the discussion above suggests that partial information only fails if a large enough fraction of the recipients of the information are wary of giving unless quality is made very clear. This may not be a problem at all in the context of charity ratings organizations. Namely, a critical difference between the game played here and the decision that real donors face is that real donors *seek out* the information provided by charity ratings organizations, presumably with prior knowledge of the nature of the information they are seeking. Thus, this important determinant of the impact of partial information may be easily satisfied in reality as a result of donors “selecting into a partial information treatment.”

The simple game that participants played in the laboratory is designed to capture the essential features of the interaction between real world information providers, such as charity ratings organizations, and potential donors. Several simplifying assumptions are made in doing so. For instance, it may not be reasonable to model information providers as purely altruistic. Charity watchdog groups may be partially motivated to attract more visitors to their websites than their competitors, which could affect the nature of the information they provide. Additionally, an important question that is ignored by this design is the appropriate behavior of informers when they must communicate the quality of multiple public goods. Suppose for instance that there are two public goods that are both socially efficient, but one is of superior quality. Under the current design, if partial information is chosen these goods would both be classified as “High” quality, when it would clearly be preferable to direct all funds to the superior public good.[10] These are valid concerns and would provide interesting extensions to the present research. However, I would argue that this experiment provides an important starting point – it is only once we verify that the benefits of an information provider are realized in this simple environment that it makes sense to test how robust this result is to increasingly realistic environments.

[10] Of course, partial information still represents an improvement over full information in this case if both of the public goods have marginal per capita returns of less than one.

References

Boucher, Vincent and Yann Bramoull. 2010. "Providing global public goods under uncertainty." *Journal of Public Economics* 94 (9-10):591 – 603.

Gordon, T.P., C.L. Knock, and D.G. Neely. 2009. "The role of rating agencies in the market for charitable contributions: An empirical test." *Journal of Accounting and Public Policy* 28 (6):469–484.

Gradstein, M., S. Nitzan, and S. Slutsky. 1994. "Uncertainty, information, and the private provision of public goods." *European Journal of Political Economy* 10 (3):449–464.

Grant, L. 2010. "The Response to Third-Party Ratings: Evidence of the Effects on Charitable Contributions." *Working paper* .

Holt, C.A. and S.K. Laury. 2002. "Risk aversion and incentive effects." *American Economic Review* 92 (5):1644–1655.

Komai, M., P.J. Grossman, and T. Deters. 2007. "Leadership and Information: An Experimental Study." *St. Cloud State University Working Paper* .

Komai, M. and M. Stegeman. 2010. "Leadership based on asymmetric information." *The RAND Journal of Economics* 41 (1):35–63.

Komai, M., M. Stegeman, and B.E. Hermalin. 2007. "Leadership and information." *American Economic Review* 97 (3):944–947.

Ledyard, J.O. 1995. "Public goods." *Handbook of experimental economics* :111–94.

Levati, M.V., A. Morone, and A. Fiore. 2009. "Voluntary contributions with imperfect information: An experimental study." *Public Choice* 138 (1):199–216.

Lowell, S., B. Trelstad, and B. Meehan. 2005. "The ratings game." *Stanford Social Innovation Review* 3:38–45.

Potters, J., M. Sefton, and L. Vesterlund. 2005. "After you - endogenous sequencing in voluntary contribution games." *Journal of Public Economics* 89 (8):1399–1419.

———. 2007. "Leading-by-example and signaling in voluntary contribution games: an experimental study." *Economic Theory* 33 (1):169–182.

Sloan, M.F. 2009. "The Effects of Nonprofit Accountability Ratings on Donor Behavior." *Nonprofit and Voluntary Sector Quarterly* 38 (2):220.

Vesterlund, L. 2003. "The informational value of sequential fundraising." *Journal of Public Economics* 87 (3-4):627–657.