

Detecting other-regarding behavior with virtual players

Paul J. Ferraro^{a,*}, Daniel Rondeau^b, Gregory L. Poe^{c,1}

^a *Department of Economics, Andrew Young School of Policy Studies, Georgia State University,
University Plaza, Atlanta, GA 30303-3083, USA*

^b *Department of Economics, University of Victoria, Victoria, BC, Canada V8W 2Y2*

^c *Department of Applied Economics & Management, Ithaca, NY 14853-7801, USA*

Received 1 August 2000; received in revised form 5 February 2001; accepted 12 February 2001

Abstract

Humans display levels of cooperative behavior that contradict the predictions of theoretical models of rational self-interested individuals. We propose a novel technique to discriminate among other-regarding behavior, self-interested strategic play, and decision errors in laboratory experiments. We introduce “virtual players” in two public goods experiments to remove the concerns of human subjects for other players. Comparing contributions across treatments, we find that other-regarding behavior elevates contributions. The results also suggest that subjects are motivated by fairness considerations. We discuss ways in which the virtual-player technique can discriminate among competing explanations of behavior observed in other experiments.

© 2002 Elsevier Science B.V. All rights reserved.

JEL classification: C91; C92; D63; D64; H41

Keywords: Altruism; Fairness; Laboratory experiments; Public goods; Group behavior

1. Introduction

Game theoretic models of bargaining and voluntary contributions to public goods make strong predictions that rely on the hyper-rationality and self-interest of players. In a dictator game, standard theory predicts that a proposer who is asked to unilaterally divide a sum of

* Corresponding author.

E-mail address: pferraro@gsu.edu (P.J. Ferraro).

¹ Visiting Fellow of Jackson Environmental Institute (JEI), Centre for Social and Economic Research for the Global Environment (CSERGE), School of Environmental Sciences, University of East Anglia, Norwich, Norfolk NR4 7TJ, UK.

money between herself and another player will offer nothing to the other player. Similarly, players should not cooperate in repeated prisoner dilemma games, nor contribute to the funding of public goods in typical voluntary contribution games.²

These predictions, however, are consistently violated in laboratory experiments. Subjects offer significant amounts of money to other players in dictator games, achieve non-Nash levels of cooperation in repeated dilemma games and voluntarily contribute to the funding of public goods when these actions are not in their immediate self-interest (Davis and Holt, 1993; Kagel and Roth, 1995). Explaining these robust empirical results is challenging. Leading suggestions can be organized into three categories: (1) confusion and decision errors; (2) strategic considerations stemming from individual motivations that extend beyond the absolute magnitude of a player's own payoff (e.g. caring about relative income), or stemming from incomplete information and uncertainty about other players' motivations, payoffs and rationality; (3) "other-regarding" preferences (altruism, fairness, warm glow, reciprocity, etc.). Discriminating among these competing explanations is difficult because non-monetary considerations cannot be directly observed, nor easily manipulated and controlled for in the laboratory.

We propose a novel technique to discriminate among other-regarding behavior, self-interested strategic play, and decision errors in laboratory experiments. The approach relies on the introduction of non-human (virtual) players in games normally played by humans only. We posit that opposing a single human subject to virtual players that do not receive payoffs neutralizes the other-regarding components of the human subject's utility function. By appropriately defining the virtual players' strategy space and behavioral rules, and providing this information to the human subject, it is possible to remove either other-regarding behavior only, or both strategic considerations and other-regarding preferences simultaneously. For instance, if virtual players are programmed to behave exactly as if they were humans, and human players are aware of this, we expect humans to behave no less strategically than humans who play with humans. Thus, one can detect other-regarding behavior by comparing the behavior of humans playing with humans in the control group to the behavior of humans playing with virtual agents in the treatment group. Alternatively, if virtual players use pre-defined decision rules that are not associated with the actions of other players, the experimental environment encourages neither strategic nor other-regarding considerations.

We present the virtual-player method and the result of its application in the two mechanisms most widely used to raise funds for public goods: the voluntary contribution mechanism (VCM) and the provision point mechanism (PPM). Our results support the view that players in these games care about the welfare of other players. We are also able to infer that the other-regarding behavior is partially motivated by a desire for fairness in the allocation of costs for the public good.

In the next section, we briefly review previous approaches to detect and measure other-regarding behavior and separate it from decision error or strategic play. In Section 3, we describe the virtual-player approach, the public goods experiments, and our results. We conclude in Section 4 by discussing the wider applicability of the virtual-player technique.

² These games and their solutions are described in most introductory game theory textbooks.

2. Efforts to detect, separate and measure other-regarding behavior

Economists have used different approaches to separate errors, other-regarding behavior and strategic considerations in public goods and bargaining games. In public goods experiments, they have relied on clever manipulations of payoffs and econometric methods (e.g. Andreoni, 1995; Palfrey and Prisbrey, 1997; Goeree et al., 1999). The reliance on manipulating payoffs, however, has resulted in incomplete separation of motives and the need to apply econometric methods to discriminate among alternative hypotheses. It also requires experimentalists to specify the elements that enter a subject's utility function and to assume a functional form.

In bargaining experiments, experimentalists have focused more directly on controlling subject motivations by adjusting the relative power of individual agents. Harrison and McCabe (1996) introduced computer automata to an ultimatum game in order to control the stimuli applied to subjects and to observe learning. Forsythe et al. (1994) used powerless human agents to transform the ultimatum game into the dictator game. The presence of a powerless receiver in the dictator game allows one to make inferences about the motivations of subjects. Güth and van Damme (1998) refined this analysis by introducing a third (powerless) player in the ultimatum game. The third player helps separate strategic play from concerns for fairness toward other players.

Our introduction of virtual players into public goods games is a natural extension of previous approaches. The method confers the ability to directly remove other-regarding concerns in addition to, or instead of, strategic considerations and is a logical step toward a better understanding of behavior in public goods games.

3. Design and applications

We employ an experimental design that focuses on the identification and removal of all forms of other-regarding behavior in two public goods funding mechanisms: (1) a variant of the VCM; and (2) a variant of the PPM studied by Rondeau et al. (1999).

By removing all but one human from the game, we neutralize the predisposition of some subjects to care about the welfare of others. We hypothesize that humans playing a one-shot public good game with virtual players will contribute less than humans playing an identical game with human players. We attribute any changes in contribution levels to the presence of other-regarding preferences.

To cleanly isolate other-regarding behavior, we must ensure that the control experiment (humans playing with humans) and the treatment (humans playing with virtual players) produce identical strategic incentives. We maintain the strategy set across treatments by making the virtual players' actions reflect past human play, and by informing subjects of this fact. We first run control experiments with human subjects and gather the individual data on contributions. The data from the all-human sessions (N players per session) generate a databank of individual choices from which the actions of virtual players in subsequent automata treatments are drawn. In virtual-player treatments (Appendix A), a single human subject plays with $N - 1$ virtual players, which are explicitly characterized in the instructions as non-human agents. The human subject also receives information on the way in which the virtual players' actions are determined.

We conducted three sets of experiments. The “Fall 1998” (PPM control with humans) and “Spring 1999” (PPM virtual-player treatment) experiments were a pre-test of the method, but we find the data sufficiently interesting to report them below. In the fall of 1999, we applied the same method to both the VCM and PPM institutions, using students recruited from a single class.

3.1. Fall 1998/Spring 1999 PPM experiments

The Fall 1998 experiment drew its subject pool from an introductory engineering economics class (similar to a managerial economics class). Students were not taught about public goods. Thirty-nine students in the class participated in an all-human PPM experiment during a weekly section meeting. Beyond the usual instructions regarding anonymity, subjects were told that the experiment was completely unrelated to the course.

The Fall 1998 PPM institution can be summarized as follows.³ Each subject receives an initial endowment $E = \text{US\$ } 12$ and divides this amount between a private account and a public investment fund (a contribution C_i by player i). Any amount deposited in the subject’s private account becomes part of the subject’s payoff. If the sum of all subject contributions is below a threshold, or provision point (PP), all contributions to the group investment funds are reimbursed (a money-back guarantee). In this case, i ’s payoff is simply his initial endowment of US\$ 12. If total contributions to the investment fund exactly match the PP, all subjects in the group receive $V = \text{US\$ } 6$ for a net payoff of $E - C_i + V = 18 - C_i$. Finally, if the sum of contributions exceeds the PP, excess contributions return to subject i proportionally to the weight of i ’s contribution to the investment fund. In this final case, i ’s net payoff is:

$$E - C_i + V + \frac{C_i}{\sum_{j=1}^n C_j} \left(\sum_{j=1}^n C_j - \text{PP} \right) = 18 - \frac{C_i}{\sum_{j=1}^n C_j} \text{PP}$$

Prior to his decision, a subject knows the group size and that all players have the same endowment (E) and value (V) for the investment.

In our design, the public good has an uncertain cost, with the PP chosen randomly from a known uniform distribution. Subjects know that, after all decisions have been made, the PP is revealed by drawing from a bingo cage containing 25 balls numbered 0 to 24 and multiplying the number drawn by US\$ 10. We chose this design to match the design of other public goods experiments conducted at Cornell as part of a larger NSF-funded research program. In the PPM game with random cost chosen from the uniform distribution $[0, C_{\max}]$, there is no dominant strategy and, with N agents in a group, the symmetric Nash equilibrium contribution is⁴:

$$\left(\frac{2}{N + 1} \right) V$$

³ Instructions are available from <http://web.uvic.ca/~rondeau/fileshare.htm>

⁴ To derive this equilibrium, simply solve the following problem for the case in which all subjects contribute the same amount ($b_i = b_j$): $\max_{b_i} \int_0^{b_i + B_{-i}} (1/C_{\max})(V_i - (b_i/\sum_{j=1}^N b_j)C) dC$, where B_{-i} are the contributions of all agents except i .

In the 1998 PPM experiment, the mean contribution to the public investment fund was US\$ 6.09 (median = US\$ 6.00; S.D. = 3.48). The high level of contributions in this one-shot environment is consistent with the results of Rondeau et al. In 1999, we replicated the Fall 1998 experiment in the same engineering course and under the same conditions except for the introduction of virtual players. Each 1999 subject played in a group consisting of the subject and 38 virtual players. As explained earlier, the contributions of the 38 virtual students were drawn randomly from the set of contributions generated in the 1998 all-human experiment and 1999 subjects were informed of this process. With the exception of additional language concerning the virtual players, the instructions and experimental environment were identical across the 1998 and 1999 sessions.

The 1999 data exhibit the expected decrease in contributions. Subjects in the virtual-player environment contributed an average of US\$ 4.60 (median = US\$ 5.00; S.D. = 3.31). Statistical differences between the human and virtual-player treatments are significant. For instance, a *t*-test on means yields a *P*-value less than 0.03 and a Mann–Whitney test of equality of the medians yields $P < 0.06$. The percentage of US\$ 0 contributions increased from 3 percent in the human treatment to 21 percent in the 1999 virtual-player environment (a test of equal proportions yields $P < 0.01$).

While these results are suggestive of the role of other-regarding behavior in this PPM environment, we must question their robustness since the subject pool was not strictly identical between the two years. Our next study draws subjects from the same pool and applies the virtual-player technique to the VCM as well as the PPM environment.

3.2. Fall 1999 VCM/PPM experiments

In the Fall 1999 experiments, 169 subjects from an introductory undergraduate economics class participated. These subjects had already participated in four other experiments during the semester including a repeated-round duopoly experiment in which deviations from the payoff-maximizing collusive strategy were common and discussed in follow-up teaching activities. Furthermore, the names of top earners to date were posted in the experimental laboratory. The top moneymakers were to receive additional cash prizes at the end of the semester, although our experiment was not included in the class competition. Thus, our subject pool can be considered an “extreme” environment in which to search for altruistic preferences: subjects were “economists in training”, operating in an environment in which self-interest was being reinforced.

Experimental sessions took place over 2 days in the class’s weekly section meetings. Subject participation was quasi-voluntary—a tiny portion of the final grade was attributed to attendance at weekly sections. About half of the subjects participated in a VCM treatment, the other half in a PPM. Roughly half of the subjects made their decisions in all-human groups. The other half made their decisions in groups of “virtual students”. Thus, the experimental design had four cells: (1) PPM with all-human groups; (2) PPM with virtual-student groups; (3) VCM with all-human groups; (4) VCM with virtual-student groups. Each cell consisted of two sessions of about 20 students each.

The PPM experiment had a design similar to the 1998 pre-test with the exception of three minor changes. The group size was reduced to 20, the individual payoff was increased to US\$ 7 and the random provision point was drawn from the uniform distribution with

Table 1

Results of the comparison of contribution levels: all-human groups vs. virtual-human groups

	VCM		PPM	
	Human	Virtual	Human	Virtual
Mean contribution (US\$)	2.14	1.16	5.30	4.27
S.D.	2.06	1.89	3.89	3.38
Median (US\$)	2.38	0.00	5.00	4.00
Induced value ^a (US\$)	7	7	7	7
Endowment (US\$)	12	12	12	12
Sample size ^b	42	43	40	38
Mood–Westenberg test	$P = 0.033$		$P = 0.035$	
Fligner–Pollicello test	$P = 0.010$		$P = 0.098$	
Mann–Whitney test	$P = 0.011$		$P = 0.099$	
<i>t</i> -test	$P = 0.012$		$P = 0.110$	

^a In the VCM session, subjects received US\$ 0.07 per share purchased at US\$ 1 each by the group, up to a maximum of 100 shares purchased or US\$ 7 in individual payoff.

^b VCM human experiments were conducted in groups of 21 individuals. All 43 participants in the virtual-player treatment were informed that they were playing with 20 virtual agents. In the PPM, the human group size was 20 and all 38 players in the virtual treatment were told that they were playing with 19 virtual players.

support [0, 140]. In the VCM design, the group size was 21 and, as in the PPM design, subjects had common endowments, payoffs, and information (see Table 1).

It should be noted that our VCM design differs slightly from the VCM design most commonly found in the literature (Isaac et al., 1984). In our design, subjects receive an individual return of US\$ 0.07 per dollar contributed to the group fund, up to a maximum of US\$ 7. Thus, contrary to previous research on the VCM, the social optimum is for the group to contribute US\$ 100 rather than the entire collective endowment. Subjects still have a dominant strategy to contribute nothing. We chose this design to match the design of other public goods experiments conducted at Cornell as part of a larger research program. The design matches real world situations in which the maximum quantity of the public good that can be provided is restricted (e.g. roads removed from a national park, Champ et al., 1997).

As predicted, subjects in these experiments behaved differently depending on whether they played with humans or with virtual players (see Table 1). The differences were most striking in the VCM environment in which the dominant strategy was to contribute nothing to the public account. In this environment, subjects in the human treatment contributed an average of US\$ 2.14 (median of US\$ 2.38) while their counterparts in the virtual-player treatment contributed US\$ 1.16 (despite the virtual-player sessions containing the two highest VCM contributions). It is worth noting that the median contribution level in the virtual treatment is exactly zero, consistent with the Nash equilibrium prediction. The percentage of zero contributions increased from 40 percent in the all-human sessions to 60 percent in the virtual-agent sessions ($P = 0.03$).⁵ We attribute the positive contributions of the remaining subjects to decision errors. Our results, like those of Andreoni, suggest that, approximately 50 percent of total contributions are a result of other-regarding behavior and the rest a result of decision error.

⁵ The number of zero bids in the PPM environment also increased, but not significantly. Since a zero contribution is not a dominant strategy, this result is not surprising.

Table 1 lists four statistical tests of significance, ranging from non-parametric tests that make few assumptions about the underlying distribution of contributions to the parametric *t*-test that makes strong assumptions about the underlying distribution. We include the non-parametric tests (Mood, 1950; Westenberg, 1948; Fligner and Policello, 1981) because of the highly irregular, skewed sample distributions generated by the experiments tests lead to a rejection of the normality hypothesis. Given such poorly-behaved distributions, we believe that the Mood–Westenberg test, a non-parametric test with few assumptions, is the most appropriate test.⁶ For the VCM data, the tests yield *P*-values ranging from $P = 0.01$ to $P = 0.03$.

In the PPM experiments, subjects playing with humans contributed an average of US\$ 5.30 compared to the US\$ 4.27 contributed by subjects playing with virtual agents. In the same order, medians are US\$ 5.00 and US\$ 4.00. The battery of statistical tests reported in Table 1 yields *P*-values ranging from 0.035 to 0.11. The weaker statistical results for the PPM can be explained by the greater variance in contribution levels observed in this mechanism compared to the VCM. The greater variance leads to a greater probability that the results are consistent with the null hypothesis and is likely attributable to the absence of a dominant strategy in the PPM. Thus, we observe a flatter distribution and a greater variance. It is notable, however, that in both the PPM and VCM experiments, the difference in mean contributions attributed to other-regarding behavior is, approximately US\$ 1.00 (US\$ 0.98 for the VCM and US\$ 1.03 for the PPM).

To add further support to our contention that other-regarding behavior causes the observed differences in contribution levels, and to demonstrate the opportunities for inquiry made possible by the virtual-player treatment, we analyzed the data for signs of concerns for fairness.

Subjects in the VCM sessions could calculate the fair cost-sharing contribution by dividing the amount of contributions required to fund the public good entirely (US\$ 100) by the number of players in the group. The fair cost-share contribution was US\$ 4.76 in all VCM sessions. If fairness plays a strong role in determining contribution levels for some subjects, we expect to observe a greater number of fair share contributions in all-human groups than in virtual-player groups. Indeed, we observe this phenomenon.

In human treatments, 14 percent of subjects contributed *exactly* the expected fair share of US\$ 4.76, while none of the subjects who played with automata contributed the fair share.⁷ This difference is highly significant ($P = 0.004$) and fair share contributors accounted for 32 percent of the total contributions. If we relax the definition of a fair share contribution to any contribution between US\$ 4.50 and US\$ 5.00, we find that 29 percent of the all-human subjects contributed the fair share, accounting for 63 percent of the contributions, while only 2 percent in the virtual-player groups contributed in this range ($P < 0.001$). The greater frequency of fair share contributions in the all-human groups also suggests that these calculations are *not* made simply because they are cognitively easier to make. Thus,

⁶ Unlike the Mann–Whitney test, which assumes that the underlying populations have the same general shape and dispersion and are symmetric about the population median, and the Fligner–Policello test, which requires symmetry about the population medians, the Mood–Westenberg test assumes only that the data are from two independent random samples, the measurement scale is ordinal, the variable of interest is continuous, and if, the two populations have the same median, the probability is the same that an observed value will exceed the grand median of the two samples combined.

⁷ We considered contributions at US\$ 4.76 or US\$ 4.77 as exact fair share contributions.

Table 2
 Characterization of subject pools by session

Session	Average age	Percent female	Percent econ major	Average #econ classes	Average anonym score ^a	Average difficult score ^b	Average altruism score ^c
VCM all-human	18.9	41	71	1.1	6.0	6.2	3.2
VCM virtual	18.8	26	84	1.1	6.1	6.0	3.1
PPM all-human	18.5	38	78	1.4	6.3	6.4	3.2
PPM virtual	18.9	37	76	0.84	6.2	5.9	3.1

^a “The procedures followed in this experiment preserved my anonymity”. Scale 1 (strongly disagree) to 7 (strongly agree).

^b “The instructions for the experiment were clear and easy to follow”. Scale 1 (strongly disagree) to 7 (strongly agree).

^c The index was generated from a series of 10 questions (contact authors for questionnaire).

virtual-player designs can also allow researchers to differentiate decisions motivated by human interactions from decisions that are made because they involve lower cognitive costs.

Finally, to ensure that the detected changes in contribution behavior between the all-human sessions and the virtual-player sessions were not generated by differences in subject groups, we collected subject information through a post-experiment questionnaire. The data are summarized in Table 2.

The data indicate that subject characteristics across experimental sessions are similar. Although the proportions of females and economics majors appear to differ in the VCM across cells, the differences are not significant ($P \geq 0.14$). In the PPM cells, however, the difference in the mean number of economics classes taken by subjects does differ significantly ($P < 0.05$); subjects in the all-human session had taken more economics classes. A simple regression of contributions on the number of economics classes taken suggests that if the number of economics classes taken has an impact on contribution behavior, it likely reduces contributions. The coefficient is negative and weakly significant (-0.69 ; $P < 0.09$). A negative impact on contributions would make detecting increases in contributions due to altruistic preferences more difficult.

There is also a difference in subject assessments of the PPM experimental instructions ($P = 0.01$). Although both groups indicate the instructions were generally easy to follow, subjects in the virtual-player session, not surprisingly, found the instructions somewhat more difficult than their counterparts in the all-human session. It is not clear what impact such difficulty may have on contributions, but regressing contributions on difficulty scores generates an insignificant ($P = 0.70$) and negative coefficient (i.e. the easier a subject thought the instructions were, the less the subject contributed). In conclusion, we believe that the differences in contribution patterns across cells cannot be attributed to differences in subject pools.

4. Concluding remarks

The indisputable evidence that individuals derive satisfaction from contributing to the welfare of others or achieving “fair” or “cooperative” outcomes poses fundamental

challenges to economic theory. Despite recent creative experimental designs (Andreoni; Palfrey and Prisbrey; Goeree et al.), our understanding of behavior in simple public goods environments remains incomplete.

By introducing virtual players, we generated data that are consistent with the presence of other-regarding behavior of identical magnitude in two public goods games: the VCM and the PPM. Our results do not support the hypothesis of Ledyard (1995, p. 171) that high average rates of contributions in VCM environments are “the unintended result of a corner non-cooperative solution and not altruism”. Although further work is needed to clarify the precise components of other-regarding preferences that are neutralized through the virtual-player treatment, we reported evidence demonstrating that a non-trivial number of subjects are concerned with fairness.

In addition to corroborating past work on the presence of non-monetary components of the utility function in a simple and direct fashion, the virtual-player technique opens up myriad possibilities for exploring individual decision-making in group settings. For example, it is often believed that in the context of repeated public goods games, the decline in contributions over multiple periods represents a learning effect (e.g. Palfrey and Prisbrey). However, the decline may also result from a decline in warm glow or altruism over time (Andreoni). Using the virtual-player technique, a researcher can isolate the effects of learning from diminishing altruistic behavior. To further discriminate between “warm glow” motivations and “pure altruism” motivations, one could run an experiment that varies the human–virtual agent mix across cells. A “warm glow-only” hypothesis predicts no change in contribution levels as the number of human agents is varied above two.

As Harrison and McCabe have demonstrated, virtual players can be programmed with specific rules of behavior. When these rules are common knowledge, the virtual-player technique can detect if human subjects respond strategically to their opponents’ motives. Thus, the technique has the potential to help us understand learning and optimization skills and strategy formulations in a broad range of games and economic situations.

Control is the essence of experimental methodology. We have demonstrated that the introduction of virtual players can be used to detect, and remove, other-regarding behavior in laboratory experiments. It can also be used to manipulate and control the strategy sets of laboratory subjects. Individuals are motivated by more complex factors than the *Homo economicus* we typically represent with models. We are hopeful that the virtual-player technique can help elucidate these motivations and provide an empirical basis to refine our theories of economic behavior.

Acknowledgements

We acknowledge the financial support of National Science Foundation grant SBR9727375, the US EPA STAR program, and USDA Hatch funds. We thank William D. Schulze for his encouragement in conducting this research and for his essential advice in designing the experiments. Thanks also to Simon Gächter and David Grether for useful comments and to Karen Grace-Martin, of the Cornell Office for Statistical Consulting, for her statistical advice.

Appendix A. Experiment instructions

In order to give the reader an idea of how the virtual-player treatment is implemented, we include the main text of the instructions (without formatting) for the virtual-player treatment of the VCM game. Instructions and timing for the all-human sessions are identical with the exception of the virtual-player language. Each session includes a 5 min oral summary and a question-and-answer period after all subjects have read the instructions. In the oral summary, the non-human nature of the virtual students and the way in which the virtual players made their decisions are emphasized. The full text of all instructions and decision sheets (PPM, VCM, all-human sessions, virtual-agent sessions) can be obtained from the authors or from <http://web.uvic.ca/~rondeau/fileshare.htm>.

A.1. Instructions: VCM virtual-agent treatment

This is an experiment in the economics of decision-making. If you follow these instructions closely and make a careful decision, you can earn money. Please do not communicate with any other student during the experiment.

In today's experiment, you are a member of a group that consists of you and 20 "virtual" students (so if there are 15 students in the lab today, there are 15 groups—each human student is paired with a group of virtual students). The decisions of these virtual students are made by a computer. To start the experiment, we give you and each virtual student an "initial balance" of US\$ 12.00. Once you have read and understood these instructions, you will be asked to enter a "bid" indicating how much of your US\$ 12 you want to invest into a "group investment fund". You can bid any value between US\$ 0 and US\$ 12.

The money that you bid to the group investment fund will be combined with the bids received from the virtual student members of your group. The fund can purchase "shares" at a price of US\$ 1 per share. The fund can purchase up to 100 shares, but no more than that. Hence, all available shares will be purchased if the sum of bids made to the investment fund equals or exceeds US\$ 100.

For every share purchased by the group investment fund, *you* will each receive US\$ 0.07 per share (7 cents per share), up to a maximum of US\$ 7 (100 shares times US\$ 0.07). Note that the virtual students will not receive a payoff because they are not real people.

The virtual students have already submitted their bids, which are in an envelope at the front of the room. Their bids come from a distribution of actual bids made by Cornell students in this exact experiment. In the previous experiments, however, all group members were human; no virtual students were used. Remember, your group is you and the virtual students. None of the other human students in your class are in your group; they are working in different groups. Your final earnings for the experiment will depend on your bid and the bids of your virtual group members.

There are two possible outcomes:

First possible outcome: the sum of bids is *less* than US\$ 100. In this case, all bids will go toward the purchase of shares at US\$ 1 per share. You will receive a personal payoff of US\$ 0.07 per share from the group investment fund. Thus, your earnings for the experiment would be your initial balance of US\$ 12, minus your bid to the group investment fund,

plus your payoff of US\$ 0.07 per share for every share purchased by the group fund (so if 80 shares are purchased, for example, each member of the group receives US\$ $0.07 \times 80 = \text{US\$ } 5.60$). Note that the virtual students are also bidding as if they were to receive a payoff per share of US\$ 0.07. In other words, their bids come from a distribution of bids submitted by human students who, like you, were to receive US\$ 0.07 per share for every share purchased by the investment fund.

Second possible outcome: the sum of bids is *equal to* or *greater than* US\$ 100. If the sum of bids is equal to or greater than US\$ 100, the investment fund will purchase all 100 available shares. Thus, you would receive the maximum payoff of US\$ 7. Your earnings for the experiment would be your initial balance of US\$ 12, minus your bid to the group investment fund, plus your payoff of US\$ 7. Note that no matter how much money is contributed to the group fund, no more than 100 shares can be purchased. Note also that the virtual students, like you, are bidding as if they were to receive a payoff per share of US\$ 0.07. In other words, their bids come from a distribution of bids submitted by human students who, like you, were to receive US\$ 7 if the sum of bids was equal to or greater than US\$ 100.

The instructions end with a one-page summary of the main points, and standard reminders of confidentiality, encouragement to ask questions, etc.

References

- Andreoni, J., 1995. Cooperation in public goods experiments: kindness or confusion? *American Economic Review* 85, 891–904.
- Champ, P.A., Bishop, R.C., Brown, T.C., McCollum, D.W., 1997. Using donation mechanisms to value nonuse benefits from public goods. *Journal of Environmental Economics and Management* 33, 151–162.
- Davis, D.D., Holt, C.A., 1993. *Experimental Economics*. Princeton University Press, Princeton.
- Fligner, M.A., Policello II, G.E., 1981. Robust rank procedures for the Behrens–Fisher problem. *Journal of the American Statistical Association* 76, 162–168.
- Forsythe, R., Horowitz, J., Savin, N.E., Sefton, M., 1994. Fairness in simple bargaining experiments. *Games and Economic Behavior* 6, 347–369.
- Goeree, J.K., Holt, C.A., Laury, S.K., 1999. Altruism and Noisy Behavior in One-Shot Public Good Experiments. Manuscript, Department of Economics, University of Virginia.
- Güth, W., van Damme, E., 1998. Information, strategic behavior, and fairness in ultimatum bargaining: an experimental study. *Journal of Mathematical Psychology* 42, 227–247.
- Harrison, G.W., McCabe, K.A., 1996. Expectations and fairness in a simple bargaining experiment. *International Journal of Game Theory* 25, 303–327.
- Isaac, R.M., Walker, J., Thomas, S., 1984. Divergent evidence on free riding: an experimental examination of possible explanations. *Public Choice* 43, 113–149.
- Kagel, J.H., Roth, A.E. (Eds.), 1995. *Handbook of Experimental Economics*. Princeton University Press, Princeton.
- Ledyard, J., 1995. Public goods: a survey of experimental research. In Kagel, J.H., Roth, A.E. (Eds.), *Handbook of Experimental Economics*. Princeton University Press, Princeton, pp. 111–194.
- Mood, A.M., 1950. *Introduction to the Theory of Statistics*. McGraw-Hill, New York.
- Palfrey, T.P., Prisbrey, J.E., 1997. Anomalous behavior in public goods experiments: how much and why? *American Economic Review* 87, 829–846.
- Rondeau, D., Schulze, W.D., Poe, G.L., 1999. Voluntary revelation of the demand for public goods using a provision point mechanism. *Journal of Public Economics* 72, 455–470.
- Westenberg, J., 1948. Significance test for median and interquartile range in samples from continuous populations of any form. *Akad. Wetensch. Afdeeling Voor de Wis* 51, 252–261.