1

Pirates Acting for the Public Good: An Experimental Study

Seung-Keun Martinez
University of California, Berkeley

January 31, 2013

Abstract

We study the structure of online peer-to-peer piracy (illicit file sharing) to show that cooperation for a public good can be established and self-sustained without exogenous intervention and despite costly entry. We use laboratory experiments to study real-time contribution strategies in a continuous public good game with endogenous group formation. We show that even when subjects must pay to access the public good and the breakeven point requires at least half of the participants contribute to it, it is still possible to achieve a cooperative outcome. For this to happen, cooperation must be established early in the game. Once established, cooperation lasts late into each session. Natural language communication is crucial: when subjects communicate via anonymous instant messaging, they achieve equivalent levels of cooperation (with and without costly entry) amounting to contributions of 80% of their total endowment. Despite the rich signaling possibilities offered by variations in real-time contribution rates, cooperation is significantly lower absent the ability to communicate with words.

*The author thanks Prof. John Morgan and Carlos Cueva for crucial guidance and insightful comments. This research is made possible by the generous financial support of Bob and Colleen Haas and the sponsorship of Prof. Jochen Menges who advocated for the author's stay at the University of Cambridge. The author also thanks Prof. Dan Friedman for allowing the use of ConG, a program developed by the UCSC Leeps Lab, and James Pettit in his assistance with this program.

1 Introduction

In recent years, an internet phenomenon known as peer-to-peer (P2P) piracy has blossomed into vast networks that have up to 18 million visitors per day.[1] A P2P piracy community consists of individual computers sharing data through online torrents. A torrent file is used to connect two or more users so that they can send data to one another. In order for these communities to operate, members must be willing to "seed" data, which allows others to download it from their computer. There is no direct benefit to seeding. Rather, seeding consumes bandwidth and puts the seeder at risk of being charged with copyright infringement. The uploaded data is freely available to anyone who wishes to download or "leech" it. Furthermore, users may freely change the rate at which they are seeding and downloading; either rate may be zero while the other is strongly positive. All P2P users are free to leave and rejoin the network at any time. That is, users choose when to connect and disconnect from torrents, which either enables or disables all seeding and leaching activity. The difference between exiting the network and contributing nothing is that one may not leech data nor be exposed to legal risk while outside of the network. Thus, P2P networks act like non-excludable public goods with free riding incentives. P2P networks are an intriguing example of public goods because not only are they self sustaining but are also endogenously created. Specifically, every member of a piracy network elected to join the network at a cost. Members also continue to join and contribute, establishing the vast networks that currently exist.

P2P piracy is a unique public good. Seeding, leaching, and messaging occur continuously, and information concerning the amount of seeding and leaching is available in real-time. Torrents in *The Pirate Bay* (the most popular P2P community) all display the number of users sharing the data and the number that are downloading it. Furthermore, each torrent also comes with its own message board where users may post free form messages to one another. Messages are sent anonymously but are visible to everyone. In this paper we investigate under which conditions the success of P2P piracy can be replicated. In order for our results to be applicable to public goods other than piracy networks, we conducted a public good game experiment. We used experiments instead of gathering data directly from piracy communities for two reasons. First, such data is not readily available. From a network administrator's viewpoint, empirical data concerning the growth and sustainment of illegal sharing is best kept secret. Secondly, behavior in P2P networks is likely influenced by anti-copyright culture. Although many torrent users are largely motivated by the availability of free content, there is also a clear anti-copyright sentiment in piracy communities. By studying the other unique properties of P2P piracy experimentally, we avoid confounding our data with any preconceived social bias. Therefore, the results presented in this paper can offer policy suggestions for other public goods.

We adapted the public good game to run in continuous time for alll treatments. Rather than having participants choose amounts of tokens to place into private or group accounts period by period, we instead asked them to choose a rate of contribution that they could alter freely. The final contribution was equivalent to the averaged

[1] Wolfram-Alpha, accessed on November 25, 2012, http://www.wolframalpha.com/input/?i=thepiratebay.se

rate according to the time spent at each rate. The remaining tokens were placed into each player's private account. Participants also accumulated points from the group account continuously. The rate of accumulation from the group account was exactly proportional to the summed rates of contribution. Our experiment had two treatment parameters: communication and endogenous entry and exit. In treatments with communication, we allowed free form messages to be sent via instant messaging. These chat rooms resembled online message boards; they were anonymous and visible to everyone. In endogenous treatments all subjects started "out" of the group account and could only contribute to and receive points from the group account when choosing to be "in" the group account. However, being in the group account bore a continuous cost.[2] Participants could leave and rejoin the group account freely. Lastly, each subject was fully informed about the decisions all other subjects faced and made.

The main innovation of our experiment lies not in the the difference between continuous and discrete public good games but in the addition of costly participation in continuous time. Our experiment permits us to examine how the combination of communication and endogenous participation affects giving and to decompose the constituent parts of the environment. Just as P2P network use has a cost of bandwidth and a risk of copyright infringement charges, our experiment has a cost of being in the group account. All participants start outside of the group account in endogenous treatments. This allows us to study how cooperative outcomes may be endogenously accomplished and sustained under costly entry and free riding incentives.

Our major results are as follows:

(1) Communication significantly increases cooperation across all treatments. Furthermore, communication allows for cooperative outcomes even with costly participation.

(2) Convergence to cooperative strategies is only observed when communication is permitted.

(3) Endogenous participation significantly lowers cooperation absent communication.

(4) Behavior observed early in the game persists throughout the session.

The above results are closely related. First, we find that communication increases contributions to the public good (1). One reason messaging increased cooperation may be that it also allows participants to pressure those who would otherwise be uncooperative into contributing. However, we find that the coordinating opportunities afforded by chat are the primary reason for high cooperation. We find that these coordinating opportunites are sufficient to overcome costly participation. In this manner, communication facilitates cooperative strategy convergence (2). *i.e.* we only observe participants converge and commit to high contribution rates when chatting is permitted. Without messaging, subjects would constantly raise and lower their contribution rates. However, communication allowed them to coordinate their strategies and agree to high contributions throughout the session. Interestingly, this high rate of contribution is either established early in the game or not at all (4). After first impressions are made, players will either continue to work with each other or will distrust each other regardless of any new attempts to foster cooperation. We find that endogenous entry and exit hinders the establishment of early positive impressions. Costly participation makes subjects

[2] To make the treatments identical, we charged the same cost in exogenous treatments as well. However, since it was unavoidable, this should have no effect on choices.

more hesitant to be among the first to give to the public good. Thus, early cooperation is not established, and this hesitance continues for the rest of the game (3).

1.1 Literature Review

To the best of our knowledge, the impact of continuous time on the public good game was first studied by Dorsey (1990). However, real-time in his experiments meant that participants could change their contribution pledge until the end of the time period, where that final pledge was contributed. Our continuous time experiment resembled that of Friedman and Oprea (2010) and Charness, Friedman and Oprea (2012). In Friedman and Oprea's "A Continuous Dilemma" (2010), they establish that having continuous play greatly increases cooperation levels in the two player prisoners dilemma game. Strikingly, the median cooperation rate reached a stunning 90% in continuous time while almost no participants cooperated in the one-shot game. Charness, Friedman, and Oprea (2012) found similar results in a four-player public good game. Their experiment contained both continuous and discrete time treatments. In the continuous time treatment, participants contributed to the public account by choosing a rate of contribution. This rate was then averaged over the time spent at each chosen rate to determine the final contribution. Subjects could observe the real-time strategies of all other players. In discrete time, each session was broken into 10 time periods where only the instantaneous contribution rate at the end of each period was recorded. Each observation was shown to every other player, and the final contribution was determined by the average of each subject's ten observations. Their results showed that continuous time treatments produced significantly higher cooperation. They argue that continuous play is more conducive to cooperation because it allows subjects to display "pulsating" behavior. That is, subjects could temporarily move to higher rates of contribution to signal a willingness to cooperate to other players. Continuous flow payments allow participants to engage in such signaling while only risking a small proportion of their endowment. In other words, real-time interaction offers an array of strategic options not available in discrete games. Since every participant observes everyone's strategies at all times, there are greater signaling opportunities available in continuous time.

Our first treatment variable is communication. We study the effects of communication in continuous time with endogeneity for two reasons: the effectiveness of communication in the previous public goods literature and the presence of message boards in P2P sharing. Isaac and Walker (1988) found that communication not only induces high levels of cooperation in the discretely repeated public good game, but also that cooperation can be sustained after communication has ceased. In their experiment communication was face-to-face, but individual contribution remained unknown. The experiments in Charness *et al.* (2012) exhibited a median cooperation rate of 100% in a four-person continuous public good game with chat-box instant messaging. Furthermore, they found no decay of contributions in treatments with communication. In our experiments we test the robustness of communication with endogenous entry and exit (our second treatment variable). We used the same instant messaging mechanism as Charness *et al.* (2012). We find that communication not only encourages cooperation, but that it may be a necessary condition. Indeed, the only sessions which displayed a convergence to sustained high contributions were those with communication.

Building off of Charness *et al.* (2012), we ran an eight person public good game in continuous time. However, the primary importance of continuous time in our experiment is that it allows us to explore our second treatment variable: costly endogenous participation. The previous literature concerning endogenous entry and exit focused on using it as a tool to encourage cooperation. Charness and Yang (2008), Ahn *et al.* (2008), and Cinyabuguma *et al.* (2005) all found that the ability to exclude players from a group increased cooperation. Each researcher used voting rules that could restrict group entry, exit, and formation, and all were seeking endogenous entry and exit mechanisms that would increase cooperation. Our implementation differs markedly from these studies. First, the group cannot exclude individuals from enjoying the benefit of the public good – not participating is a individual decision. Second, participation is not unambiguously beneficial. Because of participation costs, individuals may prefer to not participate. This is consistent with how torrenting comes with a risk of being charged with copyright infringement or how one has listen to funding campaigns on public radio. In both of these examples one has to choose to access a public good at a cost to herself. In this paper we examine how real-time coordination and communication may be used to overcome such impediments and establish cooperation.

2 Experimental Design

2.1 Basics

We conducted the experiment at the University of Cambridge. Participants, who were either undergraduate or graduate students at the university, were randomly selected using an online recruitment system. Subjects were recruited for only one session each. The general design is a 2x2 treatment public good game in continuous time with and without both endogenous entry and exit and communication.

All subjects were given 20 tokens at the beginning of each treatment. They were told that they could either put the token into their private accounts or a group account. A token put into a private account was worth one point. A token put into the group account would become four points and be split eight ways, implying a marginal per capita return (MPCR) of 0.5. All treatments had eight players playing a continuous time public good game for ten minutes. Subjects were seated at individual computer terminals with dividers placed in between them. Each subject received the same set of instructions as he arrived to the lab. We had every participant complete a comprehension quiz immediately after they read the instructions. The quiz focused on gameplay and payment comprehession, especially on how their decisions affect their own and other's payments. Subjects were asked to fix incorrect answers on their own quizzes. Fortunately, we did not have to dismiss any participants for inability to pass the quiz. Participants played using the program ConG. ConG was developed by the UCSC LEEPS lab and is the same program used by Charness, Friedman, and Oprea (2012) and Friedman and Oprea (2010). The user interface of this program is shown in Figure 1. In order for everyone to become familiar with the program, we ran a three minute simulation without pay prior to the paid game. Sessions lasted 30 to 40 minutes with an average payment of eight pounds.

Similar to how members of a P2P network may alter their contributions at any point in time, independent of others, we ran the experiment in continuous time. Continuous time meant strategies could be altered freely at any point during the experiment. Furthermore, the computer response time is less than 100 milliseconds, allowing participants to react to each others' actions almost instantaneously. There is an important distinction between the continuous and discrete public good games. In discrete games, everyone must simultaneously make an *active* decision as to how much to give. In continuous time, actions do not have to occur at the same time, and giving can be done passively. If a subject had previously set his contribution to amount x, it will remain there until he decides to change it.

Figure 1: **User Interface**

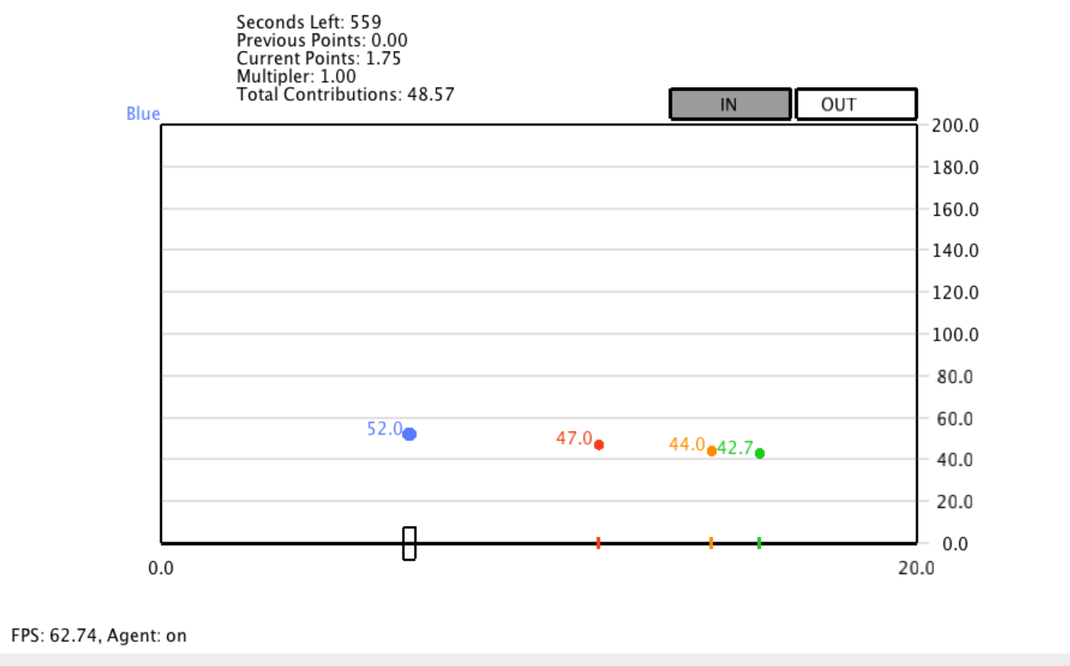

Each subject chose a rate of contribution by using the rectangular slider at the bottom of the screen (shown under the blue dot and "52" in Figure 1). The starting position of the slider was set to zero so that all initial contributions would have to result from an active choice. She would also see a corresponding colored dot exactly above her slider. Each participant was randomly assigned a unique color. The assigned color appeared on the top left corner of the program screen. The height of each subject's colored dot represented his rate of point accumulation. The contribution and accumulation rates of the other players also appeared on screen as dots according to the other participants' assigned colors. The number floating next to each dot is the rate of accumulation. Participants could change their rate of contribution at any time during the experiment. Total contributions and payoffs were calculated as flow payments by averaging rates of contribution and accumulation over the ten minute time period. This is further explained in the next section.

Having real-time interaction also allowed us to study dynamic coordinating behavior with and without endogenous entry and exit. In the endogenous treatments, two buttons appeared on the top right corner of the screen. As shown in Figure 1, these buttons were labeled "in" and "out." Selecting "in" meant that the participant was in the group account, selecting "out" meant she was out of the group account.[3] Whether a participant was in or out was indicated by which box was darkened in Figure 1. A participant must be in the group account to contribute and to accrue points from it. Thus, choosing "out" implies a rate of contribution of zero and a rate of accumulation of equal to one's starting endowment. This was graphically represented in one's own user interface by having each person's colored dot move to a fixed height (equal to their starting endowment) regardless of the position of the slider. The same dot would disapear off of other subjects' screens when one was out of the group account. Being in the group account had a flow cost of one point per minute, implying ten points if the subject were in the group account for the entirety of the game. Participants all started out of the group account in endogenous treatments. For the sake of comparison, subjects in exogenous treatments also faced the ten point cost.

Our second treatment variable was communication. In our baseline subjects could not communicate with each other. In communication treatments, participants could converse freely via instant messaging through an on-screen chat box. Chat box screen names were color coded according to each subject's randomly assigned color. That is, each person's screen name was her assigned color. This screen name also appeared in the assigned color. The chat box was only available during the 10 minutes of paid gameplay. Participants were told they may discuss anything pertaining to the experiment, but were asked to refrain from inappropriate language.

2.2 Payoffs

In this section we first describe the payoff function for exogenous treatments. In each treatment each of the n participants starts with a token endowment of *sMax*. During the course of each ten minute treatment, participants choose a rate of contribution $s_i \in [0, sMax]$. This rate of contribution is then averaged over the ten minute span according to how long was spent at each chosen rate s_i to determine the average individual contribution:

$$S_i = \frac{1}{T}\int_0^T s_i(t)\,\mathrm{d}t \in [0, sMax]$$

Each participant's payoff is then calculated as:

$$P_i = sMax - S_i - C + \frac{A}{n}\sum_{i=1}^{n} S_i \ : 0 < A/n < 1 < A$$

Where C is a fixed cost. Since A is greater than 1, it is socially optimal for every player to contribute all of her tokens to the public account. However, A/n is less than 1, which implies free riding incentives.

The difference with endogenous treatments is that players decide whether or not to be a part of the group account. Everyone being "in" the group account for the entirety

[3] Slider position was not effected by in and out choices

of the experiment is exactly the same as being in the exogenous treatment. When a participant chooses to be "out" of the group account they do not accrue payoffs from the group account nor do they incur the per minute cost of $c = C/10$. *i.e.* when choosing to be out of game, players earn at a rate exactly equal to $sMax$. Formally, S_i and C_i are now calculated as:

$$S_i = \frac{1}{T}\int_0^T 1_i * s_i(t)\,\mathrm{d}t$$

$$C_i = \frac{1}{T}\int_0^T 1_i * C\,\mathrm{d}t$$

Where 1_i is an indicator variable such that $1_i = 1$ when player i is in the group account. Thus, the payoff for player i is calculated as:

$$P_i = sMax - S_i - C_i + \frac{A}{n}\int_0^T \sum_{i=1}^{n} 1_i * s_i$$

Participants start every endogenous treatment out of the group account. It is important to note that n does not vary in the payoff equation according to how many players are in the group account. Otherwise it would be individually beneficial for a lone member of the group account to contribute all of his tokens. This would not reflect the decisions faced by members of a P2P network, nor would endogenous entry and exit act as a barrier to entry.

2.3 Chosen Parameters

All treatments had group size $n = 8$ token endowment $sMax = 20$, multiplier $A = 4$, and cost $C = 10$. The chosen parameters yield a MPCR of 0.5. We chose these parameters to emulate the basic structure of P2P piracy while imposing the worst possible conditions for cooperation. The rationale being that if cooperation succeeds under these conditions, it would likely succeed over other parameter values. The first cost a P2P user faces is the risk of copyright infringement charges. This is represented by the one point per minute cost. We chose this fixed cost because, in P2P piracy, those who benefit from the good, but do not contribute also face the risk of infringement charges. We chose these specific numerical values to reflect how P2P piracy (and many other public goods) are only beneficial if a significant number of people contribute. If a P2P piracy group is too small, it cannot be sustained because there is too little data for each person to share and download. Eventually, users would become bored of the limited amount of content and the network would collapse. Thus, our chosen parameters require that at least four of the eight participants contribute more than half of their endowment for the public good to be socially beneficial. Less than four entrants in the public account cannot yield a socially beneficial outcome with our chosen parameters.

Secondly, setting MPCR $= 0.5$ with a flow cost of one point per minute implies that a participant must have a positive rate of accumulation. If no one else is contributing, then setting one's own rate of contribution to 20 yields a payoff of:

$$P_i = 20 - 20 - 10 + .5(20) = 0$$

Thus, a MPCR = 0.5 and a cost of one point per minute is the most stringent condition under which we do not allow negative payoffs. Clearly, negative payoffs are problematic because we cannot pay our subjects negative amounts.

2.4 Theory

In this section we identify conditions where the unique subgame perfect equilibrium of our experiment is zero contribution in all treatments. First we define some notation. Let $s_i(t)$ = rate of contribution by player i at time t. Let $a = A/n$ where A is the point multiplier implying a is the MPCR. Thus player j's rate of accumulation at time t is:

$$sMax - s_j(t) + a\sum_{i=1}^{n} s_i(t) - C$$

The end payoff to player j is:

$$P_j(t) = 1/T\int_0^T sMax - s_j(t) + a\sum_{i=1}^{n} s_i(t) - C\mathrm{d}t$$

Suppose every player is using grim trigger strategies with strategy profile $s_i(t)$ = sMax if $s_{-i}(t)$ = sMax for t $\in$ [0,t-δ] and $s_i(t) = 0$ if $s_{-i}(t) <$ sMax for t $\in$ [0,t-δ], where δ is the necessary time for others to react to one's strategy deviation. It is then individually profitable to refrain from defecting at time x so long as:

$$1/T[\int_x^{x+\delta} sMax + a(\sum_{i=1}^{n} s_i(t) - s_j(t)) - C\mathrm{d}t + \int_{x+\delta}^{T} sMax - C\mathrm{d}t]$$

$$\leq$$

$$1/T\int_x^T sMax - s_j(t) + a\sum_{i=1}^{n} s_i(t) - C\mathrm{d}t$$

That is the relative gain made in the δ time before everyone else reacts is less than the relative losses sustained in remainder of the session. Notice that this inequality is trivially false at any $x \in [T-\delta, T]$. Thus, everyone should deviate at $x = T-\delta$. But if everyone deviates at $x = T-\delta$, then by the same logic, one should deviate at $x = T-2\delta$. Formally, if everyone deviates at time $x+\delta$, then the above inequality does not hold at time x. Furthermore, since T is finite, T/δ is finite. Thus, for any t $\in$ [0,T], there exists a finite m number of iterations of the above process such that $t \in$ [T-(m+1)δ,T-(m)δ]. This implies that all players should defect at all times t and that the subgame perfect equilibrium is no contributions. The proof for the endogenous case is entirely similar. The only difference being that the inequality now becomes:

$$1/T[\int_x^{x+\delta} sMax + a(\sum_{i=1}^{n} s_i(t) - s_j(t)) - C\mathrm{d}t + \int_{x+\delta}^{T} sMax\ \mathrm{d}t]$$

$$\leq$$

$$1/T \int_x^T sMax - s_j(t) + a \sum_{i=1}^{n} s_i(t) - C \mathrm{d}t$$

i.e. Player j now earns at a rate of sMax after defection rather than a rate of sMax - C.

On the contrary, cooperation can be supported as a ε-equilibrium. A profile of strategies $\{s_1(t), ..., s_n(t)\}$ is an ε-equilibrium if no player has deviation that produces more than ε additional profit. Let δ denote reaction time. Under grim trigger strategies, the largest gains from deviation occur at time $t^* = T - \delta$. Suppose not, then either (1) $t^* > T - \delta$ or (2) $t^* < T - \delta$. If (1), then a player would be forgoing a higher earnings rate when $t \in [T - \delta, t^*]$ and the other players would not be able to react in time to punish him afterwards. If (2), then a player increases her payoff for the same length of time δ, but is punished for the remaining time left in the game. For both endogenous and exogenous games, the gains for deviating at time $T - \delta$ for player i are:

$$\Delta\pi_i = 1/T \int_{T-\delta}^{T} (1-a) s_i(t) \mathrm{d}t$$

Notice that

$$\lim_{\delta \to 0} \Delta\pi_i = 0$$

Therefore, for every $\varepsilon > 0$, there exists $\underline{\delta}$ such that, for all $\delta < \underline{\delta}$, full cooperation is a equilibrium.

The first proposition illuminates the similarities between the continuous time game and the repeated discrete game. That is, given that there is an asynchronicity between one player's action and an other's reaction, we still obtain an unique subgame perfect equilibrium of zero contributions. However, this prediction is inconsistent with our experimental results. In contrast, the notion of an ε-equilibrium allows for cooperative equilibria. Furthermore, the ε-equilibrium better captures a subtle difference between continuous time and discretely repeated games. We will explain this difference by using our experiment as an example. Suppose we added another exogenous treatment where the public good game was played in ten discrete rounds – each round worth one tenth of the total payoff. With our chosen parameters, ε-equilibrium requires $\varepsilon > 1$ point (0.1pounds). However, in the continuous time treatment, if we assume $\delta = 1$ second, ε-equilibrium requires $\varepsilon > 0.017$ (0.0017 pounds). The notion of an ε-equilibrium captures how decisions made in the continuous game may have smaller payoff consequences than in the discrete game. In the ε-equilibrium, the smaller payoff consequence implies that cooperation is easier to maintain. Theoretically, one could construct a 600-period discrete game. However, this would prove impractical.

3 Results

In this section we present the main findings of our experiment. The results are presented in the reverse order of causation. Results 1 and 2 test our main hypotheses. Results 3, 4, and 5 elucidate why we observe Results 1 and 2.

Figure 2

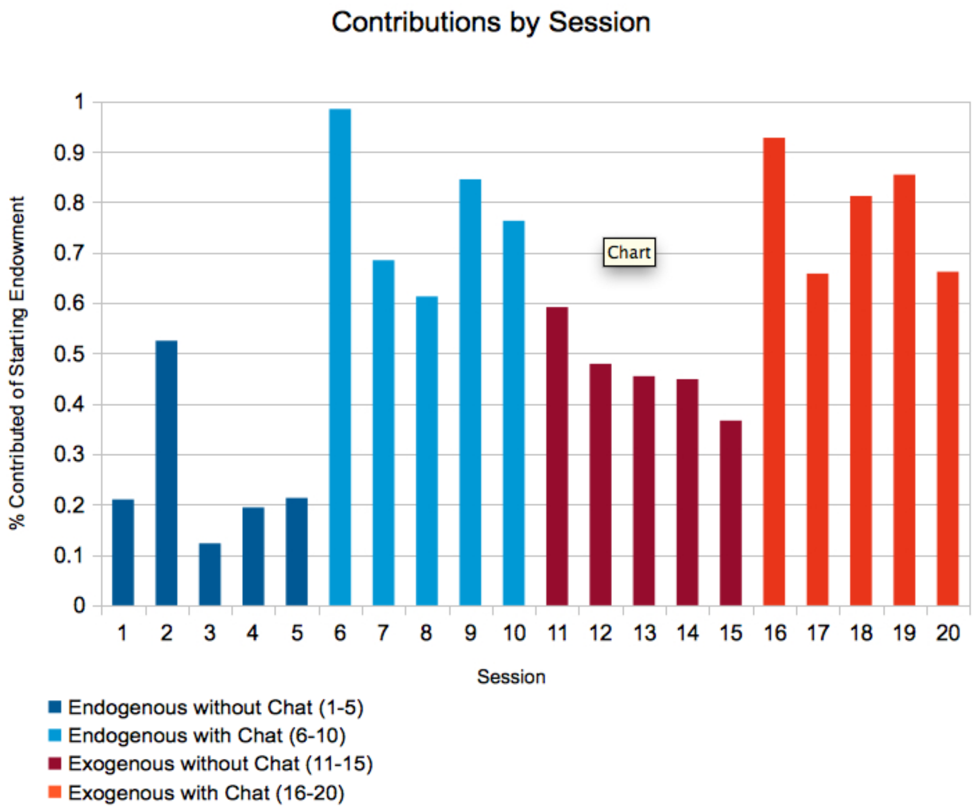

Result 1: Communication significantly increases cooperation in endogenous and exogenous treatments.

Support for Result 1 is presented in Figure 2. Columns 1 through 5 show the mean cooperation rates of endogenous treatments without chat, while columns 6 through 10 show those of endogenous treatments with chat. The overall contribution to the public good in endogenous treatments without communication is about 25% of the starting endowment. By contrast, the addition of a chat box more than triples the average contribution to 78%. The obvious economic significance of this result is supported by a Wilcoxon rank-sum test. Testing the null hypothesis of no treatment effect we find a p-value of 0.01. Columns 11 through 15 show mean cooperation rates for exogenous treatments without chat, and 16 through 20 show mean cooperation for exogenous treatments with chat. On average, participants contributed 47% in the no chat treatment and 78% in the chat treatment. The difference between the two treatments is again significant with a p-value of 0.01 under a Wilcoxon rank-sum test.

As one may expect, communication induces cooperation. Our data agrees with the findings of Charness *et al.* (2012): that this positive effect is robust over continuous time. Moreover, our data suggests that anonymous but unrestricted communication is sufficient to establish cooperative outcomes despite costly participation. Further analysis (shown in the following results) indicate that subjects are generally willing to cooperate as long as others are willing to as well. This requires that enough people be willing conditional contributors and that someone initiates the movement to cooperation. Thus, communication acts as a coordinating mechanism to establish mutual cooperation. The lack of cooperation in the absence of communication is also telling. Real-time interaction allows for continuous signaling. This signaling is observed through pulsating behavior. Participants would pulsate by frequently entering into the group and contributing for a short period of time even when it was strictly unprofitable to do so. We interpret this behavior as an encouragement for others to join and give to the public good. However, our results show that this behavior was unable to establish a socially beneficial equilibrium. This outcome is further explained in the following sections.

The use of communication to encourage investment in the public good is also observed in P2P piracy. On torrent message boards people consistently post comments to thank seeders and encourage others to cooperate. Similarly, our experimental subjects would also respond positively to cooperation and negatively to defection via chatting. For instance, in session six, shown in Figure 2, after unilateral full cooperation is achieved we observe this string of messages:

Red: From this I can surmise that you are all lovely people
Blue:that's great
Red: :)
Blue: i love you too

However, in session 9, when player "Aqua" displayed free-riding behavior we observe these messages:

Purple: lets share the prize aqua
Gray: there's always an individualist :) seriously.
Purple: united we stand :P
Aqua: Well now I feel mean
Gray: they earn less but they feel "free" :)
Green: not enough!
Blue: you should feel mean - if we all go at 20 we all earn as much as possible

As we can see, participants would try to encourage others to contribute to the public good with positive and negative language. However, we show in the following results that the primary importance of anonymous messaging lies in its coordinating possibilites.

Result 2: Costly endogenous participation lowers cooperation only in treatments without messaging.

Support for Result 2 also comes from Figure 2. The average cooperation rates in both endogenous and exogenous treatments with chat are 78%. However, the average cooperation rates of endogenous and exogenous treatments without chat differ by 22%. In the endogenous treatment we observed an average contribution rate of 25%, while the exogenous treatment had an average rate of 47%. We again used the Wilcoxon rank-sum to test the null hypothesis of no treatment effect. Using this test, we reject the null with p-value equal to 0.076. Strikingly, the same test yields a p-value of 0.92 for the treatments with messaging. From this we may extrapolate that we can reject the null hypothesis that there is a treatment effect at the 10% level. That is, subjects contributed equivalent amounts to the public good in the endogenous and exogenous game when allowed to message.

Both the difference in contributions in the treatment without chat and the lack of a difference in treatments with chat are telling. In the following results we argue that the coordination of strategies that communication allows is the primary reason anonymous chat induces cooperation. Thus, Result 2 shows that with a sufficiently strong coordinating mechanism, endogeneity has no effect on contributions. However, when no such mechanism exists, costly entry makes establishing cooperation more difficult. We further explain why cooperation is lower in the endogenous treatment absent messaging in Result 4.

Table 1: **Total Variation of Cost by Treatment**

	With Entry & Exit Variation		**Without Entry & Exit Variation**			
	Endogenous		Endogenous		Exogenous	
Sessions	Chat	No Chat	Chat	No Chat	Chat	No Chat
1	1.22	15.25	0.72	6.21	10.59	2.89
2	5.14	14.56	4.02	13.4	10.59	4.00
3	14.44	14.57	7.26	4.44	11.39	4.78
4	3.11	17.78	1.99	9.15	9.59	2.22
5	3.25	28.38	2.38	12.80	20.30	12.86

Note: TVC is a measure of variation of instantanous costs a player faces according to how much he changes his strategies. "With(out) Entry and Exit Variation" indicates that TVC was calculated with(out) the changes in costs that occur when entering and leaving the public account.

Figure 3:

Figure 4:

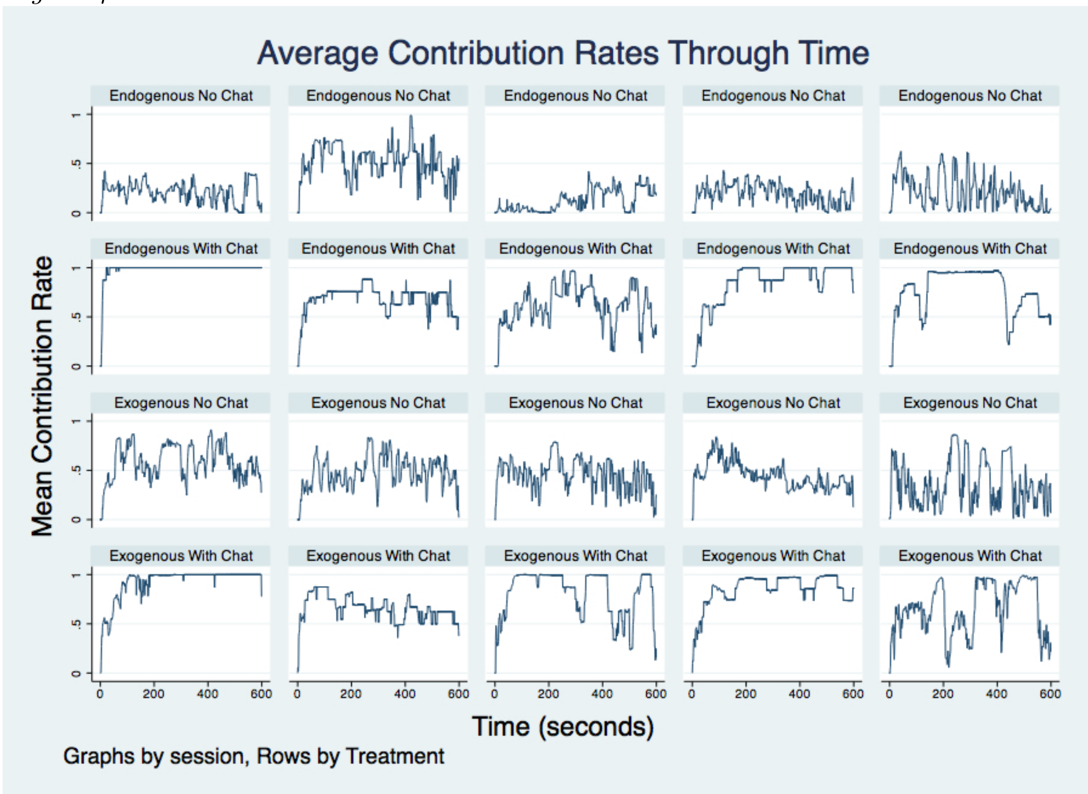

Note: "Mean Contribution Rate" indicates the average contribution rate chosen at each second.

Result 3: Convergence to cooperative strategies is only achieved with communication.

Support for Result 3 comes from Figures 3 and 4 as well as a summary measure of Total Variation. Figure 3 shows the proportion of participants in the group account. Figure 4 shows average contribution changes over time for each season. The graphs are normalized so that full contribution is denoted as a rate of 1. Figures 3 and 4 show two clear trends. First, in treatments with communication, participants generally entered into the public account from the beginning and remained in for the rest of the game. In contrast, we see frequent entry and exit in the no communication treatment. This indicates that players would leave and rejoin the public account multiple times. Secondly, we see very few players in the no communication treatment staying at fixed contribution rates for an extended period of time. This is shown by the number of subjects leaving and entering the public account in Figure 3 and by the drastic dips in cooperation shown in Figure 4. In contrast, modal behavior in any given time interval in the communication treatment is full contribution. In other words, while observed strategies are constantly changing in the no communication treatment, there is evidence of equlibria with high contribution levels when communication is allowed.

Charness, Friedman, and Oprea (2012) use the measure of total variation to quantify whether play is in a steady state. *i.e.* whether participants are staying at fixed contribution rates. Total variation measures the changes in individual strategy. Under steady state play, total variation would be zero. Higher values of total variation represent larger departures from steady state. In our setting, we must adapt total variation to account for entry and exit choices, which are absent in Charness *et al.* (2012). To quantify these results we employ the measure of Total Variation of Cost (TVC). TVC is an intuitive measure of variability computed by summing the absolute changes in individual cost which only depend on each players own choices. We define cost (C) to be the rate at which a player is losing tokens. The cost of contributing at a rate of x tokens is $0.5 * x$ since MPCR = 0.5. Where we normalize x so that $x = 1$ implies full contribution. The cost of being in the public account is 0.5 because that is the cost rate for ones presence in the account. i.e. if player i chooses a contribution rate x, then $C_i = 0.5 * x + 0.5$ If player i is out of the group account, then $C_i = 0$. Thus, TVC is given by:

$$TVCi = \sum_{t=2}^{n} |C_{it} - C_{it-1}|$$

where t indicates the time period. We calculated TVC in one second time periods since ConG outputted data in a per-second basis. We use TVC instead of total variation of contribution because TVC captures both contribution variation and entry/exit variation in a consistent manner. That is, TVC weights the variation due to entry and exit and changing contribution rates according to the exact costs to participants.

Table 1 shows TVC per session. The average TVC for endogenous treatments with chat is 5.44. Without communication, this measure more than triples to 18.11. We use a Wilcoxon rank-sum test with null a hypothesis that there exist no difference in TVC and find p = 0.01. Thus, chat shifts behavior much closer to steady state play. We find similar results in the exogenous treatments. In the exogenous treatments, we find an average TVC of 5.35 with chat and an average of 12.49 without it. Using the same Wilcoxon rank-sum test, we reject the null hypothesis of no treatment effect with p-value equal to 0.075.

In a way, the effectiveness of communication is surprising. In theory, the richness of signaling opportunities afforded in a continuous time setting might seem to be enough to establish cooperation. Indeed, subjects take advantage of modes of signaling and appear willing to cooperate. For instance, many subjects enter for a short time and contribute fully before exiting. Often, this behavior is repeated several times in succession as a kind of beacon indicating a willingness to cooperate. However, these signals only prove to be an imperfect substitute as cooperative and sustainable equilibria are never established when the chat box is removed. Specifically, other players do not respond in sufficient time and with magnitude to catalyze a cooperative equilibrium.

In terms of the P2P setting, our results suggest that rich avenues for real-time communication are essential to a thriving network. Coarser signaling mechanisms such as link seed popularity and other indirect feedback are likely to be vastly less effective. From a policy perspective, a regulator seeking to curtail these networks may achieve much of this goal by disrupting communications on a network; this may be easier than shutting down the network itself.

Table 2: ***Average Contribution Rates in No Chat Treatments***

	Average Cooperation Rates		
	Endogenous		Exogenous
Sessions	Total	While Using Public Good	Total
1	0.21	0.27	0.59
2	0.52	0.54	0.48
3	0.12	0.19	0.45
4	0.19	0.29	0.45
5	0.21	0.34	0.37

Result 4: Without chat, players use more conservative signaling strategies in endogenous treatments. i.e. players signal a willingness to contribute in both no-chat treatments, but do so with lower cooperation rates under endogeneity

As shown in Result 2, the cooperation is lower under endogenous versus exogenous treatments. However, Result 2 does not explain why cooperation is lower in the endogenous no-chat treatment than in the no-chat exogenous treatment. Curiously, Figure 4 shows that steady state cooperative play is not achieved in either no-chat treatment. Yet, the average cooperation rate in the exogenous treatment (47%) is double that of the endogenous treatment (25%).

The the reason for Result 2 is two-fold. First, the average participation rate in the no-chat endogenous treatment over the course of ten minutes is 73%. This equates to no possible contributions for 27% of the time during the endogenous treatment. Secondly, the average in game contribution rate in the endogenous treatment is 33%, while the exogenous has an average rate of 47%. In other words, not only does absence from the group account cause lower cooperation, but in account contributions are lower as well. By repeating the same rank-sum test with only the in-account data we still reject the null hypothesis with a p-value equal to 0.075.

This result illustrates why costly entry and exit has a deleterious effect on cooperation. On the surface we see that players in both treatments display pulsating behavior, but players in exogenous treatments pulse to higher contribution rates. One interpretation of this result is that players in exogenous treatments understand that they will bear the full cost of being in the group account regardless of their actions. Their only hope of overcoming this cost is by establishing sufficient contribution to the public good from multiple people. Conversely, endogenous treatment participants are not forced to bear the cost of the public good. Thus, for endogenous treatment participants, signaling high rates of contribution is simply more costly relative to their unilateral no cooperation payoffs. As a result, pulsating behavior in endogenous sessions tends to display a willingness to enter before displaying a willingness to contribute. This strategy leads to low in-account contribution rates. This, in turn, incentivizes staying out of the group account. Concisely, endogenous entry and exit exacerbates the need of an exogenous coordinating mechanism to encourage cooperation.

Figure 5

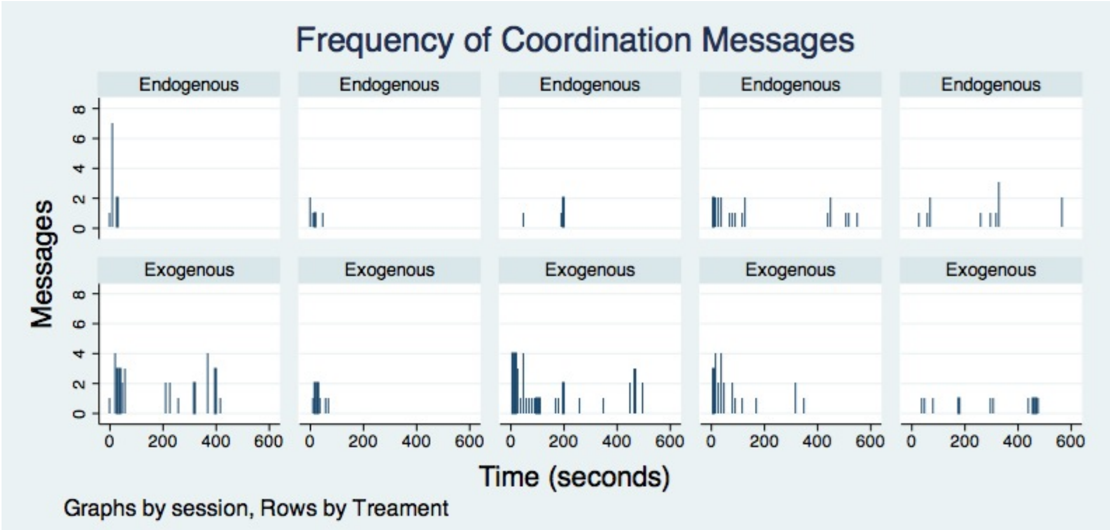

Note: "Messages" indicates the number of coordination messages sent per ten-second block.

Figure 6

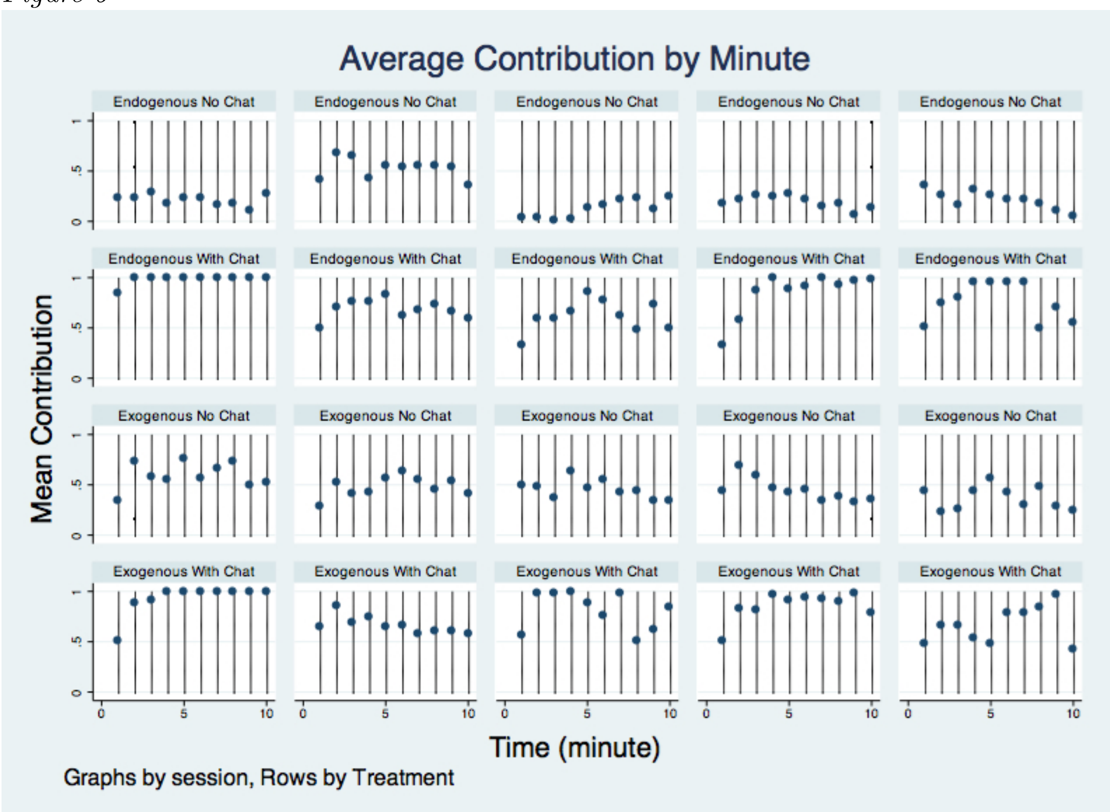

Note: "Mean Contribution" indicates the average contribution rate chosen per one-minute block.

Table 3: ***Late Contributions Regressed on First Minute Contributions***

Dependent Variable: "Following Contribution" (i.e. average contribution rate during minutes 2-10)

	Treatments			
	Endogenous		Exogenous	
	No Chat	With Chat	No Chat	With Chat
Coefficient Estimates	Following Contribution	Following Contribution	Following Contribution	Following Contribution
1st Minute Contribution	0.619**	0.538**	0.665**	0.22*
	(-4.83)	(-5.45)	(-6.49)	(-2.11)

Absolute value of t-statistics in parentheses

* significant at 5%; ** significant at 1%; standard errors are robust

Note: "1st Minute Contribution" indicates the average contribution rate during the 1st minute.

Result 5: Early behavior is sustained throughout the session. i.e. Cooperation, or lack there of, established at the beginning of a session continues through the experiment.

Support for Result 5 comes from Figure 6. In Figure 6 we see mean contribution per minute of each session. We argue that by the first or second minute of the game, the tone of the game has been set. If cooperation is not established in the first couple minutes, it does not occur later in the session. To formalize these results we regressed average contributions of participants after the one minute mark on his contribution in the first minute for each treatment. The results are shown in Table 3. From Table 3 we see that the correlation is positive is statistically significant in all treatments.

Further support is found in Figure 5. We find that in treatments with chat, sessions with high levels of cooperation also typically had high frequency of coordinating messages with the first 90 seconds. We define a coordinating message to be a message that tried to establish mutual cooperation among group members by suggesting a convergence of strategies. *e.g.* "Everyone move to the left." Regressing overall cooperation (found in Figure 2) on the total number of coordinating messages sent in the first 90 seconds yields a coefficient of 0.03 per coordination message with p-value 0.08 in the endogenous case. The same regression yields a coefficient of 0.02 with p-value 0.08 in the exogenous case. The 90 second level was chosen because that is when the first burst of coordinating messages concluded. Furthermore, when examining session 8, we see no improvement in cooperation after the second burst of coordinating messages at the 200 second mark. Similarly, subjects in the chat treatment seemed reluctant to defect from cooperation to punish free riders. Though punishment was discussed it was not committed.

One may argue that when cooperation is not established early, participants perceive that the likelihood of future cooperation is too low to continue coordinating. However, nearly all participants continue to signal a willingness to cooperate by jumping to and from a positive contribution rate. Perhaps more surprising is that subjects in session 8 failed to coordinate mid or late game cooperative equilibria even with communication. The data suggests that first impressions are a crucial part of establishing cooperation. As shown by the effectiveness of early coordinating messages, participants are very responsive to initial cooperative behavior. However, if participants fail to establish cooperation early, they may become distrusting. Therefore, all future signaling, however popular it may be, is discarded as non-credible. Players have not yet developed reputations at the beginning of the the game. Thus, participants are more willing to acknowledge good intent in words. After reputations have been established through a history of actions, messages are regarded as cheap talk.

Conversely, we argue that steady state cooperation is self-enforcing. When a cooperative steady state equilibria occurs, it exposes the free riders. The free riders will either be the only participants refusing to move towards cooperation or the only ones deviating from it. That is, because everyone else is holding a high contribution rate it becomes easier to track those who are not. This, in turn, discourages subjects concerned with self or social image from free-riding. Those concerned with self image may more easily justify free riding behavior if it is common. Those concerned with social image do not want to appear to be the only free loader in the group. Lazear *et al.* (2012) showed that subjects are often concerned with social image even when identities are kept anonymous. In their experiments they found that allowing subjects to sort out of a dictator game, by keeping the potential recipient from finding out about the game, significantly lowered giving. Avoiding the game and entering but giving nothing are equivalent with respect to payoffs and self-image. However, subjects given the option to avoid the game altogether readily took it and, thus, had significantly lower average recipient allocations.

In relation to the P2P setting, this result suggests that the ability to communicate is crucially important while such a network is in its infancy. Intuitively, the ability to communicate can allow those interested in helping to create a P2P network coordinate their contributions. However, once the network reaches maturity, communication may no longer be as vital. Rather, the previously established social norms and practices persist.

4 Conclusion

Our findings suggest that the success of P2P piracy may be replicated in a neutral setting. In our experiments we find that with sufficient coordinating mechanisms, high levels of cooperation can be established and sustained even with costly participation. One such mechanism is anonymous instant messaging. Our results not only show that messaging encourages cooperation, but also that costly participation in the public good has no effect on cooperation when messaging is allowed. With regard to policy reccomendations, our results are encouraging. It seems that it is possible to utilize real-time coordination and communication to fund public goods. However, there are several interesting possible extensions remaining. The most obvious question is whether our results can be generalized to larger populations. *e.g.* can high levels of cooperation be establish with group sizes of 800 instead of 8? There are many questions that remain unanswered in regards to how our findings can best be applied outside of the laboratory. However, our paper demonstrates some crucial factors neccessary for an endogenous public good's success. First, real-time information relay of strategies is insufficient; there must also be an coordinating mechanism that allows people to communicate commitment to contributing. Second, this mechanism must allow those involved to establish cooperation quickly. Otherwise, people may become too hesitant and pessimistic to establish cooperation later on. We suspect that the need for quick coordination is exacerbated outside of the lab, where people will have other distractions.

5 References

Ahn, T. K., Isaac, R. M., Salmon, T. C. (2009). Coming and going: Experiments on endogenous group sizes for excludable public goods. Journal of Public Economics, 93(1-2), 336351.

Charness, G., Friedman, D., Oprea, R., (2012). Continuous Time and Communication in a Public-goods Experiment, 147.

Charness, G., Yang, C. (2006). Endogenous group formation and public goods provision: Exclusion, exit, mergers, and redemption. Exit, Mergers, and Redemption.

Charness, G. (2000). Self-Serving Cheap Talk: A Test Of Aumanns Conjecture. Games and Economic Behavior, 33(2), 177194.

Cinyabuguma, M., Page, T., Putterman, L. (2005). Cooperation under the threat of expulsion in a public goods experiment. Journal of Public Economics, 89(8), 1421-1435.

Dorsey, R. (1992). The voluntary contributions mechanism with real time revisions. Public Choice, (i), 261282.

Duffy, J., Ochs, J., Vesterlund, L. (2007). Giving little by little: Dynamic voluntary contribution games. Journal of Public Economics, 91(9), 17081730.

Fershtman, C., Nitzan, S. (1991). Dynamic voluntary provision of goods. European Economic Review, 35, 10571067.

Friedman, D., Oprea, R. (2012). A continuous dilemma. The American Economic Review, 102(1), 133.

Isaac, R., Walker, J. (1988). Communication and freeriding behavior: The voluntary contribution mechanism. Economic Inquiry, XXVI(October), 585608.

Isaac, R. M., Salmon, T. C. (2008). Ndogenous roup ormation. Journal of Public Economic Theory, 10(2), 171194.

Isaac, R., Walker, J., Williams, A. (1994). Group size and the voluntary provision of public goods: experimental evidence utilizing large groups. Journal of Public Economics, 54, 136.

Lazear, E. (2012). Sorting in experiments with application to social preferences. American Economic Review: Applied Economics, 4(1), 136163.

Rege, M., Telle, K. (2004). The impact of social approval and framing on cooperation in public good situations. Journal of Public Economics, 88(7-8), 16251644.