To appear in: *Simulating Organizations: Computational Models of Institutions and Groups*. MIT Press 1996.

Fluctuating Efforts and Sustainable Cooperation

Bernardo A. Huberman and Natalie S. Glance
Dynamics of Computation Group
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304

Abstract

We show that when individuals confronted with a social dilemma contribute to the common good with an effort that fluctuates in time, they can make the group generate an average utility that actually decreases in time. This paradoxical behavior takes place in spite of the fact that typically individuals are found to be contributing at any given time. This novel phenomenon, which we verified in a number of computer experiments, is the result of an intermittent effect, whereby unlikely bursts of defection determine the average behavior of the group. Besides providing a dynamical theory for the lognormal distribution found in some sociological data, these results show that the typical behavior of individuals comprising a group can be inconsistent with its average properties.

Introduction

Social dilemmas are at the heart of the problem of generating cooperation among individuals confronted with conflicting choices. As such they have been studied for a number of years within the context of rational choice theory. As shown originally by Olson (1965), individuals having to choose between producing a public good and free riding on the effort of others face a dilemma when the benefit to be accrued by overall cooperation offsets individual costs, since rationality on the part of each individual leads to failure in achieving the common good (Hardin, 1968, Schelling, 1978, Hardin 1982, Glance and Huberman 1994.) The resolution of such dilemmas underlie successful attempts at ongoing collective action, such as the functioning of large organizations (Bendor, 1987, Glance, 1993), the mobilization of political movements (Taylor, 1976, Oliver, 1988) and the adoption of new technologies.

Beliefs and expectations are at the core of human choices and preferences. They arise from the intentional nature of people and reflect the way the future as well as the past enter into decisions that are made at present. Thus, individuals acting within the context of a larger group may take into account the effect of their actions both on their personal welfare and on the welfare of the group. In other words, individuals form their own models of how the group dynamics works based on some set of beliefs that color their preferences. Within this context, we have recently shown that for a broad set of beliefs and group characteristics cooperation can appear spontaneously in non-cooperative and diverse groups after long periods of time if delays in information are not important (Huberman, 1996). Moreover, cooperation or defection appear suddenly and unexpectedly, with diversity acting as an additional source of uncertainty, thus shortening the time to an outbreak (Huberman, 1993).

Traditionally, studies of social dilemmas have made assumptions that not always reflect the behavior of realistic groups of people, such as the existence of a utility function that is the same for all individuals, or the constancy of the benefits and costs over long periods of time. Concerning the latter, it is well known that since cooperation is based on a person's effort, one never implements precisely the amount of cooperation that one intends to. There are always a number of reasons for this to be the case. An individual may be late returning from a meeting, and so can give less time for a cooperative endeavor than he had planned. Or, in situations where the production of the collective good demands some physical effort, it is reasonable to expect that individual contributions will vary over time. The same applies to cases where intellectual effort is required, such as clerical work in an organization. In those circumstances, it can be safely assumed that the benefit from cooperation will fluctuate around a mean value, with a dispersion determined by the variance of the appropriate distribution.

In this paper we study the effects of such fluctuating benefits on the possibility of obtaining sustained cooperation from a group of individuals confronted with a social dilemma. In spite of the fact that this is the optimal case for either small groups or when individuals have long outlooks even in equilibrium some members of a cooperating group will occassionally choose to defect due to misinformation as to what the other members are doing(Glance, 1993). If in addition individual efforts also fluctuate, the benefit obtained by members of the group during those switchings could either encourage or discourage further switchings. Nevertheless, since these fluctuations lower the utility of the group, one expects that after a short period of time the group will revert to spontaneous cooperation.

If variations in individual output are normally distributed around an average value, fluctuations away from the fixed point should lead to a small and random modulation of the group's utility in time. Moreover, if the fluctuations have a vanishing time average, the average value of the utility to the group might be expected to remain constant over time. A careful study of the resulting dynamics however, shows that this is not the case. Instead, variability in effort has a dramatic effect on the overall behavior of the system. We will show that a cooperating group of individuals can generate an average utility that *decreases* in time, even though the fluctuations in effort are small and with zero mean. This striking and paradoxical behavior is the result of a social intermittency effect, whereby unlikely bursts of defection determine the average behavior of the group, in spite of the fact that individuals are typically found to be cooperating at any given time. These predictions will then be verified through a number of computer experiments.

This effect, although exhibited here within the context of social dilemmas, is likely to be encountered in many other social situations characterized by equilibria punctuated with fluctuations. In certain cases it might even be possible for a large group of defectors to generate a utility that *increases* in time. Intermittency implies that in situations where the only available data is some average utility or other productivity measure, the inference that the individuals comprising the group act in ways consistent with the average behavior might be incorrect. If one had to summarize these findings in a simple vernacular sentence it would read : *"what you see is not what you get"*.

At a different level, these results provide a dynamical theory for the lognormal distribution found in some sociological data, while showing that the typical behavior of individuals comprising a group can be inconsistent with its average properties.

Social Dilemmas

In a general social dilemma, each individual can either contribute (cooperate) to the production of a common good, or not (defect). While no individual can directly observe the effort of another, each member observes instead the collective output of the group or organization and can deduce overall group participation using knowledge of individual and group functions. In this process there is an amount of uncertainty in the relation between member's efforts and group performance. There are many causes for this uncertainty; for example a member may try but fail to contribute due to unforeseen obstacles. Another type of uncertainty might arise from the fact that individuals act with bounded rationality, occasionally making suboptimal decisions.

In a simple, but general limit, collective benefits increase linearly in the contribution of the members, at a rate b per cooperating member. Each contributing individual bears a personal cost, c. If k_i denotes whether member i is cooperating (k_i=1) or defecting (k_i=0) then the utility at time t for member i of a group of size n is given by (Bendor, 1987)

$$U_i(t) = b\frac{\hat{n}_c(t)}{n} - ck_i \qquad (1)$$

From knowledge of the functional form of the utility function, each individual can deduce the number of individuals effectively cooperating at some time t by inverting this equation and obtaning an estimate of the value of $\hat{n}_c$, or alternatively, by evaluating $\hat{f}_c(t) = \hat{n}_c(t)/n$, the fraction of individuals effectively cooperating at time t. Since the information available is not perfect, this estimate will differ from the actual number of individuals intending to cooperate.

When all members contribute successfully, each individual receives net benefits $bn/n - c = b - c$, independent of the group's size. The production of the collective good becomes a social dilemma when $b > c > b/n$. Thus, although the good of all is maximized when everybody cooperates ($b - c > 0$) the dominant strategy in a one shot interaction is to defect since additional gains from personal participation is less than the private cost($b/n - c < 0$).

This logic changes when the interaction is ongoing since future expected utility gains will join present ones in influencing the rational individual's decision to cooperate or not in the production of the collective good. In particular, individual expectations concerning the future evolution of the game can play a significant role in each member's decisions. The barest notion of expectations comes from the economic concept of horizon length. The horizon length is how far an agent looks into the future, or how long the agent expects to continue interacting with the other members of the group. Within this framework, agents believe that their present actions will affect those of others in the future. In particular, individuals expect that defection encourages defection and cooperation encourages cooperation[Quattrone and Tversky]. But the degree to which this belief holds depends on the size of the group and the present level of production.

In earlier work we stated the benefits and costs to the individual associated with the two actions of cooperation and defection, i.e. contributing or not to the social good. We also allowed for diverse beliefs and expectations about other individuals' actions in the future to influence each member perception of which action, to cooperate or defect to take (Glance, 1993). Using those preference functions, we showed that if α is the rate at which individuals reevaluate their choices, and H is the horizon length that determines the discount of future returns, there is a strategy such that an individual will cooperate if the observed fraction cooperating exceeds a value f_{crit}, which is given by

$$f_{crit} = \frac{nc - b}{H\alpha(b - c)} \qquad (2)$$

Thus, individuals engage in conditional cooperation, cooperating when the fraction perceived as cooperating is greater than a critical amount, and defecting when it is lower than the critical amount. Although the precise functional form of f_{crit} can depend on specific sets of beliefs, there is a wide range of preference functions for which a critical fraction appears as the determinant of cooperation (Huberman, 1994).

The dynamics of cooperation is determined by the equation governing the fraction of agents cooperating at a given time. This equation was derived earlier, and it is given by (Glance, 1993)

$$\frac{d\hat{f}_c(t)}{dt} = -\alpha\{\hat{f}_c(t) - \frac{1}{2}[1 + \mathrm{erf}(\frac{(<\hat{f}_c(t)> - f_{crit})}{\sigma\sqrt{2}})]\} \qquad (3)$$

where erf denotes the error function, $<\hat{f}_c> = p\hat{f}_c + (1-p)(1-\hat{f}_c)$, with p the probability that an agent intending to cooperate does so successfully, and σ determines the level of uncertainty in the system. This uncertainty is made up of two parts. One is related to the connection between intent and action that we just stated. The other refers to the variance in individual characteristics that make individuals have a diversity in their outlooks or in their benefit-cost ratios(Huberman, 1993).

The fixed points of this equation, which are obtained by setting the right hand to zero, give the critical sizes beyond which cooperation can no longer be sustained. In our earlier work we showed that there is a sharp upper limit, n*, to the size of a group which can support cooperation. Even if all members of a group larger than this size started cooperating, after a long time the whole group would switch to defection. Moreover, the outbreak of cooperation is sudden and unpredictable, taking place over a very short period of time. These predictions were observed in a number of computer experiments simulating social dilemmas.

Fluctuating Efforts

Although in the equilibrium state the number of cooperating agents does not change in time, occasional fluctuations take place whereby individuals might switch from cooperation to defection for short periods of time. These random events are due to the uncertainty that individuals have about the behavior of the others. This uncertainty can cause an individual to mistakenly perceive the amount of cooperation to be different than its actual value. Because of this misperception, an individual might act against the equilibrium condition of mutual cooperation or defection, causing the system to move away from the fixed point. The more uncertainty there is in the system, the more likely it is that there will be fluctuations away from the equilibrium state.

The situation becomes more complicated if one allows for the fact that individual efforts at contributing to the common good vary from time to time. This implies that superimposed on the occasional switchings from cooperation to defection and from defection to cooperation, there is an extra source of randomness stemming from the fact that benefits fluctuate in time, thus potentially creating a complex dynamical scenario.

When fluctuations take place, the group recovers in time according to an equation that governs the dynamics of fluctuations away from the equilibrium state. This equation can be obtained by performing a linear stability analysis of Eq. (3) around the fixed point, which we will denote by $\hat{f}_o$. For the sake of simplicity, in what follows we will consider the case where the equilibrium corresponds to overall cooperation, so that $\hat{f}_o \cong 1$.That is the optimal situation for the case of small groups or when individuals have long horizons. The opposite case, that of a large group of defectors, can be treated by similar methods. Denoting a small deviation from the fixed point by δ_c, the fraction of individuals cooperating at any given time will be given by

$$\hat{f}_c(t) = \hat{f}_o - \delta_c(t) \tag{4}$$

with $\delta_c << \hat{f}_o$. Replacing this expression in Eq. (3) and performing a linear stability analysis around the fixed point, yields an equation for the dynamics of fluctuations around the cooperative equilibrium. As shown in Appendix A, in the case of large uncertanties due to diversity one obtains

$$\frac{d\delta_c(t)}{dt} = -\alpha\delta_c(t)[A' - \frac{2}{\sigma^3\sqrt{2\pi}} f_{crit}] \tag{5}$$

with $A' = 1 + \frac{1}{\sigma\sqrt{2\pi}}(\frac{1}{2\sigma^2} - 1)$.

If the parameter values do not change with time, the expression inside the bracket of Eq. (5) becomes a constant and the solutions of this equation correspond to exponential relaxation of any initial fluctuation to the state $\delta_c = 0$, i.e. overall cooperation. The characteristic relaxation time is given by the inverse of this constant bracket times the evaluation rate, α. The situation changes drastically however, when the cooperative output that an individual can produce fluctuates in time, making the utilities change accordingly. In what follows we will assume that the cooperative output is given by

$$b(t) = \bar{b} + \beta(t) \tag{6}$$

where $\bar{b}$ is the time independent benefit that results from cooperating, and the fluctuating component, $\beta(t)$, is a Gaussian random variable such that

$$\begin{gathered} \langle \beta(t) \rangle = 0 \\ \langle \beta(t)\beta(t') \rangle = 2\sigma'\,\delta(t-t') \end{gathered} \tag{7}$$

where σ' denotes the variance in effort. Notice that the Gaussian character of the fluctuations implies that for all practical purposes the probability of observing an individual contributing a benefit above or below $\bar{b}$ by an amount exceeding three standard deviations is vanishingly small.

As shown in Appendix B, fluctuations in benefits translate into a time dependent critical fraction that is given by

$$f_{crit}(t) = f_{crit} - \frac{\beta(t)}{H\alpha(b-c)} \tag{8}$$

with f_{crit} given by Eq. (2), but with b replaced with $\bar{b}$. Substituting this expression in Eq. (5), we obtain the equation governing the dynamics of fluctuations away from the cooperative fixed point. It is given by

$$\frac{d\delta_c(t)}{dt} = -A\delta_c(t) + B(t)\delta_c(t) \qquad (9)$$

where $A = \alpha A' - \alpha f_{crit} \big/ \sigma^3 \sqrt{2\pi}$, and $B(t) = -\beta(t) / \sqrt{2\pi} H\sigma^3 (b-c)$. It also follows that $B(t)$ is a Gaussian random variable with zero mean and with autocorrelation equal to $\langle B(t)B(t')\rangle = \bar{\sigma}\delta(t-t')$, with $\bar{\sigma} = \sigma' / (\sqrt{2\pi} H\sigma^3 (b-c))$.

With these properties of the random function $B(t)$, Eq. (9) is a stochastic differential equation. In what follows we will obtain an exact solution of this equation and calculate the behavior of all the moments of δ_c.

Bursts of Defection

In order to study the role of fluctuations around the overall cooperative state, we now solve the stochastic differential equation given by Eq. (9) by standard methods (Arnold, 1992). Since the fluctuations in $B(t)$ are Gaussian, we can use the Stratonovich formalism (Stratonovich, 1967) to find the time evolution of δ_c. It is given by

$$\delta_c(t) = \delta_c(0)e^{-At+w_t} \qquad (10)$$

where $\delta_c(0)$ is the initial value of the fluctuation and w_t is a Wiener process such that $B(t) = \frac{dw_t}{dt}$, and with the property that

$$\langle w_t \rangle = 0$$
$$\langle w_t^{\,2} \rangle = \overline{\sigma} t \qquad (11)$$

As can be seen from Eq. (10), typical fluctuations away from the cooperative state relax exponentially fast to the fixed point, a result to be expected in the case of small groups and large horizons.

A paradox appears however, when computing the moments of $\delta_c(t)$, which are given by

$$\left\langle [\delta_c(t)]^m \right\rangle = [\delta_c(o)]^m \left\langle e^{-(At+w_t)m} \right\rangle \qquad (12)$$

where m is the mth moment and $\delta_c(0)$ is the initial configuration of the fluctuating system. Using the fact that $\left\langle e^{mw_t} \right\rangle = e^{m^2 \sigma t}$ this equation becomes

$$\left\langle [\delta_c(t)]^m \right\rangle = [\delta_c(o)]^m e^{-Amt+m^2\overline{\sigma}t} \qquad (13)$$

From this expression, it is clear that the higher moments grow in time with increasing order, at a rate given by $m(m\overline{\sigma} - A)$. This happens while typical realizations of the fluctuations relax exponentially fast, as indicated by Eq. (10). In particular, the average fraction of defecting agents, which is given by the first moment (m=1) of $\delta_c(t)$, *increases* in time when the fluctuations in individual efforts exceed the value of A, which amounts to the condition

$$\overline{\sigma} > \alpha\sigma^2 H(b-c)(\sigma\sqrt{2\pi}-1) \qquad (14)$$

This behavior takes place in spite of the fact that typical fluctuations decay to zero, as shown by Eq. (10). Even more striking, the average utility to an individual, which is given by $\langle U_i(t)\rangle = \langle b\hat{f}_c(t)\rangle - ck_i$ *decreases* in time as

$$\langle U_i(t)\rangle = \overline{b}[1-\delta_c(0)e^{(\overline{\sigma}-A)t}] - ck_i \qquad (15)$$

Therefore, a measurement of the utility over a given time interval would lead to the wrong inference that defection is the prevalent strategy, for most invividuals are cooperating, as follows from Eq. (10). Conversely, an increasing average utility might not necessarily signal increasing levels of cooperation in a large group.

In order to understand this paradox, it is useful to remember that the higher moments of $\delta_c(t)$ are related to the probability of very unlikely events. The fact that the moments grow in time indicates that the behavior of the system is dominated by occasional bursts in which the number of individuals suddenly found to be defecting is very large, so large that the probability of their occurence would be practically zero if the fluctuations were Gaussian in nature. These intense peaks of defection happen in spite of the fact that fluctuations in effort are normally distributed around a mean value $\overline{b}$.

If these intermittent burst are non-Gaussian, we need to compute the quasistationary distribution of fluctuations that results from their evolution over a long time period. In order to so, divide both sides of Eq. (9) by δ_c and integrate over a time interval T, to obtain, up to a constant,

$$\ln\frac{\delta_c(T)}{\delta_c(0)} = \sum_0^T B_n \qquad (16)$$

where the integral of B(t) over time has been replaced by a sum over discrete time steps, at each of which the random quantity *B(t)* acquires the value B_n.

Since the values of B_n are normally distributed random quantities with zero mean and variance $\overline{\sigma}^2$, the sum of the right hand side can be evaluated by invoking the Central Limit Theorem. Therefore, the logarithm of δ_c will be normally distributed with zero mean and variance $\overline{\sigma}^2$. In other words, δ_c is log-normally distributed according to

$$\text{Prob}[\delta_c] = \frac{1}{\delta_c \overline{\sigma}\sqrt{2\pi}} e^{-\left(\frac{(\log \delta_c)^2}{2\overline{\sigma}^2}\right)} \qquad (17)$$

.

With the mean value of δ_c, μ, given by $\mu = e^{\overline{\sigma}^2}$, and its variance by $\mu^2(e^{\overline{\sigma}^2} - 1)$.

This probabiity density function, with mean 1.0 and standard deviation 1.0, is shown in the figure, along with the Normal one with the same mean and standard deviation.

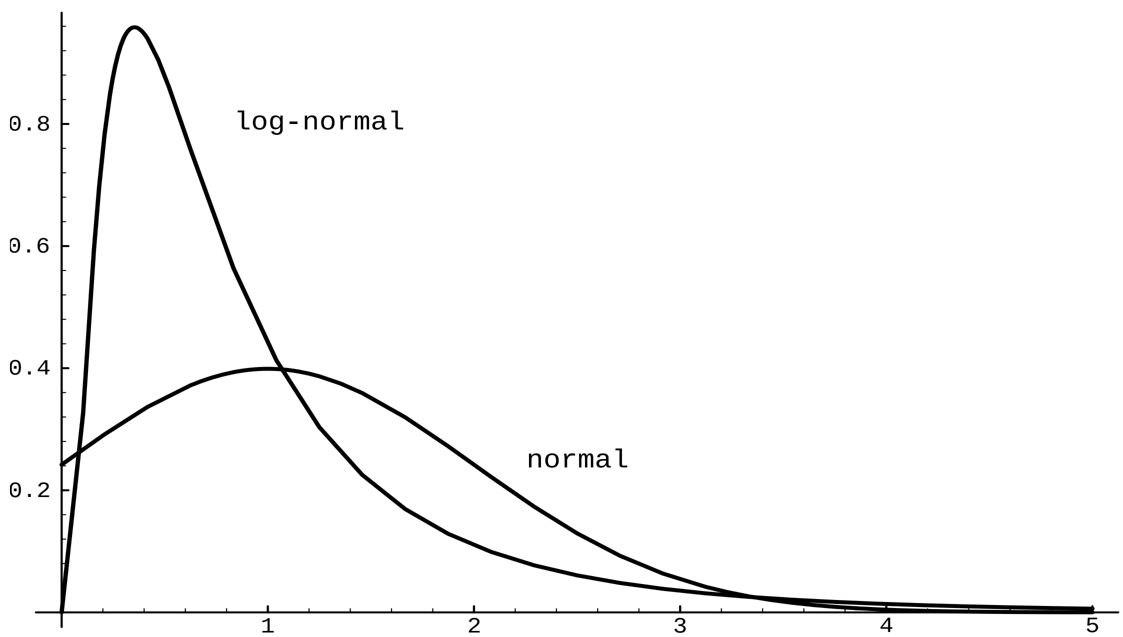

As can be seen, the log-normal distribution differs from the Gaussian one in two important aspects. The first one has to do with the fact that, unlike the normal distribution, the log-normal is quite asymmetric. Thus, typical realizations of the system do not coincide with its average behavior. Second, the tail of the log-normal distribution extends much further out than that of a normal distribution with comparable mean and variance. This can seen by the fact that for the parameters used in the Figure, the ratio of the log-normal to the normal distribution at six standard deviations is 33000, implying a finite probability for events which, when described by a normal distribution would have vanishing probability of occurring.

If fluctuations around the fixed points of the cooperative solution to social dilemmas are distributed in a log-normal fashion, they should be observable in computer simulations. But since in this case the average is dominated by rare intermittent events, care has to be exercised in the analysis of such simulations. If the simulation runs for a relatively short time, it would only sample a small fraction of the total number of possible states of the system. Thus, one would then detect the typical, or most probable, value of the observable. It is only when simulations run for long enough times so as to sample a finite fraction of all the states of the system that the measurements will converge to the true average of the observable.

Computer Experiments

We now describe a number of computer experiments that were performed to test the theoretical predictions made above.

The theory presented in this paper predicts that the distribution fo the number of individuals cooperating over time is lognormal in certain domains. In order to test this prediction, we ran experiments on the evolution of cooperation in a number of different groups characterized by different sizes, horizon lengths, uncertainty levels, and magnitudues of time fluctuations in the efforts. The algorithm that we used is described in Appendix C.

The experiments consisted of event-driven Monte Carlo simulations that were conducted in asynchronous fashion, with each agent reevaluating its choice to cooperate or defect at a rate α. The agent decides whether to cooperate or defect based on the fraction, f_{crit}, it observes cooperating and the choice criterion given by Equation (2). That is, when the fraction observed cooperating is greater than f_{crit}, the agent will choose to cooperate; otherwise it will defect. As we pointed out above, the fraction observed cooperating differs from the actual fraction intending to cooperate in a way that depends on the amount of uncertainty in the system. Because the benefit of cooperation fluctuates stochastically in time, f_{crit} does so as well. As a result, the condition for cooperation becomes more or less severe each time an agent reevaluates its decision to cooperate or defect.

For each experiment we recorded the number of individuals cooperating at discrete intervals. This data was then binned to yield the distribution of the number of cooperating individuals over the course of the recorded run. We then used the chi-squared probability distribution to test how well the experimental data could be described by either a lognormal distribution or a normal distribution.

For a number of experimental conditions, we found that a lognormal distribution could account for the observed data much better than a normal distribution could. Figure 2 shows the fit to a lognormal of the distribution of the number of individuals defecting for an experiment in which the group size was n =100, the benefit for cooperating was b=2.5, the standard deviation in b due to fluctuations=0.25, the cost of cooperation c=1, the horizon length H=69, and the probability of error p=0.9. In this case, the chi-square test yields the probability 99./4% that the predicted distribution explains the experimental data, while a normal fit to the data yields a chi-square statistic of only 4.6%.

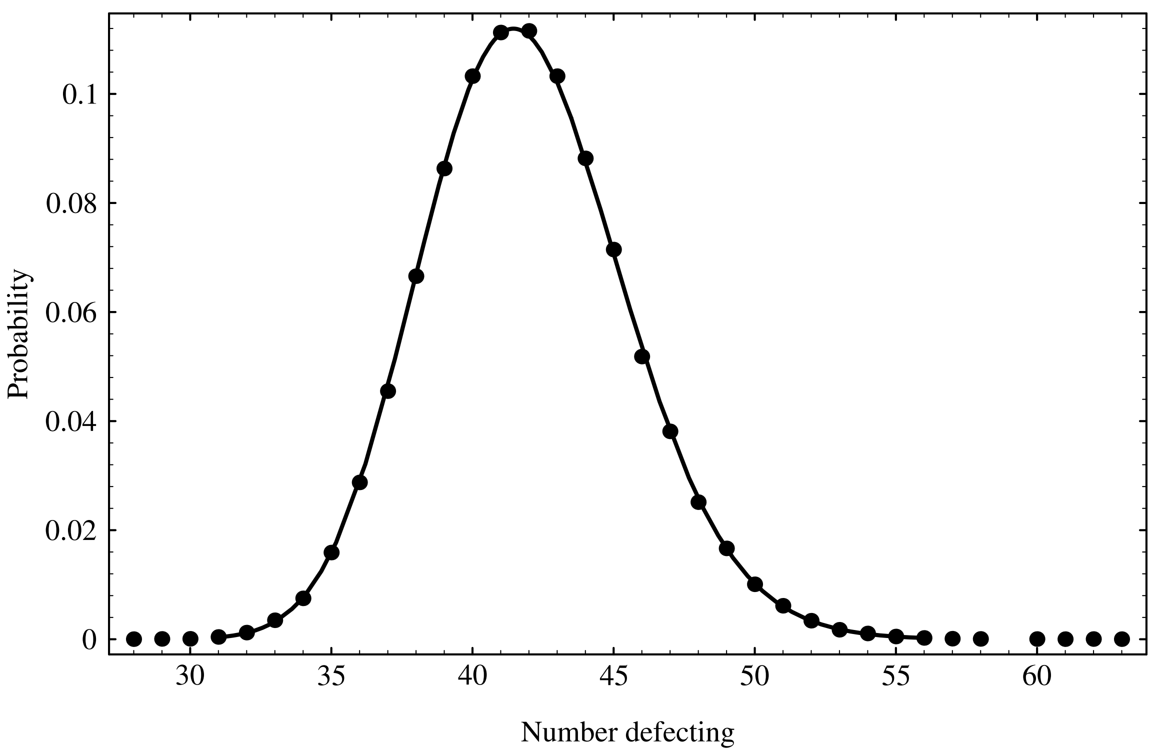

Probability
0.1
0.08
0.06
0.04
0.02
0
30
35
40
45
50
55
60
Number defecting

While it is impossible to statistically prove a theory, these results show that our predictions have a significant higher probability of being correct than a theory which would yield a normal distribution of fluctuations in the number of individuals cooperting.

Discussion

In this paper we have shown how fluctuations in the effort to produce a public good can lead to a new social phenomenon characterized by rare and intense bursts of defection superimposed on a background of overall cooperation. This phenomenon is stochastic in nature. It leads to a paradoxical behavior of the average utility accrued by any individual, for it *decreases* in time in spite of the fact that typically all inviduals are found to be contributing the same average average amount of effort. Conversely, there can be cases where an equilibrium group of defectors can generate a utility that *increases* in time, although the conditions are more restrictive than in the case of cooperation.

This paradox was observed in computer simulations of social dilemmas and should also appear in real social settings corresponding to equilibria punctuated by fluctuations. If costs or benefits fluctuate, then one would observe the fact that typical, most probable, behavior is not the same as the average. This is also reflected in the nature of the statistical distribution of the data. Rather than obtaining the familiar bell-shaped curve corresponding to a normal distribution, one would observe an asymmetric distribution with a long tail. We emphasize that the detailed form of the critical fraction used in Eq. (2) is not important for the existence of this effect. If the utility function has the general form given by Eq. (1), then fluctuations in benefits will lead to intermittent bursts in cooperation and defection.

For the purpose of verifying these predictions, computer simulations have the advantage over social experiments that the variables and parameters can be precisely specified. Although computer simulations are no substitute for real social dynamics, their relative simplicity and clear outcomes can be used as a sign that this phenomenon can appear in real social settings, in particular those situations where the average behavior of quantities reflecting collective action seems at odds with what typical member of a group are doing.

An interesting application of these results is to provide a possible explanation for the so-called paradox of group size in political movements (Oliver, 1988). It has been documented that in certain cases (Marwell, 19709) the larger the group the more likely it is to engage in collective action, a result that seems at odds with the common resolution of social dilemmas. This apparent paradox was explained by Oliver and Marwell (Oliver, 1988) by invoking the simple statistical fact that if what is needed to trigger a political movement is a minority of highly motivated individuals, a large group will be more likely to have such individuals than a smaller one.

This paper presents a dynamical alternative to that view. If large temporal variations in effort exist in a social group, rare bursts of extreme cooperation could lead to an average utility that increases in time, so that individuals confronted with a social dilemma might construe a level of cooperation which is at odds with what typically everybody is doing. A test of this hypothesis could be performed by measuring the distribution of levels of effort in large political groups. Whereas the theory of Oliver and Marwell would predict that it is normally distributed around some mean value, our theory would lead to an asymmetric, log-normal type, distribution. If this were the case, it would provide a theoretical underpinng to the folk observation that in voluntary organizations ten percent of the people do fifty percent of the work.

On a normative vein, these results suggest that in order to secure cooperation within a group or organization, the requirements of long, diverse outlooks and small groups might not be enough. Since fluctuations in the effort of individuals are unavoidable, the choice of individuals with small variance in their outputs will be more successful at generating cooperative behavior than having a group of the same size made up of individuals with large temporal variations in effort. If that were not possible, the solution would then have to rely on either making controls and monitoring more efficient, or to organize the production of the common good in such a way that the overall utility due to cooperation increases faster than linear with the number of individuals in the organization. This can be the case when cooperation is used in the solution of complex problems.[Clearwater, 1993, Huberman, 1995].

Last, but not least, this paper produces a dynamical derivation of the lognormal distribution, which has been used to fit data ranging from distributions of income in diverse populations to the size distributions of cities. Even for situations where one has a-priori expectations of normal distributions of data, these results also suggest that outliers that are routinely rejected might actually signal a deeper regularity, that of an asymmetric distribution with a long tail underlying the phenomenon being studied. In such cases, not only is it warranted to include those extreme events in analysis of the data, but also to arrive at average measures that depart from typical behavior, one of the signatures of intermittent bursts of activity.

Acknowledgments

This work profited from useful discussions with Jon Bendor and Mette Huberman. Part of the theoretical work was done while one of the authors, BAH, was a visiting professor at CONNECT, which is part of the Niels Bohr Institute, University of Copenhagen. The warm hospitality provided by Dr. B. Lautrup and his staff is gratefully acknowledged.

Appendix A

We now present the linear stability analysis of Eq. (3) around its fixed points.

Equation (3) can be written as

$$\frac{d\hat{f}_c(t)}{dt} = F(\hat{f}_c) \qquad \text{(A-1)}$$

with
$$F(\hat{f}_c) \equiv -\alpha[\hat{f}_c - \rho(\hat{f}_c)] \qquad \text{(A-2)}$$

and
$$\rho(\hat{f}_c) = \frac{1}{2}\{1 + \mathrm{erf}[\frac{(<\hat{f}_c> - f_{crit})}{\sigma\sqrt{2}}]\} \qquad \text{(A-3)}$$

with $<\hat{f}_c> = p\hat{f}_c + (1-p)(1-\hat{f}_c)$, and $0 \le p \le 1$. The fixed points of this equation are determined by $F(\hat{f}_{eq)}) = 0$ and their stability can be studied by performing a linear stability analysis of Eq. (A-1) around them. For the sake of simplicity we will study the fluctuations around the cooperative fixed point, $\hat{f}_o$, which obtains in the limit of small groups or for the case of long horizons. Writing

$$\hat{f}_c(t) = \hat{f}_o - \delta_c(t) \qquad \text{(A-4)}$$

where $\delta_c << \hat{f}_o$ is the deviating fraction of individuals away from the fixed point, replacing this expression in Eq. (A-1) and performing a Taylor series expansion of $F(\hat{f}_c)$ around its fixed point we obtain

$$\frac{d\delta_c(t)}{dt} = \delta_c(t)F'(\hat{f}_o) \qquad \text{(A-5)}$$

where $F'(\hat{f}_o)$ denotes the derivative of $F(\hat{f}_c)$ with respect to $\hat{f}_c$ evaluated at $\hat{f}_o$. Using Eqs. A-2 and A-3, and assuming $p \approx 1$ this equation can be written as

$$\frac{d\delta_c(t)}{dt} = -\alpha\delta_c(t)\{1 - \frac{2}{\sigma\sqrt{2\pi}}\exp[-\frac{(\hat{f}_o - f_{crit})^2}{2\sigma^2}]\} \qquad \text{(A-6)}$$

where, since $p \approx 1$, the uncertainty is mostly determined from diversity of outlooks among the members of the group.

For the case of a cooperative equilibrium with long horizons, $f_o \approx 1$, and f_{crit} is small. Expanding the exponential in a power series and assuming a large uncertainty gives, to lowest order in f_{crit},

$$\frac{d\delta_c(t)}{dt} = -\alpha\delta_c(t)[A' - \frac{f_{crit}}{\sigma^3\sqrt{2\pi}}] \qquad \text{(A-7)}$$

where $A' \cong 1 + \frac{1}{\sigma\sqrt{2\pi}}(\frac{1}{2\sigma^2} - 1)$.

Appendix B

We now establish how fluctuations in effort affect the time dependence of f_{crit}. As shown in Eq. (2) the critical fraction is given by

$$f_{crit} = \frac{nc - b}{H\alpha(b - c)} \qquad \text{(B-1)}$$

In the case of time dependent benefits, $b(t) = \bar{b} + \beta(t)$, so that Eq. (B-1) becomes

$$f_{crit} = \frac{nc - \bar{b} - \beta(t)}{H\alpha(\bar{b} - c)(1 + \frac{\beta(t)}{\bar{b} - c})} \qquad \text{(B-2)}$$

For small enough fluctuations such that $\beta << (\bar{b} - c)$ this can be written as

$$f_{crit} \cong \frac{nc - \bar{b} - \beta(t)}{H\alpha(\bar{b} - c)} \qquad \text{(B-3)}$$

Or equivalently

$$f_{crit}(t) = f_{crit} - \frac{\beta(t)}{H\alpha(\bar{b} - c)} \qquad \text{(B-4)}$$

with f_{crit} given by Eq. (B-1) but with b replaced with $\bar{b}$.

Appendix C

The algorithm used in the simulations consisted of the following steps.

- Initialize individual actions.

- While *current_time t* < f*inal_time*:

  Step foraward Δt in time (an exponential random deviate) to next agent reevaluation: $\Delta t = \ln(\textit{random number}) / \alpha n$, - this simulates asynchronous reevaluation.

  current_time t $=$ f*inal_time* $+ \; \Delta t$

  Pick an agent at random for reevaluation.

  Evaluate time fluctuation in benefits: $b(t) = b + \Delta b,$ with $\Delta b = \sigma * \textit{gaussian deviate}$.

  Evaluate $f_{crit}(t) \equiv \dfrac{1}{H\alpha} \dfrac{nc - b(t)}{b(t) - c}$.

  Evaluate number observed cooperating by individual at *current_time* as a function of the number intending to cooperate and the uncertaintyh parameter, p:

  $$\hat{n}_c = \textit{binomial deviate}(n_c, p) + \textit{binomial deviate}(1 - n_c, 1 - p)$$

  If the fraction observed cooperating, $\hat{f}_c > f_{crit}(t)$, then individual cooperates; if $\hat{f}_c < f_{crit}(t)$ then individual defects; if $\hat{f}_c = f_{crit}(t)$ the agent cooperates with probability 0.5 and defects with probability 0.5.

Figure Captions

Figure 1. Probability densities for both the Log-normal distribution and the Normal Distribution. They were computed for the case of mean 1.0 and standard deviation 1.0.

Figure 2. Lognormal fit to the distribution of the number of individuals defecting, sampled 5000 times at regular intervals. The parameters are: group size n=100, benefit for cooperation b=2.5, standard deviation in b due to fluctuations=0.25, cost of cooperation c=1, horizon length H=69, probability of error p=0.9. The points represent the experimental data, while the curve is a lognormal fit to the data. The chi-square test, used for testing the difference between the binned data and a continuous probability function, yields the probability 99.4% that the predicted distribution can explain the experimental data.

References

Arnold, Ludwig. 1992. *Stochastic Differential Equations: Theory and Applications*. Krieger Publications.

Bendor, Jonathan and Mookherjee, Dilip. 1987. Institutional structure and the logic of ongoing collective action. *American Political Science Review*, 81(1):129-154.

Glance, Natalie and Huberman, Bernardo A. 1994. The Dynamics of Social Dilemmas. *Scientific American,* 76-81, March.

Glance, Natalie and Huberman, Bernardo A. 1993. Organizational Fluidity and Sustainable Cooperation, in *Computational Organization Theory*, Eds. K. Carley and M. Prietula. Lawrence Elrbaum Associates.

Glance, Natalie and Huberman, Bernardo A. 1993. The outbreak of cooperation. *Journal of Mathematical Sociology*, 17(4):281-302, 1993.

Granovetter, Mark. 1978. Threshold models of collective behavior. *American Journal of Sociology* 83(6):1420-1443.

Hardin, Garret. 1968. The tragedy of the commons. *Science*, 162:1243-1248,.

Hardin, Russell.1982. *Collective Action*. Johns Hopkins University Press.

Huberman, Bernardo A. and Glance, Natalie. 1993. Diversity and Collective Action. In: H. Haken and A. Mikhailov, eds., *Interdisciplinary Approaches to Complex Nonlinear Phenomena*. pp.44-64 Springer.

Huberman, Bernardo A. and Glance, Natalie. 1995. Beliefs and Cooperation. In P. Danielson, editor. Modelling Rationality, Morality and Evolution. Oxford University Press. 1996.

Marwell, Gerald. 1970. Comment on Scott and El-Assal. *American Sociological*

Review 35:916.

Oliver, Pamela E. and Marwell, Gerald. 1988 . The paradox of group size in collective action: A theory of the critical mass. ii. *American Sociological Review*, 53:1-8.

Olson, Mancur. 1965.*The Logic of Collective Action*. Harvard University Press.

Schelling, Thomas C. 1978. Micromotives and Macrobehavior. W. W. Norton and Company.

Schmidt, Steffen, James Scott, Carl Lande and Laura Gausti. 1977. *Friends, Followers and Factions*. Berkeley. University of California.

Stratonovich, R. L. 1967. *Topics in the Theory of Random Noise*. Gordon and Breach.

Taylor, Michael. 1976. Anarchy and Cooperation, John Wiley and Sons.