

Human Motivation and Social Cooperation: Experimental and Analytical Foundations

Ernst Fehr¹ and Herbert Gintis²

¹Institute for Empirical Research in Economics, University of Zurich, Blumlisalpstrasse 10, CH – 8006 Zurich, Switzerland; email: efehr@iew.unizh.ch

²Santa Fe Institute and Central European University, 15 Forbes Avenue, Northampton, MA 01060, USA; email: hgintis@comcast.net

Key Words: Social Order, Social Exchange, Cooperation, Punishment, Strong Reciprocity.

Abstract: Since Durkheim, sociological explanations of social cooperation emphasize the internalization of values that induce norm compliance. Since Adam Smith, economic explanations of social cooperation emphasize incentives that induce selfish individuals to cooperate. Here we develop a general approach – the Beliefs, Preferences, and Constraints approach – showing that each of the above models is a special case. Our approach is based on evidence indicating that pure Homo Sociologicus and pure Homo Economicus views are wrong. We show that self-regarding and norm-regarding actors coexist and that the available action opportunities determine which of these actor types dominates the aggregate level of social cooperation. Our approach contributes to the solution of long-standing problems, including the problems of social order and collective action, the determinants and consequences of social exchanges, the micro-foundations of emergent aggregate patterns of social interactions, and the measurement of the impact of cultural and economic practices on individuals' social goals.

INTRODUCTION

THE PROBLEM OF SOCIAL ORDER

Social order as a public good

SOCIAL COOPERATION AND STRONG RECIPROCITY

Strong reciprocity in public goods provision - evidence

SOCIAL EXCHANGE AND STRONG RECIPROCITY

Social exchanges in the laboratory

The consequences of social exchanges

MEASURING THE IMPACT OF SOCIETY ON INDIVIDUALS' PREFERENCES

CONCLUDING REMARKS ON THE BELIEFS, PREFERENCES AND CONSTRAINTS APPROACH

INTRODUCTION

Over the last 20 years, prominent sociologists such as Abbott (1998), Boudon (2003), and Coleman (1986) have expressed serious reservations about the explanatory power of sociological theory and research. Boudon (2003), for example, writes: “Skepticism toward sociology has grown over the years. ... To many sociologists the state of the discipline is unsatisfactory”. We believe that one important reason for this dissatisfaction is that despite important analytical contributions such as Coleman (1990), Hedstrom (2005), and the work summarized in Macy and Willer (2002) and Hedstrom and Swedberg (1998), sociological theory has not developed a coherent, broadly accepted framework that facilitates cumulative scientific progress and explains the emergent aggregate patterns of social behavior in terms of individuals’ preferences, their beliefs, and the social and economic constraints they face. Nor has sociological research developed a parsimonious, empirically grounded view of the basic motivational driving forces of human behavior, which may be due to the limited role that controlled experiments played in the development of the discipline.

Decades ago, sociologists criticized the “oversocialized conception of man” (Wrong 1961) that played a prominent role in the work of Durkheim (1938) and Parsons (1937). They rightly questioned Homo Sociologicus, a creature who follows prevailing social norms without regard to self-interest. But they did not develop an alternative, empirically grounded, and widely accepted conception of the basic motivational driving forces of humans. This contrasts sharply with the approach taken by mainstream economics which rests on the notion of Homo Economicus, a creature who is rational and purely self-regarding. However, the Homo Economicus approach is also erroneous, as the assumption that humans are exclusively self-regarding has been decisively rejected by the evidence (Camerer 2003; Fehr and Fischbacher 2003; Gintis et al. 2003).¹ Thus, while the lack of a model of human social behavior leaves sociology *without* an anchor, mainstream economics is hitched to the *wrong* anchor, i.e., adheres to a biased view of human nature.

There is a rich tradition of experimental research in sociology (e.g., Bonacich and Light 1978; Cook and Hegtvædt 1983; Kollock 1998; Yamagishi, Jin and Kyonari 1999); controlled experiments played, in particular, an important role at the boundaries between sociology and social psychology (e.g., Marwell and Ames 1979; Yamagishi 1986; Lawler and Yoon 1993;

¹ In principle, economic methods such as optimization under constraints are strictly neutral with regard to the nature of people’s preferences. These methods certainly allow for the formalization of a richer set of motives, but in the past economists assumed that preferences are purely selfish in almost all applications.

Molm 1997; Stolte, Fine and Cook 2001; Buskens and Raub 2002; Simpson 2004; Horne 2004; Diekmann 2004). However, the experimental method is generally not prominently featured in sociological research, and experimental results are rarely published in leading journals like the *American Journal of Sociology* and the *American Sociological Review*. Most empirical work in sociology is based on non-experimental methods; experimental work by sociologists has rarely entered the mainstream of the discipline. This is unfortunate because the experimental method permits cumulative empirical and theoretical progress through the rigorous testing of alternative theories, the establishment of causal relationships between key variables, and the replicability of results. In contrast, non-experimental field data rarely permit a clean discrimination between theories; the inference of causality is typically not possible from this data, and replicability often proves difficult or impossible. The rigor and cumulative knowledge generated in the natural sciences is undoubtedly based upon their ability to conduct controlled experiments. Of course, the experimental method faces particular challenges in the social sciences, but the development of experimentation in psychology, sociology, and economics over the last two decades has clearly demonstrated the huge scientific benefits that accrue from laboratory experiments. Recent years also have witnessed the proliferation of field experiments that enable researchers to observe how experimentally controlled changes in key variables affect individual behavior in a natural social environment (Carpenter, Harrison and List 2005).

During the last two decades, there has been a surge in experimental research within economics. Much of this research has focused on topics of considerable importance for sociology, but it is our impression that this research is not widely known among sociologists. The purpose of this article is to highlight some of the research in this tradition. In addition, we will describe our own approach - the Beliefs, Preferences and Constraints model - which is characterized by a combination of micro-level experimental research and theoretical models informed by the experimental results. We will show, in particular, that one of the classic problems addressed by sociology – the problem of social order – can be better understood through this approach. In addition, we document, based on important work of sociologists such as Homans (1958), Blau (1965), Emerson (1972) and Cook (1987), how our approach can contribute to a better understanding of social exchanges. We show experimentally and in theory how social exchanges are enforced, how they shape trading relationships and competitive markets, how fairness norms shape them, how they limit the impact of supply and demand forces on price formation, and how they affect the distribution of the gains from exchange.

In the final section of our paper, we point out that experimental tools are critical for answering one of sociology's deepest questions: To what extent does society shape individuals' preferences, and how does it do so? Perhaps the foremost feature distinguishing sociology from the other social sciences consists in sociologists' emphasis on the role of socialization in general, and of the internalization of norms in particular, in constituting human behavior. However, we cannot acquire reliable knowledge about the effects of socialization without carefully controlled experiments. Questionnaires, surveys, and observation of behaviors in the field do not enable the researcher to make reliable inferences concerning motivation. In particular, one can attribute a self-regarding motive to virtually any real-life behavior, however deeply it appears driven by altruistic concerns, because observed prosocial acts can almost invariably be attributed to the selfish motive of acquiring a good reputation. By contrast, in laboratory experiments we can rule out the kind of repeated interactions on which reputational incentives are based.

Our approach rests on a large body of experimental evidence that refutes an important assumption of mainstream economics, namely, that all or most people are *exclusively* self-regarding (Camerer 2003; Fehr and Fischbacher 2003; Gintis et al. 2003). We define a *self-regarding* actor to be an agent in a social situation who maximizes his own payoff. A self-regarding actor thus cares about the choices and payoffs to other individuals only insofar as these influence his own payoff. The experimental evidence not only rejects the selfishness assumption routinely made in economics, but also suggests an alternative view about a basic predisposition of humans: strong reciprocity. Strong reciprocity is the behavioral predisposition to cooperate conditionally on others' cooperation and to punish violations of cooperative norms even at a net cost to the punisher. We will show that a substantial proportion of experimental subjects typically exhibits strongly reciprocal behavior. In addition, the evidence and our theoretical approach suggest that the interaction between strongly reciprocal and self-regarding actors drives the emergent patterns of social cooperation and social exchange. Finally, theory and evidence also indicate that the social structure of interaction plays a decisive role in shaping the emergent aggregate patterns of behavior by affecting how strongly reciprocal and self-regarding individuals interact.

Throughout the paper, the BPC approach in combination with motivational assumptions inferred from laboratory experiments guides our arguments. The BPC model (see Gintis, in press) may be considered as a variant of a game-theoretic approach. It is based on the assumption that people have consistent (transitive) preferences and beliefs about other people's behavior and about the consequences of their choices. Behavior in this approach can

be represented as choices that best satisfy people's preferences, given their beliefs and the constraints they face. In contrast to mainstream economics, our motivational assumptions are firmly grounded in empirical evidence. This evidence suggests the existence of a heterogeneous population of strong reciprocators and self-regarding individuals.

Subjects face economic incentives in all experiments discussed in this paper: they can earn money – sometimes in substantial amounts. It has been shown, for example, that strongly reciprocal behavior persists even when subjects can earn up to three months' income in a two hour experiment (Cameron 1999). The experiments also typically implement anonymous interactions between the subjects to rule out reputation effects, and sometimes even full anonymity between the experimenter and the subjects. Many experiments investigating strong reciprocity were first based on a student subject pool, but the results have since been extended to adult samples in advanced industrial societies (Falk 2004), and in a variety of cross cultural contexts (Henrich et al. 2001, 2005, 2006). Moreover, data based on nationally representative samples (Fehr et al. 2002; Bellemare and Kröger 2003) have largely replicated the original experiments with student subject pools.

THE PROBLEM OF SOCIAL ORDER AND COOPERATION

The problem of social order goes back to at least Thomas Hobbes who argued that “the life of man (is) solitary, poore, nasty, brutish, and short” in the state of nature. Hobbes concluded that social order is the product of powerful social institutions, including property rights, codified law, and a strong state. Hobbes' approach in modern times has been strongly espoused by neoclassical economic theory, which has applied general equilibrium and repeated game theory to show that these institutions permit large-scale cooperation among unrelated self-interested individuals. However, in an evolutionary time frame, these order-producing institutions came into place only very recently. Humans had to solve the problem of social order long before they invented and implemented these institutions. In fact, the very existence of these order-producing institutions is itself a result of foregoing social cooperation. We therefore must search for more basic mechanisms that could already generate social order in much simpler societies. Unfortunately, conventional repeated game theory – which is based on the assumption of Homo Economicus – has failed to produce plausible analytical models of social cooperation in a state of nature because these models do not have the required properties of dynamical stability and informational robustness (Gintis 2004)

Of course, theorists working in the sociological tradition have long been skeptical of the neoclassical model of society based on self-interest alone, and are hence not at all surprised at this failure of economic theory. At least since Durkheim (1938), sociologists have referred to the civilizing power of the internalization of social values that restrain individuals' self-interest. The role of internalized social values in the constitution of social order is most clearly stated in the work of Talcott Parsons (1937). Parsons argued that individuals internalize social values and feel a need for social approval to such a degree that there is little conflict between self-interest and social values, except for a minority of "social deviants".

The Parsonian solution to social order, however, fails to integrate the mechanisms of norm internalization and the need for social approval into a coherent model of individual choice and social interaction (Wrong, 1961; Gintis 1975). In particular, it neither explains how individuals adjudicate between satisfying personal material needs and social obligations, nor does it clarify the conditions under which individuals accept and reject alternative normative principles. Rather than attempting to repair this *lacuna* in the Parsonian framework, sociologists generally rejected the approach completely. Sociology is now in the difficult position of embracing the internalization of norms as a fundamental aspect of social life, but without a model of individual behavior to which this concept can be appropriately attached.

Integrating the internalization of norms into decision theory can only be accomplished based on extensive empirical research. The sorts of armchair speculation often found in the discussion of "human nature" simply will not suffice. Fortunately, experimental approaches in psychology (e.g., van Lange 1997, 1999), political science, (e.g., Ostrom, Walker and Gardner 1992), sociology (e.g., Marwell and Ames 1979; Yamagishi 1986; Kollock 1998; Raub and Snijders 1997, Simpson 2004) and economics (e.g., Andreoni 1988; Ledyard 1995) have already planted the seeds for a solution to this problem.

Social order as a public good

How do self-regard and social norms interact to determine individual behavior? Laboratory experiments provide a powerful tool for answering this question. One key aspect of social order can be captured by *public goods experiments*, in which self-interest and the social good are counter posed with great clarity. In a public goods experiment, each member of a group of $N \geq 2$ people is endowed with \$Y. Each group member can keep this money or invest up to Y into a group account which represents the public good. The experimenter multiplies every dollar invested on the group account by a factor M (which exceeds 1 but is smaller than N, 1

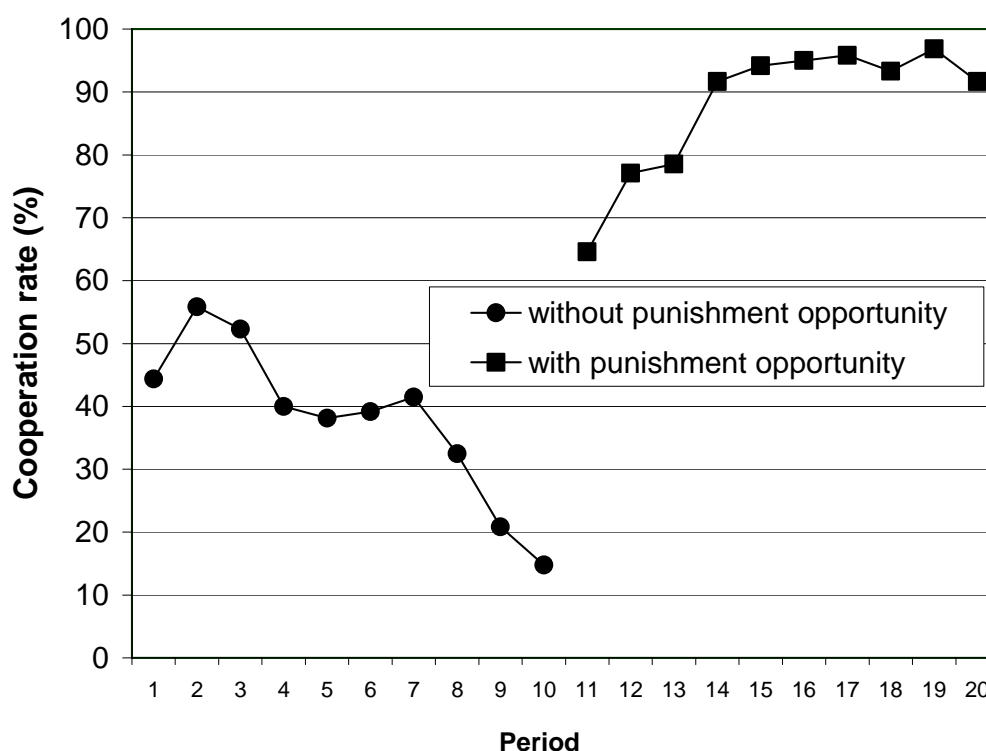
$< M < N$). When all group members have made their contributions to the group account, and when the experimenter has multiplied this amount by M , the multiplied amount is equally distributed among the group members. Thus, if a subject contributes, say, \$5 to the group account, the group as a whole earns $5 \times M$ from this contribution, and each individual, including the contributor of the \$5, receives $5 \times M/N$ from the group account. Note that because $M < N$, the investment of 5 exceeds the return of $5 \times M/N$, i.e., *the investing individual always decreases his or her economic payoff when contributing to the public good*. Thus, a self-interested individual will never contribute anything to the public good. However, if all N group members contribute their individual endowment of Y , each individual earn $N \times Y \times M/N = Y \times M$. For example, if $M = 2$ and $N = 10$, each individual can double his or her income by contributing everything to the public good relative to a situation where everybody keeps the endowment. But each individual also has a strong private incentive to free ride on the contributions of others because every \$ invested yields only a private return of \$2/10.

The incentive structure captured by the public goods experiment is ubiquitous in reality, and captures goods like public security (i.e., the absence of violence and crime), environmental pollution problems, contract enforcement, and team compensation problems, as well as the enforcement of all kinds of rules that are beneficial for the common good but costly for individuals to obey. If a stable group of individuals has the chance to repeat the experiment several times, and if the group receives feedback about the individual contributions at the end of each repetition, cooperation then typically converges to very low levels [see Figure 1 from Fehr and Gächter (2000) which is based on six groups with $N = 4$, $Y = 20$, $M = 1.6$ and ten repetitions (periods)]. In period ten, roughly 55% of the subjects contribute nothing to the public good and the remaining subjects contribute very little. Thus, it seems that – at least towards the end – individual self-interest largely dominates behavior and the view that internalization of cooperative values in itself suffices to generate social cooperation is clearly wrong. The experimental evidence (Ledyard 1995) strongly suggests that Homo sociologicus, who is socialized to obey the prevailing social order without regard to self-interest, is indeed a rare creature.

However, the first 10 periods presented in Figure 1 tells only half of the story. After period 10, the experimenters implemented the following variation of the experiment. They gave subjects the opportunity to punish each other at the end of each period, after each group member's contribution in that period was revealed. Each subject could assign punishment points to each of the other group members which decreased both the punisher's economic

payoff and that of the punished subject. Each punishment point assigned to another group member reduced the punisher's payoff by 1 money unit and the punished subject's payoff by 10 percent (i.e, by roughly 3 money units). Note that the punishment opportunity should not change the cooperation incentives within a group if all group members are purely self-interested because self-interested subjects will never punish if it is costly – at least towards the final round of the experiment. Therefore, adding this punishment opportunity should not generate stable cooperation if all subjects are only self-interested and internalized values play no role for behavior.

Figure 1: Cooperation in the absence and the presence of private punishment opportunities



The subjects' actual cooperation behavior shows a radically different pattern. In the presence of the punishment opportunity, behavior converges to almost full cooperation (see squares in Figure 1). In period 20, 83% of the subjects contribute their whole endowment to the public good, and most of the remaining subjects contribute close to the maximum, so that the average cooperation rate reaches almost 100%. Note that exactly the same subjects generated both the behavioral pattern of unraveling cooperation over the first ten periods as well as that

of high levels of cooperation in the second ten periods. In fact, there were groups in which every group member contributed nothing to the public good in period 10 and everything in period 20. Several other researchers have documented similar results to Figure 1 in somewhat different settings (Anderson and Putterman 2006, Ostrom et al. 1992, Yamagishi 1986)

How is it possible that the same subjects exhibit such radically different behavioral patterns? Whereas the observations in period 10 suggest that humans are largely self-interested, the observations in period 20 suggest that the Parsonian view of social order has an element of truth because it is hard to explain why people cooperate at almost maximal levels even in round 20 in the absence of internalized social values. The subjects knew that the experiment ends in round 20. Therefore, purely self-interested subjects will never punish in period 20 because this is costly for them and generates no future economic benefits. But if nobody punishes in the final period, why should self-interested subjects then cooperate in this period? In addition, what should we infer from the fact that a cooperation rate of roughly 50% is maintained for the first few periods in the absence of a punishment opportunity?

SOCIAL COOPERATION AND STRONG RECIPROCITY

The puzzle posed by Figure 1 can be resolved if one takes the fact that the subject population is heterogeneous in specific ways into account. Ample evidence from different types of experiments (Fehr and Fischbacher 2003; Gintis et al. 2003) indicates that a substantial share of the subjects exhibits strong reciprocity. However, the same experiments also show that a large share of subjects is indeed completely self-regarding. Recall that strong reciprocity is a behavioral propensity to cooperate conditionally on other group members' cooperation and to punish the violations of social norms even though this is costly for the punisher and causes him an economic net loss. Thus, strong reciprocators must have internalized cooperative social values because not only their self-interest shapes their behavior, but conditional cooperation and punishment motives as well. This means that strong reciprocity is not simply long-term, enlightened self-interest because the strong reciprocator cooperates and punishes even in anonymous one-shot interactions where the cooperative and punishing acts obviously reduce his economic net gain. The notion of strong reciprocity also differs from Gouldner's (1960) concept of reciprocity; in particular, Gouldner's definition does not include the punishing aspect of strong reciprocity.

A large body of research has documented important motivational forces behind strong reciprocity. Two of the most prominent forces have been termed “reciprocal fairness” (Rabin 1993, Falk and Fischbacher 2006) and “inequity aversion” (Fehr and Schmidt 1999). A reciprocally fair subject is motivated by the desire to respond to kind acts with kindness and to hostile acts with hostility. An inequity averse subject is motivated by the desire to avoid inequity and to implement equitable outcomes. Inequity averse and reciprocally fair subjects are no saints who resist unfair outcomes and punish unfair behavior under all circumstances. Rather, these subjects value equity and reciprocal fairness *in addition* to their economic self-interest, implying that if the costs of maintaining equity or of reciprocally fair behaviors increase they are less likely to engage in these behaviors (Anderson and Putterman 2006, Carpenter in press). Inequity aversion and reciprocal fairness are also motives that lead to purposeful behavior that can be mathematically represented by the appropriate utility functions. This possibility has the great advantage that the powerful apparatus of modern game theory can be applied to our problem at hand *and* to many other related problems such as social exchange or competitive market behavior. This enables the development of a unifying model that explains behavior in a wide variety of domains.

Models of inequity aversion (Fehr and Schmidt 1999) and reciprocal fairness (Falk and Fischbacher 2006) can, in particular, explain why we observe little cooperation in the absence of a punishment opportunity and why cooperation flourishes when punishment is possible. In the absence of a punishment opportunity, strongly reciprocal subjects will initially cooperate if they believe that others will also cooperate. However, they notice over time that other group members – the self-regarding ones – free ride. As strong reciprocators are only willing to cooperate if most others also cooperate, they cease to cooperate in the later periods of the experiment. Strong reciprocators also have a desire to punish free riders because they perceive free riding on their cooperation to be unfair. However, stopping cooperation is the only way to punish other group members in the absence of a direct punishment opportunity that enables the subjects to target the punishment on the free riders. Thus, both the free riders and the strong reciprocators contribute little or nothing to the public good towards the end, albeit for different reasons. The self-regarding subjects contribute nothing because this maximizes their economic payoff. The strong reciprocators contribute nothing because they are only willing to cooperate if sufficiently many others also cooperate, and because free riding is the only way available to punish the self-regarding subjects.

The argument presented above shows that the self-regarding subjects ultimately induce the reciprocators to free ride as well in the absence of a direct punishment opportunity. This

contrasts sharply with the behavior observed in the presence of a direct punishment opportunity. Here, virtually all subjects ultimately cooperate despite the fact that there is a substantial share of self-regarding subjects. The reason is that the strong reciprocators can now punish the defectors directly, creating an economic incentive for the self-regarding subjects to cooperate. Moreover, the strong reciprocators will also cooperate because they need not fear others' defection, as the self-regarding individuals are disciplined. Thus, the strong reciprocators induce the self-regarding subjects to cooperate in the presence of a direct punishment opportunity.

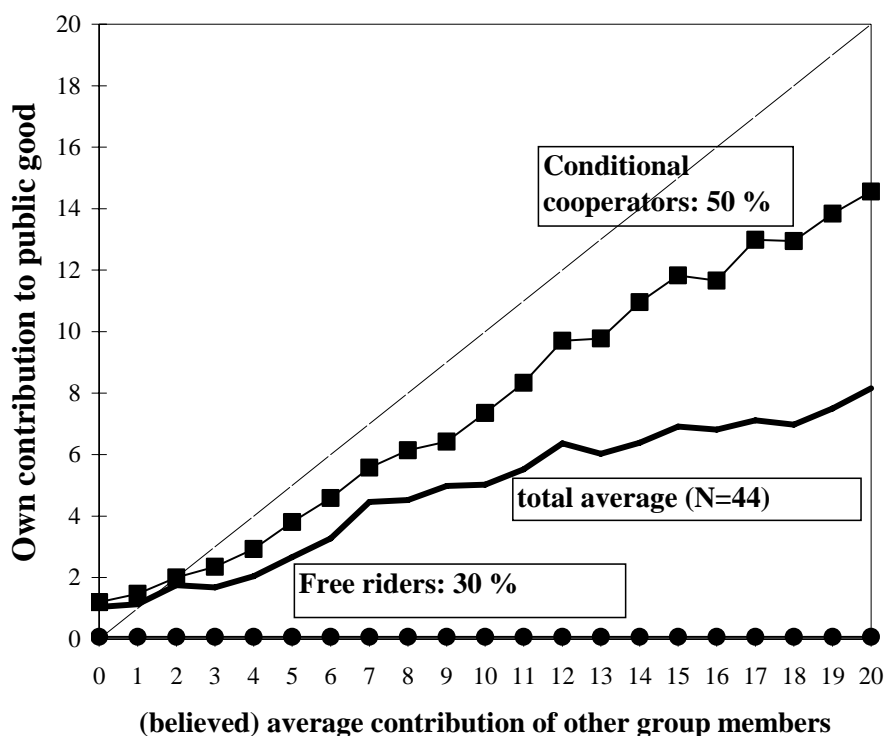
Strong reciprocity in public goods provision – evidence

To what extent can further facts enable verification of the theory advanced above? In particular, to what extent can we provide explicit evidence for the existence of self-regarding subjects and of conditional cooperators who are willing to punish free riders? In Figure 2 (from Fischbacher et al. (2001)), we show the existence of self-regarding free riders and conditional cooperators. The figure is based on the same public goods incentives (without a direct punishment opportunity) as in Figure 1 above. The squares show the conditional cooperator's average behavior. These subjects increase their cooperation level if they believe that the other group members will do so as well; 50% of the subjects behaved in this way. The figure also shows that – on average – the conditional cooperators do not completely match the others' expected average contribution, but they typically remain slightly below the level of others (i.e. below the broken 45 degree line). This indicates that the conditional cooperators' self-interest also affects their choices. In addition, the circles in Figure 2 show the existence of a large fraction of subjects who always free rides, regardless of what the others do. These subjects can be categorized as purely self-regarding; they comprised 30% of the participants. The remaining 20% of the subjects exhibited other behavioral patterns such as an increasing response to others' contributions over the first 10 expected contribution units while afterwards they show a decreasing response. Importantly, however, the vast majority of subjects can be either classified as purely self-regarding and conditionally cooperative – a fact that has since been frequently replicated (e.g., Fischbacher and Gächter 2006).

Figure 2 contains the basic rationale explaining why the maintenance of a high cooperation level is not possible in the absence of a direct punishment opportunity. Suppose, for example, that all subjects initially believe that the other group members will contribute their whole endowment of $Y = 20$ to the public good. Then the response curve of the

conditional cooperators (the squares in Figure 2) indicates that for this belief they would contribute approximately 14 units, while the self-regarding subjects would contribute nothing. Thus, the information feedback at the end of the period would reveal that the actual average contribution is considerably below the expected average contribution of 20. The subjects are therefore likely to revise their beliefs about others' expected contribution downwards which would induce the conditional cooperators to reduce their contribution level in the next period. In this way, the heterogeneous mix of self-regarding and strongly reciprocal subjects generates a downwards trend in cooperation levels.²

Figure 2: Percentage of free-riders and conditional cooperators in the public goods experiment



This argument contains a further important lesson. The fact that social order sometimes breaks down – for example, after natural disasters or at the end of a war (as in Iraq) – does not prove that all people are self-regarding. The social order can also break down even if a large share of people have internalized cooperative social values and are, thus, in principal willing

² If the self-regarding subjects rationally anticipate the response of the reciprocal subjects, they will also contribute during the initial periods of a finitely repeated game. By contributing in period t they can induce higher contribution levels of the reciprocal subjects in the subsequent periods. However, towards the end, these future (selfish) returns from current contributions decline so that rational egoists will cease to cooperate, which then induces the reciprocal subjects to also stop cooperating (Ambrus 2006).

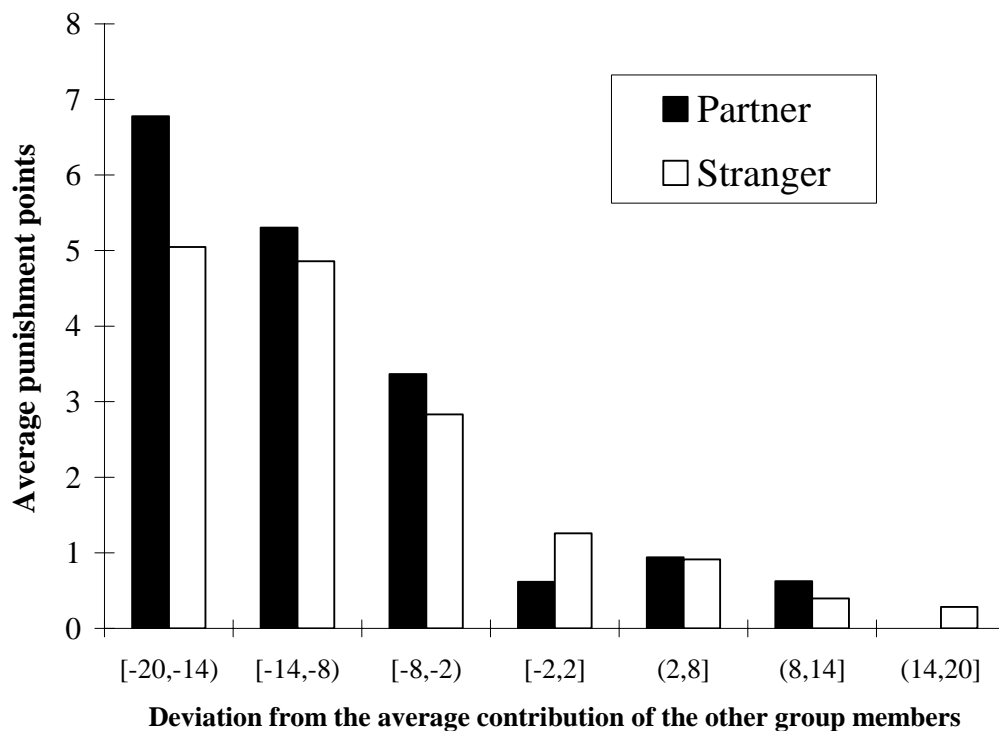
to cooperate. However, if free riders go unpunished, the cooperative individuals also tend to stop cooperating. In fact, Fehr and Schmidt (1999) have shown theoretically that a relatively small minority of self-regarding individuals suffices to generate a breakdown of cooperation in the public goods game if the free riders cannot be punished individually. Likewise, a relatively small minority of inequity averse people also suffices to generate a fully cooperative outcome if punishment can be directly targeted at the individual free rider.

Figure 3 provides further evidence for this explanation. This figure (based on Fehr and Gächter 2000) illustrates the punishment pattern in the experiment with a punishment opportunity (see black bars). The figure shows that the deviation of the punished subjects' contributions from the average contribution of the other group members is an important determinant of punishment. The more the target individual free rides relative to the group average, the higher is the punishment. For example, if the free rider's contribution deviates between -20 and -14 from others' average contribution, the free rider receives 7 punishment points which reduces his income by 70%. Thus, free riding is strongly punished when the group composition is stable (partner treatment) and, as a consequence, potential free riders have a strong incentive to cooperate.

However, the black bars in Figure 3 do not yet demonstrate that strong reciprocity drives punishment because if the same group of individuals interacts for 10 periods, there may be a selfish incentive to punish free riders. After all, a free rider who is punished increases contribution levels in the next few periods which provides also a benefit for the punisher. This problem is exacerbated if – as in the work of Yamagishi (1986) and Ostrom et al. (1992) – the subjects do not know how many periods they will interact together in a group. In fact, in this earlier work, the researcher deliberately did not tell the subjects the number of periods they will interact together because they wanted to implement a truly repeated game in which there are also selfish incentives to cooperate and to punish. In order to test whether self-regarding motives or strong reciprocity drives punishment, Fehr and Gächter (2000, 2002) conducted experiments in which the group composition changed randomly from period to period in such a way that no subject ever encountered another subject more than once. In this setting (called stranger in Figure 3), a punisher could not reap any benefit from the future cooperation of the punished free rider because the free rider was never in the punisher's group in future periods. Therefore, there cannot be any selfish benefit from punishing a free rider in the stranger setting. However, as Figure 3 shows (see white bars), punishment is still very high, indicating the existence of strong reciprocators. In fact, the differences in the punishment pattern between the partner and the stranger setting are not significant, suggesting that other-

regarding motives drive the bulk of punishment choices. This interpretation has been further confirmed by so-called third party punishment experiments (Fehr and Fischbacher 2004, Henrich et al. 2006), where “uninvolved” third parties, who observe a norm violation, can punish the norm violator at a cost.

Figure 3: Punishment pattern in stable groups (partners) and groups with randomly changing membership (strangers)



Evidence on the neurobiological underpinnings of some aspects of strong reciprocity has been procured by combining the neuroimaging tools of modern neuroscience with the tools of experimental economics. It has been shown (DeQuervain et al. 2004) that a key component of the human reward system is activated if individuals are given the opportunity to punish those who cheated in a social exchange. Thus, the punishment of cheaters seems to be associated with a direct positive hedonic impact, indicating the sweetness of revenge. In addition, those individuals who showed the highest activations in the reward system when they decided about punishment spent most money on punishing the cheaters, confirming the importance of individual differences in strong reciprocity. Singer et al. (2006) have shown that merely observing the punishment of a cheater also activates important components of the human reward system. These findings suggest that the motivational forces behind the punishment of norm violators have deep neurobiological roots.

SOCIAL EXCHANGE AND STRONG RECIPROCITY

Our theory shows how different social structures, such as the degree to which targeted punishment opportunities are available, may generate completely different aggregate patterns of interaction. Because the same individuals cause these divergent patterns, we cannot attribute the divergence solely to individual characteristics; the social structure is causally involved in this emergent pattern. However, our theory does not neglect individual characteristics. Rather, it is based on a clean characterization of empirically supported individual properties (i.e., selfishness and strong reciprocity) that provides behavioral foundations for sociological analyses. In fact, if only self-regarding individuals existed, the two social structures would generally produce the same outcomes. We can derive the emergent macro-patterns by combining the heterogeneity of individual characteristics with the constraints on individuals' action space (which is shaped by the prevailing social structure) in a game theoretic framework. In addition, our theoretical account – based on Fehr and Schmidt (1999) and Falk and Fischbacher (2006) – shows how and why the social structure generates non-obvious, perhaps counterintuitive, patterns. In the absence of a targeted punishment opportunity, even a large majority of strong reciprocators will not be able to sustain cooperation, whereas even a relatively small number of strong reciprocators can enforce a fully cooperative outcome in the presence of a punishment opportunity. Thus, our analysis shows that the BPC approach provides powerful microsociological foundations for large-scale social phenomena.

In the following, we apply this approach to the problem of social exchange. The concept of social exchange provided the basis for a movement to develop an axiomatic sociological theory starting from first principles (Homans 1958, Emerson 1962, Blau 1964, Cook 1987). A social exchange involves the mutual transfer of benefits such as goods, ideas, aid, or social approval under conditions of incompletely specified obligations (Blau 1964). The incomplete specification of the exchange partners' obligations implies that such a social arrangement cannot be enforced through binding contracts serviced by third parties (e.g., the judiciary). Sociologists noted long ago that social exchanges permeate most human interactions and are of relevance for many economic relations such as the employment relationship or the provision of complex professional services. In fact, social exchanges are likely to affect almost all long-term relationships between economic actors. This has led to a rich theoretical

literature – summarized in Emerson (1976), Cook (1987) and Molm (1997) – and many sociologists conducted powerful experiments that illuminated important aspects of the problem such as the relationship between power and network structures (e.g., Cook et al. 1983), the relationship between power and commitment (e.g., Lawler and Yoon 1993 and 1996) or the role of coercion in social exchanges (Molm 1997).

In contrast, economics largely ignored social exchange until recently. However, current progress in game theory and in experimental economics have made it possible to model social exchanges with powerful analytical methods that enable us to get a theoretical and empirical grip on understanding fundamental issues. How are social exchanges enforced? After all, the partners' obligations are ill-specified. So why should the exchange partners obey such ill-specified "terms of trade"? How do social exchanges affect social relations between the exchange partners? How do successful social exchanges get started? What is their impact on aggregate entities like prices and efficiency? How do social exchanges affect competition between actors and vice versa? Social exchange theorists such as Blau (1964) argued that the supply and demand of the traded benefits and social norms of fairness affect prices in social exchanges. Yet, it has never been made precise how exactly the laws of supply and demand interact with fairness norms.

In the following, we will show that the same methods and principles that allow us to understand multilateral social cooperation also enable us to answer many of the questions mentioned above. In particular, the specific heterogeneity of the population, i.e., the coexistence of purely self-regarding and strongly reciprocal actors, and the interaction of this heterogeneity with the prevailing social conditions, will again be the key factor in our explanation.

Social exchanges in the laboratory

Laboratory experiments are a convenient tool for the study of social exchanges. In this context, the gift exchange game (Fehr, Kirchsteiger and Riedl 1993), which is inspired by the work of Akerlof (1982), has been particularly useful. In this game, one subject takes the role of an employer who offers a wage W , the other subject takes the role of an employee who provides effort E in exchange for W . First the employer offers a *binding* wage payment W , then the employee observes W and responds with a choice of E . Because the wage offer is binding, W has to be paid regardless of how large the employee's effort E will be. This feature simplifies the situation but it is not essential for our main arguments. The employer's

monetary payoff, denoted by π , is given by $\pi = A \times E - W$ where A is a positive constant determined by the experimenter; higher effort increases and a higher wage reduces π . The employee's monetary payoff, denoted by U , is given by $U = W - C(E)$ where $C(E)$ represents the cost of effort. As higher effort levels are (by the construction of the experiment) associated with higher cost levels, a higher effort reduces U . We also assume, for the moment, that the two parties play the game only once. This game captures the key element of a simple social exchange situation because the employee is completely free to choose whatever effort level he likes if there is no third party who can enforce a contractually specified effort level. The absence of third party enforcement requires that one of the following two conditions be met: (i) the employer cannot stipulate a contract with a well specified effort level. Third party enforcement cannot function in this case because it is unclear what the third party should enforce. (ii) the third party cannot verify the actual effort level the employee chooses. In this case, even if the employer is able to stipulate a contract with a well-specified effort level, the third party simply does not know what the employee did and can therefore not decide whether the employee met or violated the contract.

In the following we always assume— unless stated otherwise — that one of these two conditions is met. This ensures the existence of the key problem in social exchanges — the enforcement problem, and it enables us to study how people behave when they face this problem. Assume, for example, that the gift exchange game takes place only once. What pattern of behaviors should we then expect? If both the employer and the employee are purely self-regarding, caring only for their own economic payoff, then the employee will always choose the lowest possible E , because this maximizes U , and the employer will choose the lowest possible W , because this maximizes π . However, if there are sufficiently many strong reciprocators among the employees, who respond to a kind act, that is, to a fair wage, with fair (non-minimal) effort levels, the employer may have an incentive to pay fair wages. The possibility of this kind of reciprocal fairness gave this game its name because if third parties cannot enforce the effort the exchange partners essentially can only exchange “gifts”. In fact, reciprocally fair exchange patterns have been frequently observed in the gift exchange game (e.g., in Fehr, Kirchsteiger and Riedl 1993): If the employer pays higher wages, the employees chose — on average — a higher effort level. This effort pattern indicates that — due to the existence of strong reciprocators — some enforcement in social exchange situations comes from people's internalized social values.

The consequences of social exchanges

The existence of sufficiently many strong reciprocators among the employees may even transform the interactions in labor markets in a profound way: profit-maximizing employers now have an incentive to pay “fair” wages because higher wages cause – on average – a higher effort level. Therefore, wages in labor markets that are characterized by social exchanges between employers and employees may not be solely determined by the laws of supply and demand but by the employers’ incentive to pay wages that the employees perceive to be fair. This possibility has been explicitly documented by Fehr and Falk (1999), who compared a bilateral gift exchange experiment with a competitive gift exchange experiment. In both experimental conditions, an employer could only employ one worker, but there was an excess supply of workers in the competitive condition, whereas the number of workers and the number of employers was the same in the bilateral condition. Thus, the employees had to compete for getting a job offer in the competitive case because there were always fewer jobs than workers. According to standard economic principles, the excess supply of workers should drive wages in the competitive conditions below those in the bilateral condition. However, as Fehr and Falk (1999) showed, the firms in both conditions had an incentive to pay “fair” wages due to the existence of strong reciprocators and, therefore there was no difference in wage levels across conditions. Thus, this experiment lends support to the views expressed by social exchange theorists such as Homans (1958) and Blau (1964). Fairness norms are likely to have an important effect on price formation in social exchanges by weakening the impact of supply and demand forces. The analysis above also goes beyond the work of these authors, however, by making precise how fairness norms shape price formation: the internalized social values of strong reciprocators induce them to respond to fair wages with fair effort levels which then provides an economic incentive to pay fair wages.

However, social exchanges may even transform the pattern of market exchanges in a more fundamental manner than simply by affecting wages or prices. They may radically change the way in which trading partners interact with each other. Kollock (1994) and Brown et al. (2004) have shown that the enforcement problem inherent in social exchanges causes long-term relationships between the trading partners. If third parties enforce contractual obligations, i.e., if the possibility for social exchanges is absent, there is no necessity for the traders to trust each other and, therefore, they do not care about their partner's identity. As a consequence, trade predominantly takes place in one-shot interactions. However, traders need to trust each other in the absence of third party contract enforcement and, as a consequence, long-term relationships emerge. Traders care a lot about the trustworthiness of their partner

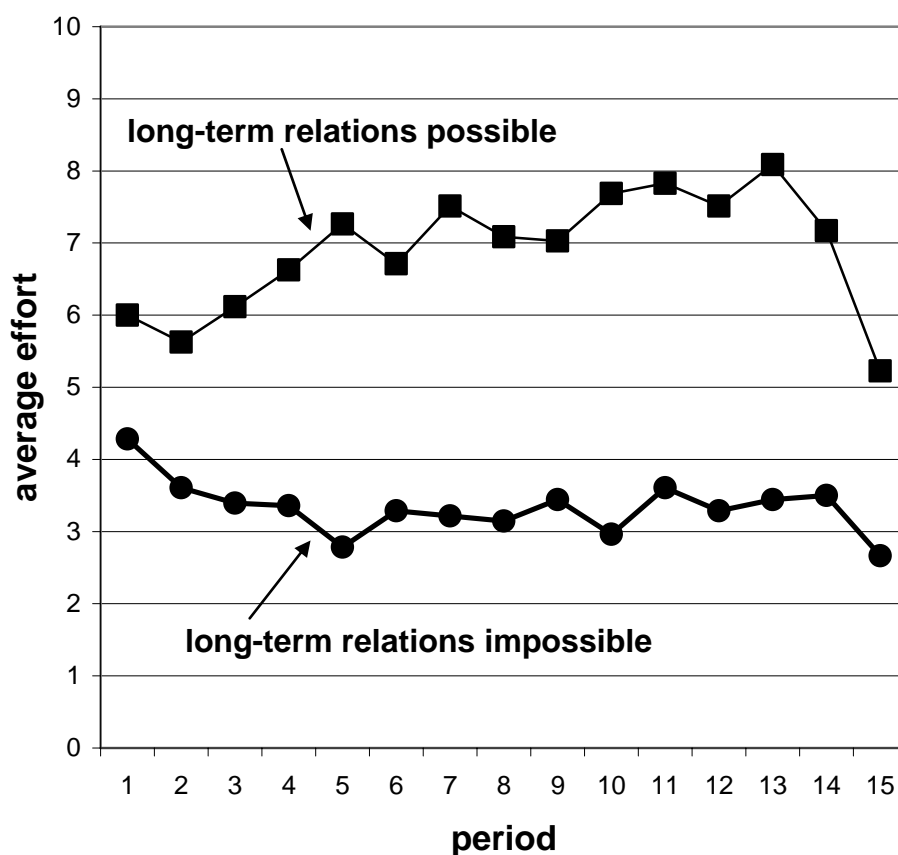
under social exchange conditions. It is therefore essential for them to know with whom they are trading while the partners' identity plays no role under third party enforcement of contracts. Trade is therefore initiated under social exchange conditions by making offers to *specific* potential trading partners who have a reputation for being trustworthy, whereas offers are usually made to *the whole group* of potential trading partners under third party enforcement because this maximizes the chances that one of the potential partners accepts the offer.

In Figure 4 – based on Brown et al. (2004) – we show that these trust-enhancing strategies of the trading partners are associated with an enormous increase in effort enforcement relative to a situation where only one-shot interactions are possible. The experiment lasted 15 periods in both the “one-shot” and the “relations” treatment, and the subjects in the role of an employer could make wage offers to the workers in each period. There was always an excess supply of workers, i.e. some workers were unemployed in every period. These wage offers could either be targeted to specific workers or to the whole group of workers. Once a worker had accepted a wage offer he chose an effort level between 1 (the minimum effort) and 10 (the maximum effort). Also, the employers and the workers were assigned identification numbers in both treatments. However, the identification numbers were randomly assigned to the subjects in every period in the “one-shot” treatment, whereas the assignment of the identification numbers was fixed throughout the 15 periods in the “relations” treatment. Thus, if employer no. 5 traded with worker no. 8 in period t in the relations treatment, and if the employer was satisfied with the performance of worker no. 8, he could make again a wage offer to the same individual (i.e., worker no. 8) in the next period. This possibility of repeated interaction did not exist in the one-shot treatment, where a different individual was assigned the “worker no. 8” label in every period.

Figure 4 shows that the effort enforcement is quite limited if long-term relations between the trading parties are ruled out, whereas rather high levels of effort can be enforced if the possibility of long-term interactions is added. The key for explaining this treatment difference again lies in the interaction between self-regarding and reciprocal types. If no long-term relations are possible, the self-regarding individuals will never choose non-minimal effort levels. Only the strong reciprocators will provide non-minimal effort depending on the wage they are offered. However, if long-term relations are possible, the existence of strong reciprocators among the workers generates the possibility to discipline the self-regarding workers. Due to the existence of strong reciprocators, it is profitable for employers to pay fair wages – in excess of the income of being unemployed – even in the final period of the

experiment. Thus, the employer can implicitly threaten to punish the worker for all non-final periods by not renewing the relationship in the next period. The refusal to renew a relation always imposes a loss on the worker because – due to the excess supply of workers – the worker remains unemployed with positive probability. Therefore, the self-regarding workers have strong incentives to provide non-minimal effort levels in the “relations” treatment, incentives that are based on the existence of the strong reciprocators.

Figure 4: Average effort with and without the opportunity to form long-term relations



Social exchanges not only transform the interaction patterns between the trading partners in radical ways, but also cause a very different distribution of the earnings from trade. Brown et al. (2004) show this by comparing the relations treatment with a treatment in which third parties enforce the effort level. The two treatments are identical in all other respects; in particular, there is an excess supply of workers in both treatments, putting the workers in a weak bargaining position. As Figure 5 shows, the workers’ weak bargaining position causes a very uneven distribution of earnings when third parties enforce the contract, whereas the

distribution is much more equal under social exchange conditions (in the relations treatment). Therefore, the existence of the enforcement problem, which gives reciprocal workers the possibility to vary effort according to their fairness preferences, counteracts the weak bargaining position that is generated by their excess supply.

Figure 5: Distribution of the earnings per trade between employer and employee

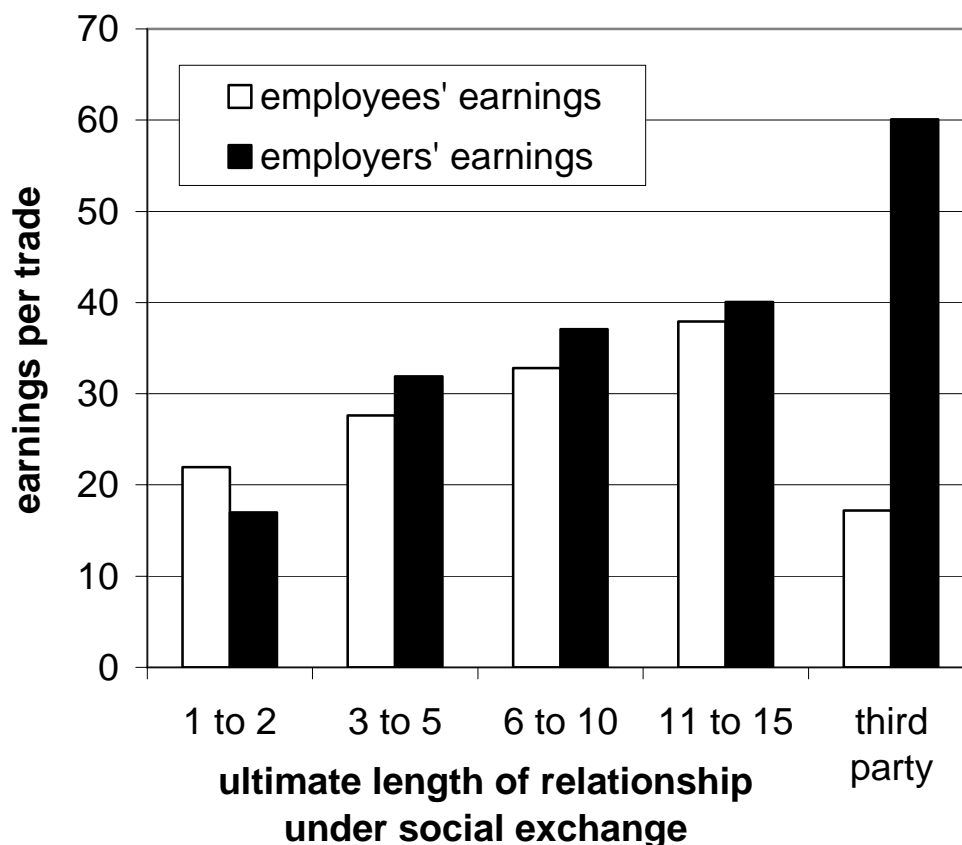


Figure 5 also nicely illustrates the efficiency gains that accrue in relationships with a longer duration. The longer a relationship ultimately lasts in the relations treatment, the higher are the earnings of both the employer and the employee. These efficiency gains are a result of the fact that the effort level is higher in longer lasting relationships. Employers did not renew the relation with those workers who did not provide sufficient effort. Therefore, relationships which survive longer have – on average – a higher effort which also causes higher total earnings in these relationships. In fact, in those relationships that last longest (11-15 periods), the sum of the earnings is as high as under third party enforcement.

MEASURING THE IMPACT OF SOCIETY ON INDIVIDUALS' PREFERENCES

Economists often take individual preferences as given and ask how the interplay between given preferences and institutions shapes social and economic outcomes. In contrast, many sociologists (e.g., Durkheim 1951, Benedict 1934) and psychologists (e.g., Grusec and Kuczynski 1997) have pointed out that “society” shapes individuals’ preferences through socialization processes. However, the other behavioral sciences often ignore and implicitly reject these observations, which are part of the core of sociological theory. This situation is due in large part to the difficulties in providing clean behavioral measures of preferences – measures that must not be confounded by individuals’ beliefs or by reputational and repeated interaction incentives.

Laboratory experiments provide a powerful tool for verifying the overarching importance of internalized values in modeling social interaction. We shall illustrate this point with several examples of controlled experiments. We begin with the cross-cultural experiments of Henrich et al. (2001, 2005) in 15 small scale societies. The experimenters conducted ultimatum games in all 15 societies; they also conducted public goods and other experimental games in some societies. In an ultimatum game, two players – a proposer and a responder – bargain about the distribution of a given sum of money (or some other valuable resource) according to the following rules: the proposer makes one proposal how the money should be distributed among the two; the responder observes this proposal and accepts or rejects it. In case of acceptance, the money will be split according to the proposal. In case of a rejection, both players receive nothing. If both players are purely self-regarding, the responder will accept any positive amount, however small it is. Therefore, a proposer who anticipates the responder’s behavior will offer the smallest positive amount to the responder. However, numerous experiments in Western countries have shown (Güth and Tietz 1990) that responders frequently reject a sizeable fraction of the available money even if the stakes are as high as three months income (Cameron 1999). Note that self-regarding preferences cannot explain a rejection of a positive offer in an anonymous one-shot ultimatum game; instead, rejections most likely reflect internalized fairness norms. Due to the high rate of rejections, even self-regarding proposers have an incentive to make relatively fair offers.

How does the pattern of proposals and rejections vary across societies and how are these differences related to salient aspects in the societies’ social life? Henrich et al. (2001, 2005) find that differences in societies in “market integration” and “cooperation in production” explain a substantial portion of the behavioral variation between groups. The societies were

rank-ordered in five categories – “market integration” (how often do people buy and sell, or work for a wage), “cooperation in production” (is production collective or individual), “anonymity” (how prevalent are anonymous roles and transactions), “privacy” (how easily can people keep their activities secret), and “complexity” (how much centralized decision-making occurs above the level of the household). Using statistical regression analyses, only the first two characteristics, market integration and cooperation in production, were significant, and they together accounted for 66% of the variation among societies in mean ultimatum game offers.

Another important result from this study is that experimental behavior mirrors patterns of interactions found in everyday life in that society. For example, among the Papua New Guinean tribes of the Au and the Gnau, many proposers offered more than half of the available sum of money, and many of these “hyper-fair” offers were rejected! This reflects the Melanesian culture of status-seeking through gift-giving. Making a large gift is a bid for social dominance in everyday life in these societies, and rejecting the bid is a rejection of being subordinate. Another example is given by the Hazda, a tribe located in Tanzania. They made low offers and had high rejection rates, which mirrors the tendency of these small-scale foragers to share meat, but with a high level of conflict and frequent attempts of hunters to hide their catch from the group. This behavior contrasts sharply with those of the Ache, a tribe located in Peru. Almost all offers made by the Ache are close to 50% and no rejections occurred. In daily life, the Ache regularly share meat, which is distributed equally among all the households, irrespective of which hunter made the catch.

The impact of prevailing social practices on internalized social values can also be studied within advanced societies. Burks et al. (2006) examined the impact of piece rate incentive schemes, hourly payment schemes, and team based pay on the degree of conditional cooperation exhibited by workers of several bike messenger companies in Switzerland and the US in a sequentially played prisoners’ dilemma (PD) game. They measure conditional cooperation by the frequency with which the second movers in the sequential PD cooperate in response to first mover cooperation. Conditional cooperation is significantly more frequent in firms that pay hourly wages or where the revenue is shared among the workers than in firms with piece rate incentives. The lower degree of conditional cooperation in firms with piece rate incentives could be due to the possibility that self-regarding workers self-select into these firms. Yet, it could also be caused by the incentive system itself; piece rate incentives renders mutual help costly for the workers and may, therefore, make workers more self-regarding. Burks et al. show that the latter is the more likely reason because firms with piece rate

incentives are not located in the same cities as firms that pay on an hourly basis or share the revenue among the workers. Because the cities are relatively far from each other, the bike messengers typically cannot choose to work in firms with different incentive schemes, rendering the self-selection explanation unlikely.

Another exciting example of the use of experiments for studying the impact of social practices on preferences is given in Gneezy et al. (2006). Their work builds on the paper by Niederle and Vesterlund (in press) who show that women in the US are competition averse. When given the choice to enter a piece rate incentive scheme (that rewards absolute experience) and a tournament incentive scheme (that rewards only the winner of the tournament) women predominantly prefer the piece rate scheme while the vast majority of US men prefer the tournament scheme. This preference pattern across men and women even prevails if they control for risk preferences and equity preferences. Gneezy and List hypothesized that women's competition aversion might have to do with the fact the US culture is still dominated by males. Therefore, they conducted the same experiments in a paternalistic culture, among the Maasei in Tanzania, and in a matrilineal society, the Khasi in India. Among the paternalistic Maasei, men were more likely to prefer the tournament incentive than were the women, but in the matrilineal society of the Khasi, the gender results were turned upside down: the majority of the women (54%) preferred the tournament while a majority of the men (61%) preferred the piece rate incentive. These results are consistent with the notion of culturally determined gender roles are an important determinant of preferences for or against competition.

CONCLUDING REMARKS ON THE BELIEFS, PREFERENCES AND CONSTRAINTS APPROACH

We kept the mathematical apparatus of the BPC approach in the background throughout this paper. However, formal models, such as those of inequity aversion (Fehr and Schmidt 1999) and reciprocal fairness (Falk and Fischbacher 2006), are indispensable tools of social analysis. We would not have been able to pierce the bewildering puzzles posed by the experimental regularities without the help of such models. Moreover, these models provide a unifying explanation of a vast amount of experimental regularities, many of which we could not present here due to space limits. They help us understand why a minority of selfish individuals may make sustaining cooperation in the absence of a punishment mechanism

impossible, whereas a minority of strong reciprocators may permit cooperation to flourish when a punishment option is available. The models also show how a heterogeneous population of selfish actors and strong reciprocators enforces and transforms anonymous trading relationships in rich social exchanges. Furthermore, they help us understand the vast differences in the distribution of the gains from trade between standard neoclassical markets with third party enforcement on the one hand, and social exchange markets with endogenous enforcement facilitated by strong reciprocators.

The same models also explain, for example, why fairness concerns play a prominent role in bilateral bargaining whereas their role is very limited in competitive markets with third party enforcement (Fehr and Schmidt 1999), why the prisoners dilemma game may be better viewed as an assurance game with multiple equilibria rather than as a prisoners dilemma game, or why employers may deliberately generate social exchanges by offering *incomplete* contracts to their employees rather than complete ones (Fehr, Klein and Schmidt 2007).

The BPC approach rests on decision theoretic insights that show that transitive individual behavior can be represented by the maximization of an objective function subject to the individual's beliefs and the available action space. The particular models that we used in this paper (Fehr and Schmidt 1999, Falk and Fischbacher 2006) go one step further by confining attention to correct beliefs and equilibrium behavior. Sociological theory has widely, and wisely, rejected the wholesale adoption of these assumptions, but in view of the insights generated we believe that a central place should be reserved for them in sociological theory because they play an important, though not hegemonic, role in a unified explanation for the facts mentioned above.

The BPC approach in general, and the particular approaches used in this paper, are obviously not perfect. Confining attention to equilibrium behavior is not always justified. People sometimes hold irrational beliefs and the understanding of belief formation constitutes a difficult problem. However, these problems do not speak against the proper microfoundation of aggregate phenomena in terms of individuals' beliefs, preferences, and constraints. Quite the contrary, we believe that the BPC approach – in combination with careful experimentation – offers a disciplined way to solve these problems. The BPC approach highlights the necessity to examine the distribution of preferences and beliefs in the population empirically and to understand belief and preference formation processes. This empirical knowledge can then be fed back into the model to better predict and understand individual behaviors and aggregate-level phenomena.

While we have proposed the BPC model as a foundation concept for sociological theory, it is clear that the exact same model can serve as a foundation concept for economic theory as well. Indeed, a major thrust of experimental economics has been to convince economists of the need for a far richer notion of human preferences than those given in traditional economic theory. Sociologists need no such convincing, so our task here has been to assert the scientific benefits of framing human action in terms of the BPC model, and stressing the necessity of controlled experiments as a powerful tool for enhancing sociological research.

We thus suggest a thorough unification of sociological and economic theory that is based on the combination of the BPC model with controlled (laboratory and field) experiments. Talcott Parsons (see Parsons and Shils 1951) made the first serious attempt to unify sociological and economic theory. The Parsonian synthesis did not have the benefit of the empirical data now available to us, and hence foundered on Parson's attempt to rely on pure theoretical constructs (pattern variables, structural-functionalism) where data were lacking. Recognizing the potential power of the rational actor model to unify sociological theory, James Coleman (1990), Michael Hechter (Hechter and Kanazawa 1997), Gary Becker (1976) and others made great headway, but were again limited by the lack of solid evidence concerning the nature of preferences. For this reason, these contributions tended to slight the rich menu of human motives recognized in sociological theory, and hence were marginalized within sociological theory. We believe that our proposed principles of unification, by contrast, do not suffer from these weaknesses, and the close interaction between theory development and careful experimentation offers a route for a cumulative elimination of remaining ones.

ACKNOWLEDGEMENTS

This paper benefited a lot from comments made by Michael Biggs, Andreas Diekmann, Jon Elster, Kimmo Eriksson, Peter Hedstrom, Christine Horne, Martina Morris, Trond Peterson, Brent Simpson, and Toshio Yamagishi.

References

- Akerlof G A. 1982. Labor contracts as partial gift exchange. *Quarterly Journal of Economics* 97:543-569
- Anderson C M, Putterman L. 2006. Do non-strategic sanctions obey the law of demand? The demand for punishment in the voluntary contribution mechanism. *Games and Economic Behavior* 54:1-24
- Andreoni J. 1988. Why free ride? Strategies and learning in public goods experiments. *Journal of Public Economics* 37:291-304
- Becker G. 1976. *The economic approach to human behavior*. Chicago: The University of Chicago Press
- Bellemare C, Kröger S. 2003. On representative trust. *Working Paper Tilburg University*
- Benedict R. 1934. *Patterns of culture*. Boston: Houghton Mifflin
- Blau P. 1964. *Exchange and power in social life*. New York: Wiley
- Bonacich P, Light J. 1978. Laboratory experimentation in sociology. *Annual Review of Sociology* 4:145-170
- Boudon R. 2003. Beyond rational choice theory. *Annual Review of Sociology* 29:1-21
- Brown M, Falk A, Fehr E. 2004. Relational contracts and the nature of market interactions. *Econometrica* 72:747-780
- Burks S, Carpenter J, Goette L. 2006. Performance pay and the erosion of worker cooperation. *Discussion paper 2013, Institute for the Study of Labor, Bonn*
- Buskens V, Raub W. 2002. Embedded trust: Control and learning. In *Group cohesion, trust, and solidarity*, ed. E. Lawler, S. Thye, pp. 167-202
- Camerer C F. 2003. *Behavioral game theory - experiments in strategic interaction*. Princeton, New Jersey: Princeton University Press
- Cameron L A. 1999. Raising the stakes in the ultimatum game: Experimental evidence from indonesia. *Economic Inquiry* 37:47-59
- Carpenter J. in press. The demand for punishment. *Journal of Economic Behavior & Organization*
- Carpenter J, Harrison G, List J A, eds. 2005. *Field experiments in economics*, Vols. 10. Greenwich and London: JAI Press
- Coleman J. 1990. *Foundations of social theory*: The Belknap Press of Harvard University Press
- Cook K S. 1987. *Social exchange theory*. Beverly Hill CA: Sage
- Cook K S, Emerson R M, Gillmore M R, Yamagishi T. 1983. The distribution of power in exchange networks: Theory and experimental evidence. *Am J Sociol* 89:275-305
- Cook K S, Hegtvædt K A. 1983. Distributive justice, equity, and equality. *Annual Review of Sociology* 9:217-241
- De Quervain D J F, Fischbacher U, Treyer V, Schelthammer M, et al. 2004. The neural basis of altruistic punishment. *Science* 305:1254-1258
- Diekmann A. 2004. The power of reciprocity - fairness, reciprocity, and stakes in variants of the dictator game. *Journal of Conflict Resolution* 48:487-505
- Durkheim E. 1938. *The rules of sociological method*. Chicago: University of Chicago Press
- Durkheim E. 1951. *Suicide: A study in sociology*. New York: The Free Press
- Emerson R M. 1962. Power-dependence relations. *Am J Sociol* 27:31-41
- Emerson R M. 1976. Social exchange theory. *Annual Review of Sociology* 2:335-362
- Falk A. 2004. Charitable giving as a gift exchange: Evidence from a field experiment (iza dp 1148). *Discussion paper 1148, Institute for the Study of Labor*
- Falk A, Fischbacher U. 2006. A theory of reciprocity. *Games and Economic Behavior* 54:293-315

- Fehr E, Falk A. 1999. Wage rigidity in a competitive incomplete contract market. *Journal of Political Economy* 107:106-134
- Fehr E, Fischbacher U. 2003. The nature of human altruism. *Nature* 425:785-791
- Fehr E, Fischbacher U. 2004. Third-party punishment and social norms. *Evolution and Human Behavior* 25:63-87
- Fehr E, Fischbacher U, Rosenblatt B, Schupp J, Wagner G. 2002. A nation-wide laboratory examining trust and trustworthiness by integrating behavioral experiments into representative surveys. *Schmollers Jahrbuch* 122:519-542
- Fehr E, Gächter S. 2000. Cooperation and punishment in public goods experiments. *American Economic Review* 90:980-994
- Fehr E, Kirchsteiger G, Riedl A. 1993. Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics* 108:437-459
- Fehr E, Klein A, Schmidt K M. 2007. Fairness and contract design. *Econometrica* 75:121-154
- Fehr E, Schmidt K M. 1999. A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics* 114:817-868
- Fischbacher U, Gächter S. 2006. Heterogeneous social preferences and the dynamics of free riding in public goods. *IEW Working Paper No. 261, University of Zurich*
- Fischbacher U, Gächter S, Fehr E. 2001. Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters* 71:397-404
- Gintis H. 1975. Welfare economics and individual development: A reply to talcott parsons. *Quarterly Journal of Economics* 89:291-302
- Gintis H. 2004. Modeling cooperation among self-interested agents - a critique. *Journal of Socio-Economics* 33:697-717
- Gintis H. in press. Towards a unified behavioral science. *Behavioral and Brain Sciences*
- Gintis H, Bowles S, Boyd R, Fehr E. 2003. Explaining altruistic behavior in humans. *Evolution and Human Behavior* 24:153-172
- Gneezy U, Leonard K L, List J A. 2006. Gender differences in competition: The role of socialization. *Manuscript, University of Chicago*
- Gouldner A. 1960. The norm of reciprocity. *American Sociological Review* 25:161-178
- Grusec J E, Kuczynski L. 1997. *Parenting and children's internalization of values: A handbook of contemporary theory*. New York: John Wiley & Sons
- Güth W, Tietz R. 1990. Ultimatum bargaining behavior - a survey and comparison of experimental results. *Journal of Economic Psychology* 11:417-449
- Hechter M, Kanazawa S. 1997. Sociological rational choice. *Annual Review of Sociology* 23:199-214
- Hedstrom P. 2005. *Dissecting the social: On the principles of analytical sociology*. Cambridge: Cambridge University Press
- Hedstrom P, Swedberg R, eds. 1998. *Social mechanisms: An analytical approach to social theory*. Cambridge UK: Cambridge University Press
- Henrich J, Boyd R, Bowles S, Camerer C, et al. 2001. In search of homo economicus: Behavioral experiments in 15 small-scale societies. *American Economic Review* 91:73-78
- Henrich J, Boyd R, Bowles S, Camerer C, et al. 2005. "Economic man" In cross-cultural perspective: Behavioral experiments in 15 small-scale societies. *Behavioral and Brain Sciences* 28:795-+
- Henrich J, McElreath R, Barr A, Ensminger J, et al. 2006. Costly punishment across human societies. *Science* 312:1767-1770
- Homans G C. 1958. Social behavior as exchange. *Am J Sociol* 63:597-606
- Horne C. 2004. Collective benefits, exchange interests, and norm enforcement. *Social Forces* 82:1037-1062

- Kollock P. 1994. The emergence of exchange structures. An experimental study of uncertainty, commitment and trust *Am J Sociol* 100:313-345
- Kollock P. 1998. Social dilemmas: The anatomy of cooperation. *Annual Review of Sociology* 24:183-214
- Lawler E, Yoon J. 1993. Power and the emergence of commitment behavior in negotiated exchange. *Am. Soc. Rev.* 58:465-481
- Lawler E, Yoon J. 1996. Commitment in exchange relations: Test of a theory of relational cohesion. *Am. Soc. Rev.* 61:89-108
- Ledyard J. 1995. Public goods: A survey of experimental research. In *Handbook of experimental economics*, ed. J. Kagel, A. Roth, pp. 111-194: Princeton University Press
- Macy M W, Willer R. 2002. From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology* 28:143-166
- Marwell G, Ames R E. 1979. Experiments on the provision of public-goods. Resources, interest, group-size, and the free-rider problem. *Am J Sociol* 84:1335-1360
- Molm L. 1997. *Coercive power in social exchange*. Cambridge UK: Cambridge University Press
- Niederle M, Vesterlund L. in press. Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics*
- Ostrom E, Walker J, Gardner R. 1992. Covenants with and without a sword - self-governance is possible. *American Political Science Review* 86:404-417
- Parsons T. 1937. *The structure of social action*. New York: McGraw-Hill
- Parsons T, Shils E. 1951. *Toward a general theory of action*. Cambridge, MA: Harvard University Press
- Rabin M. 1993. Incorporating fairness into game theory and economics. *American Economic Review* 83:1281-1302
- Raub W, Snijders C. 1997. Gains, losses, and cooperation in social dilemmas and collective action: The effects of risk preferences. *Journal of Mathematical Sociology* 22:263-302
- Simpson B. 2004. Social values, subjective transformations, and cooperation in social dilemmas. *Social Psychology Quarterly* 67:385-395
- Singer T, Seymour B, O'doherty J P, Stephan K E, Dolan R J, Frith C D. 2006. Empathic neural responses are modulated by the perceived fairness of others. *Nature* 439:466-469
- Stolte J F, Fine a G, Cook K S. 2001. Sociological miniaturism: Seeing the big through the small in social psychology. *Annual Review of Sociology* 27:387-413
- Van Lange P a M. 1999. The pursuit of joint outcomes and equality in outcomes: An integrative model of social value orientation. *Journal of Personality and Social Psychology* 77:337-349
- Van Lange P a M, Otten W, De Bruin E M N, Joireman J A. 1997. Development of prosocial, individualistic, and competitive orientations: Theory and preliminary evidence. *Journal of Personality and Social Psychology* 73:733-746
- Wrong D H. 1961. The oversocialized conception of man in modern sociology. *American Sociological Review* 26:183-193
- Yamagishi T. 1986. The provision of a sanctioning system as a public good. *Journal of Personality and Social Psychology* 51:110-116
- Yamagishi T, Nobuhito J, Kiyonari T. 1999. Bounded generalized reciprocity: Ingroup boasting and ingroup favoritism. In *Advances in group processes*, ed. S.R. Thye, E.J. Lawler, M.W. Macy, H.A. Walker, pp. 161-197. Stamford: JAI Press