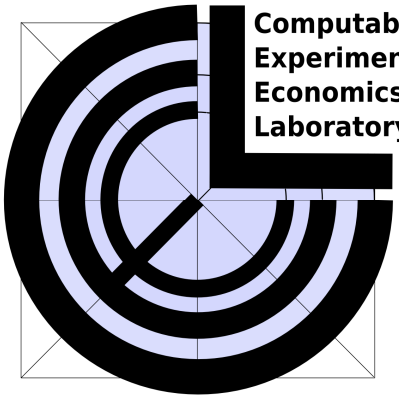

Francesco Guala and Luigi Mittone

Paradigmatic Experiments: the Dictator Game

CEEL Working Paper 7-08

Computable and Experimental Economics Laboratory

Via Inama, 5 38100 Trento, Italy

http://www-ceel.economia.unitn.it
tel. +39.461.882246

Paradigmatic Experiments: the Dictator Game*

Francesco Guala
(Department of Sociology & Philosophy, University of Exeter,
and CEEL, University of Trento)

Luigi Mittone
(Department of Economics and CEEL, University of Trento)

Abstract

Recent experiments with the Dictator Game (and the ensuing discussions) have been affected by considerable confusion regarding the purpose of this design. A common complaint is that the design gives rise to fragile regularities and therefore is of little use for theory-testing. We take issue with this view, and instead argue that the Dictator Game is potentially a very useful tool for experimental game theory, if properly used. It is particularly useful for investigating social norms, but economists have failed to take advantage of the Dictator Game because they still lack an adequate theory of norms.

1. Introduction

In the Postscript to the second edition of *The Structure of Scientific Revolutions* (1971), Thomas Kuhn highlights the role played by "exemplars" in defining and driving research within a scientific paradigm. Exemplars are paradigmatic examples of how "good science" ought to be done, and can take various forms: a theoretical model, a mathematical proof, a methodological device, an experimental set-up can all be "exemplars" in Kuhn's sense. A *paradigmatic experiment*, as the term will be used in this paper, is an exemplary experimental design in Kuhn's sense. Paradigmatic experiments are important at least for three reasons

* We thank participants to the Behavioral Economics workshop in Trento for their comments and suggestions. The usual disclaimers apply. Guala's work was supported by a research grant of the Computable and Experimental Economics Laboratory, University of Trento.

(Guala 2008): they have (1) a pedagogic function, by showing students how a good experiment is to be designed and run; (2) a reference function, by generating robust regularities from which new effects can be detected once some details of the experimental design have been varied; and (3) a sociological function, by setting standards that differentiate practice in one discipline from what is done in neighboring fields.

Studying paradigmatic experiments is thus particularly instructive, for what they can tell us about the research ethos and epistemic commitments of a scientific community. Experimental economics is no exception from this respect. Among widely replicated experiments in this field, the Double Oral Auction, Public Goods, and Ultimatum games are perhaps best known to economists and social scientists in general. Over the last decade, however, a new design – known as the "Dictator Game" (DG) – has attracted considerable attention. The DG is probably the simplest experimental setting one can think of. Indeed, the interaction between players in a DG is so minimal that one wonders whether the term "game" is appropriate at all: two players are to divide a sum of money provided by the experimenter (say, 10 euros, for simplicity). Only one player, however, can determine the size of the shares – she is the "dictator", and the other player can only accept the proposed division.

The DG is often presented as the latest addition to an already long list of "anomalous" games in the experimental literature (cf. Camerer and Thaler 1995). According to standard economic theory – based on the assumptions of rationality and self-interest – a dictator should keep 100% of the cake, and give nothing to the other player. As we shall see, this is not what happens in the laboratory, when subjects play the DG for real money. The rationality assumption is hardly questionable in a simple setting such as the DG: in the absence of major disturbances, it would be perverse to suggest that players do not understand the logic of this game. This leaves the selfishness hypothesis as the weak link in the standard economic model. Thus, unsurprisingly, much of the debate has focused on what the DG can teach us about human preferences and motives.

2. Altruism, fairness, and robustness

In a "standard" DG, only 40% of the experimental subjects playing the role of dictator keep the whole sum allocated by the experimenter. The majority of individuals prefer to give something to the other player, and the amounts offered (including zero-offers) average at

about 20% of the cake (Forsythe et al. 1994). These results seem to contradict the selfishness assumption, and are usually read as evidence of the importance of other-regarding preferences in economic behavior.[1] Two such preferences stand out as potentially relevant for the DG: dictators may be willing to give up part of their gains because they care about others' welfare, but also out of a concern for the equality of the resulting allocation. To distinguish, we shall label them the *altruism* and *fairness* hypotheses, respectively.

Under this interpretation, the DG may be a useful device to elicit and measure altruistic and fair preferences. That fairness and altruism are real phenomena is rather uncontroversial. Volunteering and donating to charities, for example, are important aspects of the non-profit sectors of all advanced economies, helping to solve market failures of various kinds and fulfilling other important social functions. What is controversial, or at least poorly understood, is the *extent* to which such fairness and altruism motives counteract selfish ones, and the *conditions* under which non-standard preferences may become important for predictive and policy purposes. The importance of the DG for experimental economics and economic theory would depend accordingly on the relative magnitude and robustness of the phenomenon that has been observed. As a consequence, much recent work has been concerned with testing the robustness of other-regarding preferences to changes in the details of the experimental design. This is a standard procedure across all the sciences, and should not be dismissed merely as a dogged attempt by dogmatic economists to get rid of an unwelcome result. By varying the initial conditions of a design, one can explore the range of circumstances in which a phenomenon can be observed, and better assess its importance for theoretical, practical, and engineering purposes.

Robustness testing is at the origins of some of the seminal and celebrated results of experimental economics. The convergence to efficient equilibria of markets with double oral auctions was demonstrated by Vernon Smith (1962) while testing the robustness of classroom experiments conducted by Edward Chamberlin at Harvard in the 1940s. Chamberlin (1948) observed that experimental buyers and sellers fail to discover the price of market clearing, and interpreted this result as a confirmation of his own theory of monopolistic competition.

[1] "Economic behavior" from now on is to be taken broadly as any kind of behavior having to do with the allocation of scarce resources, rather than narrowly as human behavior in the context of competitive markets. The broad definition is by far the most entrenched in both neoclassical and heterodox traditions.

Smith thought that Chamberlin's experiment had failed to give the standard theory a fair shot, and ran a replication changing a few key parameters of the experimental task. Crucially, he provided traders with a richer information environment by posting asks and bids on a public blackboard. Furhtermore, he gave them the opportunity to learn by running repetitions under the same market parameters. In this new environment, Smith discovered that markets do generate prices that approximate (with time) the theoretical equilibrium.

Notice that some of Smith's modifications of the Chamberlin design were motivated by the desire to instantiate in the laboratory the conditions that the theory deems necessary for market clearing. Perfect information regarding prices is clearly a case in point. But not all such conditions can or indeed ought to be instantiated in the lab: if markets cleared only when there is an infinite number of traders or perfectly smooth supply and demand curves, for example, the theory could not even be tested in the laboratory. Moreover, some of the other changes made by Smith were not suggested by the theory: the theory does not say for example that repetition and learning are important for the efficiency of markets.

It is worth keeping this in mind when assessing current experiments on the DG. Like Smith's seminal experiments, some of the recent literature seems to be driven by the desire to "give the standard theory a better shot". In a widely cited experiment, Hoffman et al (1996) report a significant decrease of donations in the DG. The decrease is obtained by imposing strict anonymity ("double-blind")[2] conditions and by decreasing the "social distance" between subjects and experimenter, as well as among the subjects themselves. In this environment, 60% of dictators decide to keep the whole amount for themselves, and the proportion of subjects donating more than thirty percent of the cake decreases from about 40 to less than 10% of the sample.

But DG behavior can be pushed even closer to the prediction of standard theory. Cherry et al. (2002) have observed that 95% of subjects donate nothing when a "legitimacy" factor is added to double-blind anonymity. Legitimacy over the assets to be shared is induced by

[2] In most economic experiments subjects do not know each other's identity but, potentially at least, their identity is known to experimenters ("single blind" design). When appropriate procedures are put in place to preserve anonymity also to experimenters, economists speak of a "double blind" design.

making the dictators earn money by answering the questions of a GMAT[3] quiz. In an inexact science like economics, a 5% deviation from theoretical predictions is an impressive result indeed.

Other studies however suggest that DG behavior can be pushed in the opposite direction too. Mittone and Ploner (2008) use an environment that is identical to the Cherry et al. (2002) design, except that the recipients are asked to exert the same effort as the dictators (by answering a quiz) while not being rewarded by the experimenter (i.e. the recipients' effort does not contribute to the size of the cake to be shared). In this case, "asset legitimacy" has a much weaker impact on the level of donations, because it is counter-balanced by equity of effort considerations. Up to 80% of dictators now are willing to give something, and the average level of donations is *tripled* compared to the treatment with asymmetric effort.

A similar phenomenon had already been highlighted by Ruffle (1998), who observed the effect of asymmetric effort in a game where only recipients were asked to contribute to the size of the cake by answering a quiz. In a significant number of cases (about 20% of the time) dictators went so far as to offer more than half of the sum that recipients had earned. This effect was absent in the baseline condition where the recipient's earnings were determined by a random device. As expected, exerting effort increased the average level of donations too.

Other studies report an increase in donations when the recipient is identified with a "reputable charity" like the Red Cross, compared to an anonymous subject. Even in a double-blind environment, the Red Cross attracted donations from over 73% of the dictators, as opposed to only 27% in the anonymous recipients condition (Eckel and Grossman 1996). The average level of contributions was tripled. Brañas Garza (2006) similarly observed that recipients from poor countries attract higher levels of donation than anonymous individuals drawn from the same pool of subjects as dictators, and that specifying the way in which money is going to be spent (e.g. to purchase medicines) further enhances this effect.

[3] Graduate Management Admission Test.

3. What is the DG for, after all?

We have gone to some length in surveying this literature, mainly to highlight the diversity of results obtained with the DG. This diversity is not unique to the DG – anthropologists for example have found quite a lot of variation across different cultures in the results of Ultimatum and Public Goods experiments (Henrich et al. 2004). But it is quite striking that so different results were observed in similar pools of subjects, in countries with similar social structures. As a consequence, it is not uncommon nowadays to hear or read complaints about the intrinsic instability of the DG. The implication is that instability implies *unreliability*, and hence that one should better abandon the DG as a research tool.

In the rest of this paper we want to put this implication under scrutiny, and try to disentangle the considerable confusion that surrounds current research on the DG. The purpose is not to challenge the conclusion that the DG is a poor experimental environment – it may indeed be the case, and anyway only time will tell whether important insights can be gained with this experimental design or not. Rather, we want to argue that learning from the DG is possible only conditional on a radical reorientation of economists' research goals and focus. We will claim the DG can be an interesting device for the student of *social norms*, and therefore can be useful to economists to the extent that social norms influence economic behavior both within and outside markets.

But we are running too far ahead here. First of all, it will be necessary to identify some of the reasons that cause economists to be suspicious about the DG. Economics has always been a distinctively theory-driven scientific discipline. Although experimental economics was partly developed as an antidote against this bias, it was nevertheless inevitably affected by theory-centrism. Many experimental economists, for example, still see their business as being mainly concerned with theory testing. Of course there is nothing wrong with testing theories: focusing on sharply defined predictions that have been rigorously derived from a set of transparent postulates is indeed an advantage that other neighbor disciplines (such as psychology) lack. And yet, theory-centrism carries the risk of interpreting every experimental result as either a falsification or a confirmation of an existing model, while neglecting its relevance for other issues regarding which no formal, general economic theory is currently available.

This is, in our view, happening with the DG. Most experimental work over the last decade seems to be concerned with answering one of the following questions: How seriously should one take the evidence of willingness to donate reported in early DG experiments? Is the phenomenon robust enough to warrant a revision of the standard theory? If so, can the evidence be explained by existing models of altruism or fairness?

The answer to the latter question is obviously not unrelated to the answer given to the first two questions. If the results of the DG are very volatile and sensitive to small variations in the experimental set-up, one may be tempted to conclude that the evidence is *both* too ephemeral to count as a falsification of the standard model, and too heterogeneous to be captured by models of altruistic or fairness preferences. Such models typically represent other-regarding motives as modifications of the self-interested utility function of standard economic theory. According to the self-interest model, each individual i acts so as to maximize her own utility $U_i(.)$, where the utility function varies over her own monetary payoffs π_i only. It is relatively simple to "complicate" the utility function by adding factors other than self-interest; a concern for altruism for example can be represented in the context of a two-person interaction as

$$U_i = \alpha\pi_i + \beta\pi_j, \text{ for } i \neq j,$$

where α and β are free parameters to be determined empirically and may vary from person to person. A preference for fairness – intended as equality – may take the following form:

$$U_i = \alpha\pi_i - \beta \mid \pi_i - \pi_j \mid.$$

In other words: an individual who is averse to inequality loses an amount of utility that is proportional to the difference between her monetary gains and those of another player. There are various models of this kind, differing in detail but sharing this basic modeling strategy (e.g. Fehr and Schmidt 1999, Bolton and Ockenfels 2000, Charness and Rabin 2002). The authors of such models are well aware that they are simplifying what is probably a very complex set of mechanisms, but nevertheless defend their models appealing to their simplicity and predictive power. Prediction, however, has proven to be their Achilles' heel. In spite of their valiant attempts, the identification of robust parameters that can predict data

across different environments has been problematic. The DG provides perhaps one of the most vivid examples of such difficulty: small variations in details that are not even represented in these models lead to huge variations in behavior. Overall, it seems unlikely that such models can explain what has been observed in DG experiments so far.

4. The origins of the Dictator's Game

Daniel Kahneman, Jack Knetsch and Richard Thaler (1986) are usually credited for running the first DG in experimental economics. Interestingly the DG is *not* presented in their paper as a new experimental setting worthy of independent interest – as a setting that is likely to generate robust phenomena for economic theorizing, in other words. Kahneman et al. use a "mini-DG"[4] as a *control treatment*, i.e. as a set-up that helps investigating *another* game, which is the main focus of their study. The DG is used in this seminal paper only as a means "to obtain an indication of the prevalence of unenforced fairness" in the Ultimatum Game (Kahneman et al. 1986: 106). In an Ultimatum Game a "proposer" offers a division of the cake that a "respondent" can accept or reject. When proposers offer more than a minimal amount, it is impossible to say whether this is because they care about the fairness of the allocation, or because they are playing strategically to anticipate possible rejections.[5] The observation of donations in a DG that mirrored closely the structure of the UG, but where the strategic element had been eliminated by design, enabled Kahneman et al. (1986) to conclude that a concern for fairness is likely to play some role also in the Ultimatum Game.

Subsequent replications used the DG in a similar fashion. Forsythe et al. (1994) again compared it with an Ultimatum Game, with the aim of investigating the effect of procedural considerations that go beyond concerns for inequality. The authors conclude that fairness cannot explain all the offers made in the Ultimatum Game. Notice that this methodology is entirely orthodox from a game-theoretic point of view. The theory of rational play at the core of game theory is an "if ... then ..." theory: it says that *if* their preferences and beliefs are so-and-so, *then* (rational) players will behave in such-and-such a way. But in principle the theory does not impose any restriction on the content of people's preferences. People may

[4] In the "mini-version" of a DG the Dictator can choose only among a limited number of allocations (usually two), rather than the full range as in the standard DG.

[5] In a standard Ultimatum Game played in Western countries, about 20% of the offers are rejected by Responders. Unfair offers, in particular, face a high chance of rejection.

well be entirely self-interested, or care for equality, as in the theories of fairness reviewed above. The standard models of economic theory thus combine two distinct assumptions (rationality and self-interest) that must be kept separate in the interpretation of experimental games.

Von Neumann and Morgenstern (1947) famously introduced a method for measuring utility independently of strategic situations. Once people's preferences have been observed using this method, one can proceed and check whether the theory of rational play issues correct predictions using the measurements as input data in the testing procedure. Consider the mini-Ultimatum Game on the right-hand side of Figure 1. One cannot test the rationality postulate of game theory in a game like this, without first gathering information about players' preferences. The fact that 94% of the B-players choose the Pareto-superior distribution ($750, $400) (see Charness and Rabin 2002) is not anomalous *per se*. It is perfectly possible that they care about others' welfare, and there is nothing irrational in that kind of preference (*de gustibus non est disputandum*). This behaviour is anomalous only in light of the choices that B-players make when offered a straight choice between the same allocations in the non-strategic decision on the left-hand side of Figure 1. Here subjects' concern for Pareto efficiency turns out to be significantly weaker. Intuitively, a feeling of gratitude or reciprocity shifts 25% of B-players towards ($750, $400) in the mini-Ultimatum Game. They choose this outcome in order to reward A-players who could have chosen ($750, $0) at no risk for themselves (but didn't).

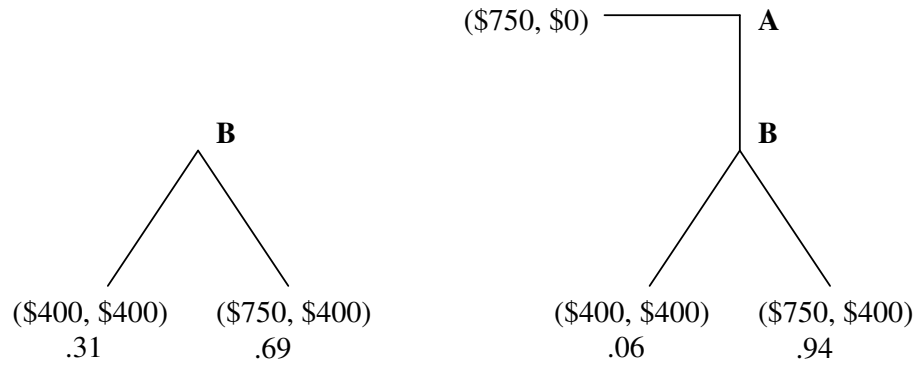

Figure 1: Mini-Dictator and mini-Ultimatum games, from Charness and Rabin (2002). The percentage figures under each pair of payoffs represent observed frequencies of choice in the experiments.

Reciprocity, to put it differently, violates consequentialism, a key postulate of rational choice, and thus *is* anomalous for standard game theory. But this is not the point we want to make in this paper.[6] Our main point is that the von Neumann-Morgenstern method involves essentially the use of mini-DG's. The DG is "kosher" game theory, when used this way. In this sense, it should be used *more*, not less often, because it is the best access we have to the observation of preferences over "pure" distributions that are undistorted by non-consequentialist (procedural) concerns.

5. Looking for norms

If our interpretation is correct, the DG was never supposed to be a design of independent interest, but rather a methodological tool for the investigation of non-standard preferences. It was not meant to become a "paradigmatic experiment" in the same sense as the Prisoner's Dilemma or the Ultimatum Game. And there are good reasons to believe that it *cannot* become an experimental set-up of that kind. In order to see why, it is necessary to ask what makes Ultimatum and Prisoner's Dilemma games so interesting and fruitful as loci of experimental investigation.

Prisoner's Dilemma and Ultimatum games allow the observation of *social norms* in action. Cooperation and fair splits are effects of our compliance with social norms that prescribe such behavior in settings of a given kind. What kind of settings fall under the domain of a norm is, as we shall see, an important empirical question. But the key point is that such games have the capacity to elicit normative behavior from players, even in the highly abstract and purified settings that experimental economists typically implement. As Ken Binmore (1998, 1999) has argued, subjects bring inside the laboratory a whole set of experiences and social cues that help them coping with what is an otherwise unfamiliar and puzzling situation. Anthropologists who have played such games with hunter-gatherers in Africa, Asia, and South America (Henrich et al. 2004) largely confirm this insight. The Orma, a group of pastoral-nomadic people living in Kenia, for instance do not fail to notice that the Public Goods game is structurally similar to the *harambee* – an institution used for fund-raising and other collective projects – and behave accordingly (Ensminger 2004).

[6] See Guala (2006) for a full analysis of the significance of these designs for experimental game theory.

Public Goods and Ultimatum Games elicit norms of this kind in virtue of their *structure*, and of the fact that games with such a structure are relatively *common* in many societies.[7] The same, however, cannot be said of the DG. The DG has a remarkably simple structure – indeed too simple (we should perhaps say "poor") to elicit a specific norm. It is an unusual situation too, for in real life one rarely deals with "windfall money" to be shared with an anonymous stranger. It is important to stress both aspects: structure and familiarity. Subjects have to deal with windfall money and anonymous partners in Ultimatum Games too, but the structure of the UG is rich enough to focus their attention on power asymmetries, and thus elicit the fairness norms that apply in such circumstances. The DG in contrast is too "thin" for that, and experimental subjects are left to puzzle over which behavior is deemed appropriate for a situation of that kind.[8]

That's probably why behavior in the DG is so volatile. In the absence of a norm, small changes in the experimental design can prompt major variations in observed behavior. If, to paraphrase John Ledyard (1995), experimenting with Public Goods games is like doing free-fall experiments with a table-tennis ball, in a DG one is experimenting with a soap bubble that may be blown away by the smallest whiff, or even burst. Out of metaphor, subjects in the DG are keen to exploit even the slightest cues that may help them identifying a norm of conduct. Consider the Cherry et al. (2002) experiment reviewed in section 2. There is a strong norm in our society that an individual should be entitled to consume what she has produced. At the same time, there is no obligation to share the fruits of one's labor with idle individuals. It is unsurprising then that in the experiment of Cherry and colleagues 95% of dictators decide to keep the whole sum for themselves. But in commenting on this result,

[7] Which does not mean that the norms that apply to such contexts are the same in all cultures. As the studies in Henrich et al. (2004) demonstrate, there can be significant variation in behavior depending on local institutions and norms.

[8] A survey conducted by Cristina Bicchieri and Jason Dana largely confirms this point. In a questionnaire they asked a sample of college students what normative expectations they associate with a DG. About half of the respondents (56%) answered that no outcome can be said to be unfair in such a game. When explicitly asked what a *fair* outcome would be, 68% indicated the equal split, but 21% of the sample thought that keeping the whole sum was fair. Bicchieri (2006: 126) rightly concludes that these answers reveal a great deal of uncertainty and confusion about what is appropriate in the DG. In our society there just isn't a social norm that unambiguously identifies what ought to be done in circumstances of this kind. It would be interesting to conduct the same survey in the context of an UG, and to compare results.

Cherry and his coauthors miss the most important and interesting point: the reason why the data fit the prediction of the selfish rationality model is that their design has triggered a powerful normative mechanism that invites people to behave in a self-interested manner.

In experiments that introduce information about effort and reward of the recipient (Ruffle 1998, Mittone and Ploner 2008) the level of sharing is increased, because different rules of conduct are triggered. In richer environments people have to deal with different normative cues, and their behavior is going to reflect a trade-off between norms pushing in different directions (as well as the very human tendency towards selfishness that is always at work, to some extent). Thus a "pure" DG can be mostly useful as a benchmark, a contrast case to measure the relative effect of introducing such cues. It is implausible to expect that it can elicit a strong effect that is robust across a variety of environments.

So why have economists focused on testing theories of non-standard preferences instead? Partly, because they presently lack a theory of norms. And although norms are at the centre of much research in the social sciences, neighbor disciplines like sociology have consistently failed to provide simple but general formal models that can be tested in the laboratory. This, in our view, goes a long way towards explaining the misunderstandings and unrealistic expectations that have confused the debate on the DG. The theory-centered approach prompted many experimental economists to look at the DG as a source of data for theory-testing and theory-building. But in doing so they were looking at the *wrong theories*.

6. Theories of norms

There currently are three major theoretical approaches to the study of norms: the first one represents norms as unconditional rules of behaviour; the second one as rules that are conditional on a system of expectations; and the third one as joint commitments produced by team thinking. In this section we review them briefly and assess their explanatory potential.

(1) *Unconditional rules*: the most popular approach in economics and the social sciences is to model social norms as Kantian imperatives of the form "in C, do S", where C is a set of relevant conditions and S is an available strategy (in game-theoretic sense). Thus, for example, a norm of cooperation would prescribe to cooperate unconditionally in every situation that has the structure of a Prisoner's Dilemma or a Public Goods game. This concept

of norm has the advantage of being easy to model in mathematical terms: one may represent individuals as "dumb" agents who mechanically play the same strategy (as in evolutionary game theory), or alternatively add an other-regarding element in their utility function (as in models of fair preferences). The problem, however, is that all models of this kind are "born refuted": we know for example that unconditional cooperation explains only a minority data in Public Goods games (Burlando and Guala 2005, Gachter and Thoni 2005), and that subjects behave quite differently depending on whether the other players expect them to be fair or not. In games where the payoffs are expressed in "chips" with different monetary values to each player, for example, subjects are quite happy to simply *pretend* to be fair (by sharing the chips, but not the money, equally) if the other players do not know how much the chips are worth (Kagel et al. 1996).

(2) *Conditional rules*: the main alternative to the Kantian view is to model norms as conditional rules of behavior, of the form "in C, do S if X is the case". Theorists disagree regarding the nature of X: according to an approach that has gained some prominence in economics, individuals condition their behavior on the *intentions* of other players. Thus, for example, they are ready to cooperate if others also cooperate in PD games (respond "nice" to "nice"), but reject unequal offers in the UG (respond "nasty" to "nasty"). The main advantage of this approach is that it is consistent with a large body of literature emphasizing the importance of reciprocity both in biological and in cultural evolution (e.g. Trivers 1971, Gintis et al. 2005). Its main shortcoming is that "psychological games" with players who care about intentions are difficult to model and have a large number of equilibria (Dufwenberg and Kirchsteiger 2003, Falk and Fischbacher 2006). Alternatively, some theorists (Pettit 1990, Bicchieri 2006) interpret X as sets of expectations: individuals condition their behavior on other players' beliefs concerning what they will or ought to do in situation C. While under the first interpretation pro-social behavior in, say, Public Goods games is prompted by *unconditional norms of conditional cooperation*, under this interpretation individuals are influenced by *conditional norms of unconditional cooperation*. The key difference between the two theories is in the role played by expectations.

(3) *Team reasoning*: finally, a substantial body of literature has emerged in psychology, economics, and philosophy, that emphasizes the importance of group identity for social behavior and the emergence of institutions (e.g. Tajfel and Turner 1986, Sugden 2000, Bacharach 2006, Gilbert 1989, Searle 1990). Institutions such as conventions and norms,

according to this approach, result from a *joint commitment* by the members of a group to pursue collective rather than individual goals. Such commitments are normative roughly in the sense that a social contract is: they are mutually binding and cannot be breached unilaterally by individual members of the group. The main difference between this approach and the previous alternatives is that no individual reasoning is supposed to take place (of the form "If the other player does/believes X, I prefer to do S", etc.) and action is explained instead by desires and beliefs in "we-mode" ("*we* ought to do S in C").

7. Thick and thin designs

Research on norms – what they are and how they function – is only beginning to take off. To the extent that social norms influence economic behavior both within and outside market exchange, economists obviously have a stake in this research project. Experimental economists, moreover, have provided some important methodological tools that are now widely used across the human sciences – from social psychology, to anthropology, and evolutionary game theory. The final point we would like to make in this paper is that experimenters must now abandon some entrenched methodological habits and learn to adapt their experimental designs to the task at hand. If the key to explaining cooperation and fair allocation of resources lies in understanding social norms, it will be necessary to tailor the experiments accordingly. The days of testing the assumption of selfish rationality are gone. We now know that behavior can both diverge and conform to the standard model, and we know that in different situations experimental subjects can be pushed towards one or the other extreme of the self-interested/other-regarding spectrum.

The classic precepts of experimental economics (Smith 1976, 1982), if strictly implemented, can be used to induce behavior that is very close to the predictions of the standard model. This is not surprising, because the precepts were first devised in the context of market experiments aimed at testing the effects of different institutional rules governing the exchange and allocation of goods. It was very important in the context of such experiments to achieve complete control on individual preference profiles, and abstract environments with strict anonymity conditions help a great deal in this respect. Over time, however, the precepts have been extended to decision-making and game theory experiments, where they are not

always suited to the research goals.[9] In running such games experimenters are often interested in *investigating* the motivations that prompt subjects' behavior, rather than in *controlling* such motives as in market experiments.

Of course control is an important scientific goal in itself, and there may be situations in which economists are interested in finding ways to constrain agents' behavior in such a way as to make it conform more closely to the axioms of economic theory. (Consider an oligopolistic market, for instance: there may be very good reasons to make sure that certain forms of communication – and hence collusion – are forbidden in running an auction.) But it is important to keep in mind that the constraints are not going to be neutral with respect to the agents' interpretation of the game they are playing. Self-interested and competitive behavior, however common and perhaps even innate for humans, is regulated by a network of norms that can be triggered by various contextual cues. In the Hoffman et al. (1996) experiment reviewed in section 2, for instance, the double-anonymity design is likely to act like a trigger of this kind. Stressing in the experimental instructions that great care has been taken to make sure that absolutely *nobody* (not even the experimenter) will ever know what each subject has done, is not unlikely to point dictators towards one particular kind of behavior.

In this respect we find it useful to repeat George Loewenstein's (1999) reminder that every experimental context is a social context of some kind: even the most "purified" design, where every element of sociality has been removed by the experimenter, must be interpreted by the experimental subjects (as a game, as a competition, or simply as an experiment). And indeed, as we have seen, by removing *every* cue that may guide the players in forming normative expectations, we may end up with extremely fragile results that have little ecological validity.

Another important, related issue concerns the degree of ecology of experimental designs. When a social (norm) driven behavior is analyzed is important to control how effective the experimental design is in activating the proper social norm. This is not only a matter of choosing the right payoffs or the right strategic interaction frame but is also a problem of choosing the correct stimuli. Economic experiments often use very artificial settings like presenting a matrix of numbers/payoffs. The claim is that such experimental designs are "clean": in other words, any uncontrollable factor induced by a specific set of stimuli – e.g.

[9] This point is made in more detail in Guala (2005, Ch. 11); see also Cubitt et al. (2001).

the effects produced by a specific scenario – is eliminated through the adoption of a totally artificial setting. There is extensive evidence however that stimuli of this kind activate specific areas of the brain which play only a limited role in social reasoning (see e.g. Sanfey et al. 2003, Camerer et al. 2005). Imagine that one is interested in analyzing the activation of a social norm which typically drives cooperation in small group interaction. In this kind of situations human beings are driven by a mix of selfish driven calculus and emotionally driven reactions. Designing the problem as if it were only a matter of computing the "best" payoff in a game theory setting can minimize the emotional part and thus generate a poor ecological result. Such designs are usually defended by pointing out that even in anodyne settings the effect of norms and emotions is strong enough to be detected by statistical means – the implication being that their effects must be even stronger in settings that are not biased towards triggering subjects' "calculative" attitudes. Once the very existence of norms is beyond dispute, however, this argument loses much of its appeal. Most experimental designs nowadays are devoted to investigating the factors that enhance or hinder the influence of norms. There is no guarantee that such factors exert the same effect in highly abstract as in socially richer designs. For this reason it is important that economists introduce more variation in designing the stimuli used in the experiments.

We do *not* find it useful, then, to insist that experiments be *always* performed in standardized environments with simple tasks, high monetary incentives, and repetition of the task. This is *not* equivalent "to use clean test tubes in chemistry experiments" (Binmore 1999: F17) – not, at least, if we are interested in studying norms and their effect on economic behavior. It is particularly surprising that such a piece of advice should come from an economist who has devoted much time and effort to studying social norms (Binmore 1998). But again, the main focus of Binmore's research has been standard economic theory, and how to use it to explain normative behavior. It must be natural for him then to conceive of experimental economics as primarily aimed at theory testing rather than at the direct observation of the causes and effects of norms.

8. Conclusion

We hope that these remarks will help economists to appreciate the virtues and limitations of the DG in experimental research. We believe that our suggestions transcend the divide between behavioral and experimental, as well as heterodox and neoclassical economics. We

have argued that much confusion and misunderstanding would be avoided by returning to the original methodology advocated by von Neumann and Morgenstern: DG's are particularly useful as tools for the observation of preferences and the estimation of utility functions in simple non-strategic contexts *before* we engage in testing the theory of rational play that is at the core of game theory. Moreover, the DG is likely to be useful as in studying the effect of contextual cues that determine the framing and triggering of social norms. As such, it will continue to be used by those economists who are willing to re-orient their research, away from the testing of preference theories, and towards theories of social norms.

References:

Bacharach, M. (2006) *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton: Princeton University Press.

Bicchieri, C. (2006) *The Grammar of Society*. New York: Cambridge University Press.

Binmore, K. (1998) *Game Theory and the Social Contract II: Just Playing*. Cambridge, Mass.: MIT Press.

Binmore, K. (1999) "Why Experiment in Economics?", *Economic Journal* 109: F16-24.

Bolton, G. and Ockenfels, A. (2000,) "ERC: A Theory of Equity, Reciprocity and Cooperation", *American Economic Review* 90: 166-93.

Brañas Garza, P. (2006) "Poverty in Dictator Games: Awakening Solidarity", *Journal of Economic Behavior and Organization* 60: 306-20.

Burlando, R. and Guala F. (2005) "Heterogeneous Agents in Public Goods Experiments", *Experimental Economics* 8: 35-54.

Camerer, C., Loewenstein, G. and Prelec, D. (2005) "Neuroeconomics: How Neuroscience Can Inform Economics", *Journal of Economic Literature* 43: 9-64.

Camerer, C. and Thaler, R. (1995) "Ultimatums, Dictators, and Manners", *Journal of Economic Perspectives* 9: 209-19.

Chamberlin E.H. (1948) "An Experimental Imperfect market", *Journal of Political Economy* 56: 95-108.

Charness, G. and Rabin, M. (2002) "Understanding Social Preferences with Simple Tests", *Quarterly Journal of Economics* 117: 817-69.

Cherry, T.L., Frykblom, P., and Shogren, J.F. (2002) "Hardnose the Dictator", *American Economic Review* 92: 1218-21.

Cubitt, R.P., Starmer, C. and Sugden, R. (2001) "Discovered Preferences and the Experimental Evidence of Violations of Expected Utility Theory", *Journal of Economic Methodology* 8: 385-414.

Dufwenberg, M. and Kirchsteiger, G. (2003) "A Theory of Sequential Reciprocity," *Games and Economic Behavior* 47: 268-98.

Eckel, C. and Grossman, P.J. (1996) "Altruism in Anonymous Dictator Games", *Games and Economic Behavior* 16: 181-91.

Ensminger, J. (2004), "Market Integration and Fairness: Evidence from Ultimatum, Dictator, and Public Goods Experiments in East Africa", in Henrich et al. (2004), 356-81.

Falk, A. and Fischbacher, U. (2006) "A Theory of Reciprocity", *Games and Economic Behavior* 54: 293-315.

Fehr, E. and Schmidt, K. (1999), "A Theory of Fairness, Competition and Co-operation", *Quarterly Journal of Economics* 114: 817-68.

Forsythe, R., Horowitz, J.L., Savin, N.E., and Sefton, M. (1994) "Fairness in Simple Bargaining Experiments", *Games and Economic Behavior* 6: 347-369.

Gachter, S. and Thoni, C. (2005) "Social Learning and Voluntary Cooperation among Like-Minded People", *Journal of the European Economic Association* 3: 303-314.

Gilbert, M. (1989) *On Social Facts.* London: Routledge.

Gintis, H., Boyd, R., Bowles, S. and Fehr, E. (eds. 2005) *Moral Sentiments and Material Interests*. Cambridge, Mass.: MIT Press.

Guala, F. (2005) *The Methodology of Experimental Economics*. New York: Cambridge University Press.

Guala, F. (2006) "Has Game Theory Been Refuted?", *Journal of Philosophy* 103: 239-63.

Guala, F. (2008) "Paradigmatic Experiments: The Ultimatum Game from Testing to Measurement Device", *Philosophy of Science*, forthcoming.

Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E. and Gintis, H. (eds. 2004), *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies.* Oxford: Oxford University Press.

Hoffman, E., McCabe, K. and Smith, V.L. (1996) "Social Distance and Other-Regarding Behavior in Dictator Games", *American Economic Review* 86: 653-60.

Kagel, J., Kim, C. and Moser, D. (1996) "Fairness inUltimatum Games with Asymmetric Information and Asymmetric Payoffs", *Games and Economic Behavior* 13: 100-110.

Kahneman, D., Knetsch, J. and Thaler, R. (1986) "Fairness and the Assumptions of Economics", *Journal of Business* 59: S285-300.

Kuhn, T.S. (1971) *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press (2nd edition).

Loewenstein, G. (1999) "Experimental Economics from the Vantage Point of Behavioral Economics", *Economic Journal* 109: F25-34.

Mittone, L. and Ploner, M. (2008) "Asset Legitimacy and Perceived Equity in the Dictator Game", CEEL Working Paper, University of Trento.

Pettit, P. (1990) "*Virtus Normativa*: Rational Choice Perspectives", *Ethics* 100: 725-55.

Ruffle, B.J. (1998) "More Is Better, but Fair is Fir: Tipping in Dictator and Ultimatum Games", *Games and Economic Behavior* 23: 247-65.

Sanfey, A., Rilling, J., Aaronson, J., Nystrom, L., and Cohen, J. (2003) "The Neural Basis of Economic Decision-Making in the Ultimatum Game", *Science* 300: 1755–58.

Searle, J. (1990) "Collective Intentions and Actions". In *Intentions in Communication*, edited by P. Cohen, J. Morgan, M.E. Pollack. Cambridge, Mass.: MIT Press.

Smith, V.L. (1962) An experimental study of competitive market behavior, *Journal of Political Economy* 70: 111-137.

Smith, V. (1976) "Experimental Economics: Induced Value Theory", *American Economic Review* 66: 274-7.

Smith, V. (1982) "Microeconomic Systems as an Experimental Science", *American Economic Review* 72: 923-55.

Sugden, R. (2000) "Team Preferences", *Economics and Philosophy* 16: 174-204.

Tajfel, H. and Turner, J.C. (1986) "The Social Identity Theory of Inter-Group Behavior", in S. Worchel and L. W. Austin (eds.), *Psychology of Intergroup Relations*. Chicago: Nelson-Hall.

Trivers, R.L. (1971) "The Evolution of Reciprocal Altruism", *Quarterly Review of Biology* 46: 35-57.

Von Neumann, J. and Morgenstern, O. (1947) *The Theory of Games and Economic Behavior*. New York: Wiley (2nd edition).