

Jamison, Julian C.

Working Paper

Games with synergistic preferences

Working Papers, No. 11-15

Provided in Cooperation with:

Federal Reserve Bank of Boston

Suggested Citation: Jamison, Julian C. (2011) : Games with synergistic preferences, Working Papers, No. 11-15, Federal Reserve Bank of Boston, Boston, MA

This Version is available at:

<https://hdl.handle.net/10419/55552>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Games with Synergistic Preferences

Julian C. Jamison

Abstract:

In economic situations a player often has preferences regarding not only his or her own outcome but also regarding what happens to fellow players, concerns that are entirely apart from any strategic considerations. While this can be modeled directly by simply writing down a player's final preferences, these are commonly unknown *a priori*. In many cases it is therefore both helpful and instructive to explicitly model these interactions. This paper, building on a model due to Bergstrom (1989, 1999), presents a simple structure in the context of game theory that incorporates the “synergies” between players. It is powerful enough to cover a wide range of such interactions and model many disparate experimental and empirical results, yet it is straightforward enough to be used in many applied situations where altruism, or a baser motive, is implied.

Keywords: altruism, interdependent preferences, fairness, cooperation

JEL Classifications: C72, D03, D64

Julian C. Jamison is a senior economist in the Research Center for Behavioral Economics at the Federal Reserve Bank of Boston. His e-mail address is julian.jamison@bos.frb.org.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at <http://www.bostonfed.org/economic/wp/index.htm>.

The author would like to thank Ted Bergstrom, Lynn Connell-Price, Elizabeth Murry, Lones Smith, and an audience at the Massachusetts Institute of Technology for useful comments and feedback. He apologizes for any remaining errors.

The views expressed in this paper are those of the author and do not necessarily represent those of the Federal Reserve Bank of Boston or the Federal Reserve System.

This version: July 22, 2011

1. Introduction

Frank (1987) states: “Our utility-maximization framework has proven its usefulness for understanding and predicting human behavior. With more careful attention to the specification of the utility function, the territory to which this model applies can be greatly expanded.” This is a particularly germane observation with respect to game theorists, who simply tend to assume that they are given the full and correct final preferences of players in a game—and that their object is to analyze the resulting strategic interactions. Where these preferences come from, and especially what differences might arise between the payoff to an individual and his or her ultimate preference over outcomes, has generally not been considered to be within the purview of game theory. However, as Frank points out, this necessarily limits game theory’s scope. For instance, it is probably not an exaggeration to say that all game theorists feel that no rational player should ever knowingly play a strictly dominated strategy. And yet this is exactly what robustly occurs in the one-shot Prisoner’s Dilemma. The fault lies not with the theory, but with the inattention as to its application.

This paper attempts to provide a general and formal theoretical link between the base payoffs in a game and the resulting final utilities or preferences. The discrepancy is due to the fact that players care about the utilities of the other players in the game, for example, due to altruism. The main reason to formalize this link is to provide applied and experimental economists with a model for this pervasive interaction, so they are not forced to come up with new (and ad hoc) formulations every time it is relevant. A second reason lies in the theorists’ stock-in-trade: to gain a better understanding of the process.

The jump from payoffs to the calculation of final utilities goes on all the time in almost all games, so we should have a model (or, better yet, several competing models) of how it happens and what it implies.

We introduce a general definition of games with synergistic utility. Synergistic utility functions capture the idea that utility increases in one's own payoff, and may increase or decrease in others' utilities. Sufficient technical conditions are imposed for the concept to be well-defined, but otherwise the formulation is general enough to allow maximal variety in specific applications. All players are fully rational (including being expected-utility maximizers) and no new equilibrium concepts are introduced. A specific example, the linear synergistic utility function, is introduced and analyzed in greater detail. Several applications of the theory are given, including: cooperation in the Prisoner's Dilemma, overproduction in Cournot oligopoly, extended play in the centipede game, and interior solutions in the dictator game.

The paper proceeds to section 2, in which some of the related literature, both applied and theoretical, is discussed and compared with the synergistic utility concept. In section 3, the formal model, including the central definition, is given. Next, section 4 illustrates the theory with examples both of different synergistic utility functions and of their application to different games of interest. Section 5 addresses several topics from game theory, such as incomplete information, in the context of synergistic games. Finally, section 6 briefly concludes.

2. Literature

The literature relating to altruism and interdependent preferences is wide and diverse, with each paper seemingly taking its own course depending on the topic at hand and the immediate goals of the authors. The first broad category can be construed as considering the various applications of altruistic-like tendencies in specific situations. In the overlapping generations (OLG) macroeconomics literature, this includes the famous paper by Barro (1974) on Ricardian equivalence, and the subsequent paper by Kotlikoff, Razin, and Rosenthal (1990) which disputes the finding. The models in these papers have “dynasties” in which the current generation cares about their descendants’ consumption as well as their own. Bisin and Verdier (1998, 2001) study the Prisoner’s Dilemma in the context of cultural transmission, modeling altruism with the addition of a positive constant. All of these papers model altruism in one direction only, that is, there is no feedback effect between the players. In labor economics, Rotemberg (1994) studies workplace relations. He determines under what conditions cooperation can be obtained and when this benefits the employer, but defines altruism only insofar as an employee’s utility is the sum of payoffs to the group. Rotemberg states that, “Cooperative outcomes for *either* individual in the Prisoner’s Dilemma obtain only when *both* individuals feel altruistic toward each other.” As we shall see, this claim contradicts the conclusions of a synergistic utility model, in which an altruistic player may desire to cooperate even when facing a nonaltruistic opponent.

Altruism within the family has been studied since Becker (1974) and his Rotten Kid Theorem. He models interdependent utilities using a basic additive form. Bruce and Waldman (1990) compare this line of work to the Samaritan’s Dilemma and Barro-Ricardian equivalence in a similar framework. Other work applying some degree of

altruism includes Coate (1995), who studies insurance with rich and poor agents, and Collard (1975) working in a general equilibrium framework. Note that these studies represent only a small sample of the work that employs altruism or interrelated utilities in some form or other. In addition to the various subfields of economics already mentioned, these types of models have been used in disciplines ranging from law to philosophy to political science.

The second general class of papers, on evolution and biology, are also closely tied to the theoretical psychology literature. Frank (1987) belongs to this vein when he studies the commitment problem. He finds that if one can choose to be a guilty type (perhaps through an evolutionary process) and show it, one can commit credibly. This can be of great benefit, for instance in the provision of public goods. Wolpert *et al.* (2010) develop a general theoretical model along similar lines. Bergstrom (1995) studies genetically predetermined behaviors, which is to say there is no free choice on the part of the players. He finds that cooperation in the Prisoner's Dilemma can be a stable outcome when players have preferences that take into account the payoffs (not the utility) of others and that genetic propagation occurs through imitation of successful strategies. Once again, these are only a sample of the papers that consider this sort of evolutionary fitness paradigm. All are distinguished from the present work in that this paper concerns rational players in a nondynamic setting, but it is interesting to note that some of the conclusions reached are similar.

A third category consists of the relatively large number of experimental economics papers which have looked at a number of different games and found results that diverge from those predicted by the basic equilibrium concepts. Dawes and Thaler

(1988) study experiments with public goods, ultimatum games, and the Prisoner's Dilemma. They discuss altruism as a general explanation but do not suggest a model. Palfrey and Rosenthal (1988) also study public goods provision where altruism consists of a single lump-sum addition to payoffs (from "doing the right thing") when a player contributes. Cooper et al (1992) consider altruism in the setting of cheap talk and coordination games. One of the complications that arises from explaining the data in these and other games in this way is that it requires not only positive emotional interactions, such as altruism, but also negative interactions, such as spite (or at least retribution). For instance, it is otherwise impossible to rationalize rejected offers in the ultimatum game. Levine (1998) creates a relatively simple theory with utility linear in one's own and one's opponent's payoffs (with a possibly negative weight on the opponent). The parameters of his model are pinned down by matching data on ultimatum and centipede games. Levine then tests the model, with some success, on public goods games and on market games. The two main distinctions between his theory and the synergistic utility theory are that Levine's players care about the payoffs rather than the utilities of their opponents and his inclusion of a reciprocity factor, so that how much a player cares about others depends on how much they care about him. It turns out that a lot of the observed behavior can be explained without introducing this additional slight complexity, and that synergistic utilities can also rationalize some behavior (for example, in the dictator game) that Levine's model, as it stands, cannot. Charness and Haruvy (2002) experimentally test several models of altruism within a single framework.

The final group of related papers are those from the game theory literature. Geanakoplos, Pearce, and Stacchetti (1989) introduce the concept of psychological games

(and psychological equilibrium) in which utility is a function not only of actions but also of beliefs over actions. Among other things, this concept allows utility to depend on reactions of pleasure or anger, although only with respect to expected actions in a particular game. Players do not explicitly care about the welfare of their opponents, although in theory it can always be incorporated into their preferences. This is an extremely powerful and all-encompassing structure, but this breadth precludes finding a common thread from which to deduce or to explain the results observed across a variety of different games. Rabin (1993) instantiates this idea by introducing a *fairness equilibrium*, a more inherent concept which begins with a *kindness function* between the two players. Because of the special nature of the equilibrium concept, Rabin's results depend on the absolute level of the base payoffs and apply only to two-person games. Nevertheless, he is able to draw several fairly general conclusions. Sally (2001) takes a similar but somewhat more extended approach, building on the "psychological distance" between players. He develops the *sympathetic equilibrium* concept, and finds that it is sometimes possible to choose cooperation in the one-shot Prisoner's Dilemma. As in Rabin's paper, reciprocity is the starting point and again, essentially because of reciprocity, it is unclear how to extend the results to more than two players. Returning to the traditional equilibrium concepts, the approach used in Bergstrom (1989,1999) is perhaps closest to the present paper. He presents a general model in which a player's utility is an increasing transformation of the given player's payoff and the other players' utilities. One distinction with the synergistic utility concept is that Bergstrom takes a fixed-point rather than a limit-point approach. He is able to explain cooperation in the Prisoner's Dilemma, although this approach does lead to some rather counterintuitive

conclusions in other situations. Note that this paper is not pursuing a behavioral approach, since all players in the synergistic model are fully rational with standard preferences over final utility (and we use standard equilibrium concepts), but it does explore some of the same questions.

3. Model

One way to introduce an altruism-like aspect in a formal game-theoretic model is to add a positive constant to payoffs following a “good” action, such as contributing in a public goods game or cooperating in the Prisoner’s Dilemma. This is plausible in some circumstances, but does not capture the positive or negative benefits that a player may receive depending on the welfare of his or her opponents.¹ These benefits can be captured most simply by adding a proportion of the opponents’ payoffs to that of the player in question. This approach, however, has an inherent inconsistency: for instance, if the benefit arises not just from doing good, but instead from being glad that a fellow player is happy, then it should be the other player’s utility and not payoff that matters.² That is, rational players will be farsighted and will think through more than one step of the process. In general, then, final utilities will be a function of one’s own payoff and of the final utilities of the other players.

It is not unreasonable to ask why utilities should not be a function of own *utility* and others’ utilities. The short answer is that this too is inconsistent: preferences are utilities, they are not defined subsequent to utilities. As an example, consider an altruistic

¹ Throughout the paper “opponent” will be used interchangeably with “other player,” whether or not the particular relationship happens to be adversarial.

² One caveat is that this may not apply as fully in a corporate setting, where firms are likely to care about the profits of other firms and not the indirect feelings caused by profits.

player with an indifferent (entirely self-concerned) opponent. The opponent will necessarily always have final utility equal to base payoff. If the altruist has utility equal to a weighted average between own payoff and the other's utility, her final utility will lie somewhere in between the two original payoffs. If, however, her utility is a weighted average between own *utility* and the other's utility, her final utility must equal that of her opponent no matter what her original payoff. In fact, under these assumptions it is not uncommon that the final utilities of both players will depend only on their altruism types and will be wholly independent of their original payoffs, an undesirable feature.³

Before proceeding to the formal model, the interpretation of the base payoffs needs to be clarified. The base payoffs are already objects in utility space, so these should not be thought of as monetary payoffs or profits. Rather, the base payoffs can be considered as the utility resulting from the outcome if it were in a one-person setting, or in a setting where the effects of the outcome on other players are unknown. Alternately, the base payoffs are the utilities of thoughtless players, to whom it has not yet occurred that there are other players and/or what implications that consideration might entail. We assume, as ever, that players already include in their base preferences any positive feelings that arise from simply doing good or being fair, or on the flip side any negative feelings directly arising from an act of, say, betrayal. What they do not include are preference changes due to the realized utility of one's opponents in a particular outcome of the game.⁴

We are given a game \mathbf{G} with I players and payoffs v_i . A **synergism type** for a player i is an element θ_i drawn from a type-space Θ . Denote the vector of synergism

³ The author has worked considerably with this alternate model and is more than willing to share the results of these pursuits with anyone who is interested.

⁴ Note that we are assuming, as we must, the possibility of interpersonal comparison of utility.

types for the I players by Θ . Let f be a real-valued function taking as arguments I real numbers (interpreted as welfare measures for oneself and one's opponents, respectively) and as parameters the elements of Θ . Hence $f : \mathbf{R}^I \times \Theta \rightarrow \mathbf{R}$. So f is the same for all players, but each has a separate synergism type. The base payoff for player i is $v_i = u_i^0$. Following the motivation above, we define $u_i^1 = f(v_i, v_{-i}; \theta_i) = f(v_i, u_{-i}^0; \theta_i)$ and in general $u_i^n = f(v_i, u_{-i}^{n-1}; \theta_i)$. At each round, players recalculate their opponents' utility levels and then adjust their view of their own utility in response, continuing ad infinitum. Finally, let $u_i^\infty(v_i, v_{-i}; \theta_i) = \lim_{n \rightarrow \infty} u_i^n$. Of course this limit may not always exist.

Definition: Given Θ , a function $f : \mathbf{R}^I \times \Theta \rightarrow \mathbf{R}$ is a **synergistic utility function** if

- (i) f is everywhere both continuous and strictly increasing in its first argument
- (ii) f is everywhere both continuous and either strictly increasing, decreasing, or constant in each of its other real arguments
- (iii) there exists $\theta_E \in \Theta$ such that for all vectors \mathbf{v} in \mathbf{R}^I , $f(\mathbf{v}; \theta_E) = v_1$
- (iv) for all $\theta \in \Theta$, $f(\mathbf{0}; \theta) = 0$
- (v) for all $\theta \in \Theta$ and all \mathbf{v} in \mathbf{R}^I , $u^\infty(\mathbf{v}; \theta)$ exists (as defined above).

Stated otherwise, requirement (i) posits that utility must be increasing in one's own payoff. Requirement (ii) asks that utility, if it is affected by someone else's payoff, always be affected in the same direction. This could be weakened, but imposes no

untoward restrictions⁵. Requirement (iii) imposes that there exists a traditional game-theoretic type, that is, one who has utility equal to own payoff regardless of the other players in the game.⁶ Requirement (iv) is a moderately weak normalization that rules out adding arbitrary constants to the utility: essentially it states that you can't get something for nothing. And finally, requirement (v) insures that utilities exist in all cases and are well-defined.

Definition: If \mathbf{G} is a game with payoffs v_i , then we say (\mathbf{G}, f, θ) is a **game with synergistic utility** (a synergistic game) if it is identical to \mathbf{G} except that utility is given by $u_i = u_i^\infty(v_i, v_{-i}; \theta_i)$ for all i , and f is a synergistic utility function.

Proposition 1: If (\mathbf{G}, f, θ) is a synergistic game, then $u_i = f(v_i, u_{-i}; \theta_i)$ for all i .

Proposition 1 says that the limit utilities, which necessarily exist, satisfy a fixed-point property. The proof follows straightforwardly from the definitions and the continuity of f . One can imagine defining synergistic utilities directly as solutions to the fixed-point equation, but this definition has several weaknesses. First, the motivation for synergistic utilities—that players update their own welfare by taking into account the welfare of the other players—leads directly to the limit process. Secondly, the fixed-

⁵ Note, however, that it does not allow sufficient flexibility for very much reciprocity. This is by design: we see how much can be accomplished in as simple a setting as possible.

⁶ E stands for economist or egotist, as the reader prefers.

point solution may exist even if the limit does not.⁷ For example, suppose that we have two altruistic players of the same type; in particular we assume $f = v_i + 2u_{-i}$ for both.⁸ If $v_1 = v_2 = 1$ then the limit diverges, as would be expected (utilities go to infinity as each player gets happier and happier contemplating the situation). The fixed-point solution, on the other hand, yields $u_1 = u_2 = -1$, which appears unreasonable. Thus the limit is central to the definition, but Proposition 1 may provide a short-cut in explicit calculations.

Proposition 2: In a synergistic game, utilities u_i are continuous in payoffs \mathbf{v} .

Proof: Let $\mathbf{v} \in \mathbf{R}^I$ have associated synergistic utilities $\mathbf{u} \in \mathbf{R}^I$. Take any sequence $\{\mathbf{v}_n\}_{n=1}^\infty$ such that $\lim_{n \rightarrow \infty} \mathbf{v}_n = \mathbf{v}$. We wish to show that $\lim_{n \rightarrow \infty} \mathbf{u}_n = \mathbf{u}$. If not, there exists $\varepsilon > 0$ such that $B(\mathbf{u}, \varepsilon) \cap \{\mathbf{u}_n\}_{n=1}^\infty = \emptyset$. From the definition of synergistic utility, $\mathbf{u} = \lim_{m \rightarrow \infty} \mathbf{u}^m$ and hence there exists M such that $d(\mathbf{u}, \mathbf{u}^m) < \frac{\varepsilon}{2}$ for all $m \geq M$. But since f is continuous, we know that $\mathbf{u}^1 = \lim_{n \rightarrow \infty} \mathbf{u}_n^1$, and iterating $\mathbf{u}^2 = \lim_{n \rightarrow \infty} \mathbf{u}_n^2$, ... so that in particular $\mathbf{u}^M = \lim_{n \rightarrow \infty} \mathbf{u}_n^M$. Thus we can choose N with the property that $d(\mathbf{u}^M, \mathbf{u}_n^M) < \frac{\varepsilon}{2}$ for all $n \geq N$. But now $d(\mathbf{u}_N^M, \mathbf{u}) \leq d(\mathbf{u}_N^M, \mathbf{u}^M) + d(\mathbf{u}^M, \mathbf{u}) < \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon$, implying $\mathbf{u}_N^M \in B(\mathbf{u}, \varepsilon)$.

This is a contradiction, and so we're done. □

⁷ In general, of course, there may be several fixed-point solutions, while there is necessarily at most one limit point. This is another reason to choose the limit definition, although in synergistic games as defined multiplicity won't be a problem.

⁸ Note that since f is simply a function of bound variables, whether we write the other players' welfare functions as v or u is a matter of clarity and convenience only.

Proposition 2 gives us another general property of synergistic utility functions, but this is about as much as can be said in complete generality. It may be helpful at this point, in part to clarify the definitions, to consider some examples of potential synergistic utility functions. We say potential because for the moment we ignore condition (v), and we leave Θ unspecified. The most obvious is probably the linear formulation $f = av_i + \sum_{j \neq i} b^j u_j$. Here $\theta = (a, \mathbf{b})$ and $\Theta \subseteq \mathbf{R}^I$. On the other hand, $f = av_i + b(u_{-i})^2$ is impermissible, for instance, because it violates (ii). The effect of an increase in the other player's utility on one's own should be independent of the absolute levels involved. Thus, $f = av_i + b(u_{-i})^3$ is once again acceptable. Cobb-Douglas formulations, more common in macroeconomics, look like $f = (v_i)^a (u_{-i})^b$ and require that "consumptions" be nonnegative. However, upon taking logs, this is equivalent to the original linear form.⁹ All of the above satisfy condition (iii) by choosing $a = 1$ and $b = 0$, and satisfy condition (i) if $a > 0$. Examples of applications of these utility functions to particular games, along with an additional nonlinear formulation, are given in Section 4.

To apply the theory in a specific situation, one must choose an appropriate (f, Θ) pair and show that this pair yields a synergistic utility function. We do this now for the two-player linear case, though it is easy to extend it to more players.

Proposition 3: Let $\Theta = (0, \infty) \times (-1, 1)$ and $\theta = (a, b)$. Then $f(v_i, u_{-i}; \theta) = av_i + bu_{-i}$ is a synergistic utility function.

⁹ Note that we cannot then independently choose the cardinalization for taking expected utilities.

Proof: We have the recursive equations $u_1^n = a_1 v_1 + b_1 u_2^{n-1}$ and $u_2^n = a_2 v_2 + b_2 u_1^{n-1}$, or

$$\begin{bmatrix} u_1^n \\ u_2^n \\ 1 \end{bmatrix} = \begin{bmatrix} 0 & b_1 & a_1 v_1 \\ b_2 & 0 & a_2 v_2 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u_1^{n-1} \\ u_2^{n-1} \\ 1 \end{bmatrix}.$$

We may write this as $u^n = \mathbf{M}u^{n-1}$, and hence $u^n = \mathbf{M}^n u^0$, where $u^0 = [v_1 \ v_2 \ 1]^T$.

Then multiplying out the powers of \mathbf{M} yields

$$\mathbf{M}^{2m} = \begin{bmatrix} (b_1 b_2)^m & 0 & (a_1 v_1 + b_1 a_2 v_2) \sum_{i=0}^{m-1} (b_1 b_2)^i \\ 0 & (b_1 b_2)^m & (a_2 v_2 + b_2 a_1 v_1) \sum_{i=0}^{m-1} (b_1 b_2)^i \\ 0 & 0 & 1 \end{bmatrix}.$$

But by assumption $|b_1 b_2| < 1$ so

$$\lim_{n \rightarrow \infty} \mathbf{M}^n = \frac{1}{1 - b_1 b_2} \begin{bmatrix} 0 & 0 & a_1 v_1 + b_1 a_2 v_2 \\ 0 & 0 & a_2 v_2 + b_2 a_1 v_1 \\ 0 & 0 & 1 - b_1 b_2 \end{bmatrix}.$$

Now $\lim_{n \rightarrow \infty} u_i^n$ is simply the i^{th} row of the 3rd column of the matrix above so it too exists (and in fact this gives an explicit formula for it). Naturally, this is the same solution one would find from solving the system of two fixed-point equations. It is clear that conditions (i)-(iv) also hold. □

Note that the perverse example mentioned earlier, which had $b = 2$, is not allowed in this scenario. Nonlinear synergistic utility functions will have their own

requirements for Θ .¹⁰ Turning to another question that can be answered given a specific synergistic utility function, it is well known that positive linear transformations of any player's payoffs leaves the strategic structure (that is, the preferences over final outcomes) of a game unaffected. This result carries over to synergistic games as much as possible (it is clear that multiplying only one player's payoffs by some constant may substantively change utilities in an interdependent setting).

Proposition 4: In a linear synergistic game, preferences over outcomes are unaffected if

- (a) all player's payoffs are multiplied by the same positive constant, or
- (b) any or all players have a constant added to their payoffs

Proof: (a) Since f is linear in v_i (or in fact more generally whenever f is homogeneous of degree one in v_i), utilities all along the limiting sequence, and hence also final utilities, will also be multiplied by this constant. So then, by the standard result, preferences remain the same.

(b) Adding a constant to one player's payoffs affects all players, but only to the extent of adding some constant to each of their payoffs. Although this constant may be different for each player, it is the same for a given player across his or her outcomes. This is clear from the explicit formulas found in the proof of Proposition 3. But now, once again, the standard result applies. □

¹⁰ For example, we might imagine that more generally one would require the derivative of f with respect to opponent's utility to be bounded by 1.

Although this result doesn't hold in general for all synergistic games, it will hold in other particular settings. We now turn to illustrating the theory with a spectrum of examples.

4. Examples

The proof of the pudding lies in the taste, and the believability of synergistic utilities lies in its potential applications. For the time being, we confine ourselves to the linear synergistic utility function analyzed above, $f = av_i + bu_{-i}$. We first define three types of players to convey some idea of the range of possibilities. Although unnecessary, it is convenient to choose them such that $a + |b| = 1$; this keeps the magnitude of the utilities directly comparable to those of the base payoffs.¹¹ The first type is the one required by part (iii) of the definition, $\theta_E = (1, 0)$. This type always has final utility equal to base payoff regardless of the other players. The second type is an altruist, denoted by S for socialist: $\theta_S = (\frac{1}{2}, \frac{1}{2})$. This type approximately treats the two players equally. Finally, we define an unfriendly type: $\theta_J = (\frac{1}{3}, -\frac{2}{3})$. In the game theory literature, this general type has been called spiteful, but that is perhaps too strong a condemnation for these preferences. Rather, this player simply enjoys doing better than his or her opponent; the notation is thus Jones, for "keeping up with the Joneses."¹² Note that since we apply the theory to single games, it is possible to switch types over time or in differing situations or against different players. The model does not require players to be

¹¹ Most of the previous literature has instead chosen (in its own context) $a = 1$.

¹² A similar Jones type appears in the macroeconomics consumption literature, so this is conceivably an example of micro keeping up with the macro Joneses.

one intrinsic type. Also, it is fairly easy to see how to come up with multi-player analogies for these types.

		Player 2	
		C	D
Player 1	C	0,0	-9,3
	D	3,-9	-6,-6

The basic Prisoner's Dilemma can be written as:

Here C stands for cooperate and D for defect, as usual. Of course the unique Nash Equilibrium (NE) is (D,D) . If two type E 's (economists) play against one another, the payoffs remain as these began and the game is unchanged. So the unique NE is also the same. We next consider an economist player 1 opposing a Jones player 2. E 's utilities are the same as ever, while J 's may then be calculated using f (it takes only one step in this case). We arrive at the following game form:

		type J	
		C	D
type E	C	0,0	-9,7
	D	3,-5	-6,2

The unique NE is again for both players to defect. What is interesting, however, is that this outcome is no longer Pareto inefficient, as it was previously. The economist is so unhappy that it makes the Jones player happy. Of course, this outcome depends on the exact payoff structure and player 2's type, but holds over a wide class. Consider next a socialist player 1 against a Jones type:

		type J	
		C	D
type S	C	0,0	-3,3
	D	0,-3	-3,0

This game now has two pure NE, in both of which type *J* defects (unsurprisingly it turns out that types *E* and *J* always defect). Type *S* is completely indifferent, and is thus willing to cooperate. Of course this willingness is knife-edged; types near to *S* will be pushed in one direction or the other, some of them always cooperating. The (C,D) equilibrium is [weakly] Pareto efficient in this case. We now change player 2 to a type *S* as well:

		type S	
		C	D
type S	C	0,0	-5,-1
	D	-1,-5	-6,-6

Cooperation is a dominant strategy here for both players; it is also the optimal outcome in the game. This is the stereotypical result of altruistic cooperation in the Prisoner's Dilemma. The final combination of players that we consider is when player 1 is a type *E* once more:

		type S	
		C	D
type E	C	0,0	-9,-3
	D	3,-3	-6,-6

The unique and strict NE is (D,C) . The surprising observation here is that it requires less inherent altruism to cooperate with a type E than with a type S .¹³ This result can be explained by noting that defection hurts a type E opponent more than it does a type S opponent (who is consoled by the fact that one's own payoff has been improved). Hence a type S player will have a stronger incentive not to defect when playing against a type E . Recall that we have tried to put aside any issues of reciprocity.

Turning next to an example of a continuous game, we consider Cournot duopoly. In the simplest case with linear unit demand and zero marginal cost, price $p = 1 - q$, where q is the total quantity produced. Payoffs are simply net profits, so $v_i = q_i(1 - q)$. The unique NE with standard (that is, type E) players is for both to produce at $q_i = \frac{1}{3}$. It is plausible, however, to model the firms as type J . Perhaps it is a small market so that profits themselves are not important but beating the rival firm is critical for advertising. Or perhaps the managers are paid with yardstick competition incentives, so again what is important is to do better than the other firm. The symmetric equilibrium in this case is that both produce $q_i = \frac{3}{7}$. In the end of course neither firm actually does any better than the other, but each is willing to overproduce ("sacrificing" profits) in order to try to do so.

¹³ Contrast this once again with the quote from Rotemberg (1994) in Section 2.

Note also that this is much closer to the zero profit outcome of Bertrand competition, and in fact it converges to that outcome if the firms converge toward the full Jones type.

Experimental game theory has worked extensively not only with the Prisoner's Dilemma but also with other games such as ultimatum, dictator, centipede, and public goods. As in the case of the Prisoner's Dilemma, the results are often quite disparate from those predicted by standard theories. For instance, no positive quantity should ever be rejected in an ultimatum game, yet this result is often observed in experiments. This outcome can be explained using synergistic utilities: players similar to the Jones type will reject all offers up to some level (which will depend on the exact type chosen and on the type of the opponent). Of course altruism alone, without some sort of negative analogue, can never rationalize these rejections. Recall that it is possible to extend the theory to include reciprocity if desired, so a player's type need not remain constant. As has been documented previously (see section 2), altruism can explain extended play in a centipede game or a contribution made in a public goods game. The point is that a simple theory, such as synergistic utilities, is sufficient to explain these wide-ranging outcomes.

In the so-called dictator game, player 1 simply decides how to divide an amount of money (typically around \$10 in experiments) between himself or herself and an often anonymous opponent. Player 2 has no action other than to accept the split as dictated. Traditional equilibrium concepts predict that player 1 should keep the entire amount, and previous models of altruism have not altered this prediction. For instance, continuing with the types as defined above, if an altruistic type S opposes another type S , the optimal action is still to give nothing away. No linear model can predict an interior solution, although in practice this is what the data clearly support. We turn, then, to a nonlinear

synergistic utility function. For simplicity we assume that player two is a type E , so that as always $u_2 = v_2$. For player 1, we assume the altruistic formulation $u_1 = \sqrt{v_1 u_2}$. In this case the optimal allocation is an even split, that is, \$5 for each player. This outcome is occasionally, though rarely, observed in experiments. If we assume instead the slightly less magnanimous utility $u_1 = \frac{1}{2}v_1 + \frac{1}{2}\sqrt{v_1 u_2}$, then we find $v_1^* \approx \$8.54$. In fact this result agrees remarkably well with the observed average division. Naturally, this is meant only to illustrate the potential applicability of the theory, in addition to the fact that nonlinear functions do not simply provide generality but in fact may be necessary in practice.

5. Topics

Despite the fact that the game structure remains the same in synergistic games (only the payoffs have changed), there are several topics that take on new meaning in this context. For instance, cooperative games with transferable utility will be difficult to analyze since some players may actually prefer having a smaller total surplus to divide (think of the type J above). Evolutionary game theory, a popular subject of recent study, offers another example. In the present setting, it is possible to discuss the evolutionary strengths not just of different strategies but also of different synergistic types. What is unclear, however, is what to use as a measure of reproductive fitness. One could argue that players with the highest welfare (final utility) will be the most productive and successful. On the other hand, it may be that success is determined by physical rather than mental well-being, so that base payoffs (food or money leading to direct consumption) should enter the calculation of the dynamics. A player might be happy that others do well, but this does not necessarily grant an increased chance of survival. The

appropriate measure may depend on the particular situation. Note that in the Prisoner's Dilemma example of section 4, altruistic players, type S in the notation there, fare relatively poorly under either system.

A related consideration, though more in the mode of full rationality, is the idea of segregation. Since players are of different types, they may prefer to play against one type of opponent rather than another, and thus selectively associate. Of course, they may not have the opportunity to make this choice, but if they do then it has long-term welfare (and hence possibly evolutionary) implications. Returning once again to the Prisoner's Dilemma example of the previous section, note that while types E and S always prefer an altruistic type S opponent, this is not necessarily true of type J players, who like to play type E 's (since the latter end up so unhappy). So a plausible scenario is that S types play against themselves, while J s and E s pair off against one another. This leaves the self-centered economist types quite unhappy; their only hope is to run across extremely altruistic players, who will actually want to make them happy by cooperating (in effect, happily sacrificing themselves). Recall that all players are fully utility maximizing at all times.

There is no doubt at least some element of reciprocity in almost all human interactions. Synergistic utilities, as defined, make no allowance for this: a player's degree of altruism is independent of the attitudes of the other players. The work of Rabin (1993) and Sally (2001) depend explicitly on these added interactions, and similar constraints can be added to synergistic games. One approach would be to require that players enter a game with their own individual synergistic type θ , but then require that all of the players play the game using the average θ of the group (if Θ is such that this

has meaning). Another possibility is to add a reciprocity player, type R , who takes on the θ of whomever he or she is playing. As always, this is difficult to implement with more than two players. The point is that players being subject to altruism, jealousy, and so on makes sense independently of any reciprocity arguments, so the simplest models of such behavioral tendencies will not include these traits as building blocks. Yet including these tendencies may be necessary in order to fully explain either our own introspective assessments or all empirically observed behavior.

Finally, games with incomplete information take on an added dimension if it is possible to allow for players of synergistic type. In general, there is no reason to assume that all players know each of their opponents' type, synergistic or otherwise. Fortunately, the entire game-theoretic apparatus developed to analyze this eventuality is still perfectly applicable. In particular, the Bayesian equilibrium concepts apply just as well here. As the presence of synergistic types are certainly payoff-relevant, signaling will be an important component to playing extensive-form synergistic games. It may or may not be beneficial for a player in a given situation to reveal his or her synergistic type (consider, for instance, the discussion of segregation above). In fact, aspects of synergistic games involving incomplete information seem perhaps to be the most fruitful line for future theoretical research using this model.

6. Conclusion

Game theorists assume that the payoffs in a game indicate players' true preferences, which is to say that they already take into account welfare interactions between the players. But in real-life situations, often the only information available is

about base payoffs, for example, firm profits or monetary payoffs in an experimental setting. It is useful to have a specific model of altruism and other emotional aspects that influence preferences and welfare in order to link these payoffs to the ultimate utilities in a game. The concept of synergistic utilities attempts to do this by providing a simple framework in which to address these concerns in various applied contexts. In this model each player's utility is a function of his or her own payoff and of the other players' utilities. Standard equilibrium concepts are sufficient, and since the process involves a transformation of payoffs only, the theory can be applied to arbitrary games with any number of players. One special case, a linear formulation, was given and analyzed in more detail, followed by examples such as how both cooperation in the Prisoner's Dilemma and positive gifts in the dictator game can be rationalized.

The main distinction between the present work and previous literature lies in the simplicity of synergistic games. There is nothing new imposed on the game structure or analysis, since the only change made is in the numerical values of the payoffs. Nor is an idea of reciprocity inherent or necessary to the model. Nevertheless, many observed behaviors can be explained within this paradigm. Note in particular that standard theories, ignoring any subtleties regarding preferences, have done exceptionally well in predicting behavior in market situations. In these games, by definition, a player cannot influence the payoff of any other participant in the game (or at least has this impression). Hence a player with synergistic utility will behave exactly as a standard player would, which provides a robustness check on the theory. Surely there will be more such checks to come.

References

- Barro, Robert J. 1974. "Are Government Bonds Net Wealth?" *Journal of Political Economy* 82(6): 1095–1117.
- Becker, Gary S. 1974. "A Theory of Social Interactions." *Journal of Political Economy* 82(6): 1063–1093.
- Bergstrom, Theodore C. 1989. "Love and Spaghetti, The Opportunity Cost of Virtue." *Journal of Economic Perspectives* 3(2): 165–173.
- Bergstrom, Theodore C. 1995. "On the Evolution of Altruistic Ethical Rules for Siblings." *American Economic Review* 85(58–81).
- Bergstrom, Theodore C. 1999. "Systems of Benevolent Utility Functions." *Journal of Public Economic Theory* 1(1): 71–100.
- Binmore, Ken. 1994. *Game Theory and the Social Contract Volume 1: Playing Fair*. Cambridge, MA: The MIT Press.
- Bisin, Alberto, and Thierry A. Verdier. 1998. "On the Cultural Transmission of Preferences for Social Status." *Journal of Public Economics* 70, 75–97.
- Bisin, Alberto, and Thierry A. Verdier. 2001. "The Economics of Cultural Transmission and the Dynamics of Preferences." *Journal of Economic Theory* 97(2): 298–319.
- Bruce, Neil, and Michael Waldman. 1990. "The Rotten-Kid Theorem Meets the Samaritan's Dilemma." *Quarterly Journal of Economics* 105(1): 155–165.
- Charness, Gary, and Ernan Haruvy. 2002. "Altruism, Equity, and Reciprocity in a Gift-Exchange Experiment: An Encompassing Approach." *Games and Economic Behavior* 40(2): 203–231.
- Coate, Stephen. 1995. "Altruism, the Samaritan's Dilemma, and Government Transfer Policy." *American Economic Review* 85(1): 46–57.
- Collard, D. 1975. "Edgeworth's Propositions on Altruism." *Economic Journal* 85(338): 355–360.
- Dawes, Robyn M., and Richard H. Thaler. 1988. "Anomalies: Cooperation." *Journal of Economic Perspectives* 2 (3): 187–197.
- Frank, Robert H. 1987. "If Homo Economicus Could Choose His Own Utility Function, Would He Want One with a Conscience?" *American Economic Review* 77(4): 593–604.

Geanakoplos, John, David Pearce, and Ennio Stacchetti. 1989. "Psychological Games and Sequential Rationality." *Games and Economic Behavior* 1(1): 60–79.

Kotlikoff, Laurence J., Assaf Razin, and Robert W. Rosenthal. 1990. "A Strategic Altruism Model in Which Ricardian Equivalence Does Not Hold." *Economic Journal* 100(403): 1261–1268.

Levine, David K. 1998. "Modeling Altruism and Spitefulness in Experiments." *Review of Economic Dynamics* 1(3): 593–622.

Palfrey, Thomas R., and Howard Rosenthal. 1988. "Private Incentives in Social Dilemmas: The Effects of Incomplete Information and Altruism." *Journal of Public Economics* 35(3): 309–332.

Rabin, Matthew. 1993. "Incorporating Fairness into Game Theory and Economics." *American Economic Review* 83(5): 1281–1302.

Rotemberg, Julio J. 1994. "Human Relations in the Workplace." *Journal of Political Economy* 102(4): 684–717.

Sally, David .2001. "On Sympathy and Games." *Journal of Economic Behavior and Organization* 44(1): 1–30.

Wolpert, David H., Julian Jamison, David Newth, and Michael Harre. 2010. "Strategic Choice of Preferences: The Persona Model." *FRBB Working Paper 10-10*.