Selfishness vs Altruism vs Balance

Pradeep Dubey[*] and Yair Tauman[†]

18 April 2017

Abstract

We give examples of strategic interaction which are beneficial for players who follow a "middle path" of balance between pure selfishness and pure altruism.

JEL Classification: C70, C71, C72, C79.

1 Introduction

For the most part, game theory is silent about how payoffs of players are defined in a game. Indeed, as a rule, payoffs are taken to be exogenous; and the focus of the theory is on strategic behavior that emerges endogenously from the given payoffs (mostly by invoking Nash Equilibrium (NE), but sometimes by means of other solution concepts).

In this note, we shall commit the sacrilege of comparing different payoffs, or more precisely different “standpoints” (in the jargon, “types”) of the players, which in turn determine their payoffs.

Let us first think of two players named $1,2$ whose interaction can be modeled as if they are in one of several possible “games of money”, picked by nature according to a common prior. Depending on the particular game-form at hand, as well as the strategies chosen by the players, the *outcome* is a pair of numbers (x_1, x_2) where x_α denotes the dollars accruing to player $\alpha = 1,2$.

[*]Center for Game Theory, Stony Brook University; and Cowles Foundation for Research in Economics, Yale University

[†]Center for Game Theory, Stony Brook University; and Interdisciplinary Center, Herzliya, Israel.

The *payoff* (ζ_1, ζ_2) that arises from the outcome (x_1, x_2) depends on the *type* of the players. We focus on three types:

Selfish (S)**:** each values only his own reward[1], i.e., $\zeta_1 = x_1$ and $\zeta_2 = x_2$.

Altruistic (A)**:** each values only his opponent's reward, i.e., $\zeta_1 = x_2$ and $\zeta_2 = x_1$

Balanced (B)**:** each values both rewards equally, i.e., $\zeta_1 = \zeta_2 = (x_1 + x_2)/2$

For the time being, assume that both players are of the *same* type. This type, moreover, is common knowledge among the players, as is the game form realized by nature and the prior π. Thus both play the game of complete information that occurs after nature's move, and calculate their overall expected payoff in accordance with π.

Our observation is that if the domain of game-forms is "unbiased" and "varied" enough, there exists an NE selection for B-games whose expected payoff is higher than those of *any* NE-selection for the other two games, not only if these payoffs are calculated from the B-standpoint, but *even if they are calculated from the S-standpoint or the A-standpoint*.

Later we shall consider prisoner's dilemma, in which players have the strategic freedom to program their computer to play according to one of the standpoints S, A, B. We shall pinpoint conditions under which the B-program emerges as the unique "trembling-hand perfect" NE.

All our observations are in the nature of examples rather than any general theory and, once made, they are completely obvious. Nonetheless we think they are worthy of being placed on record, in order to stimulate further inquiry.

[1]We may replace x_i by $u(x_i)$ in defining ζ_i, where u is the utility of money. This does not affect our analysis but just addes a layer of unnecessary notation. The important assumption we are making is that the *same* u applies to both players. There is no interpersonal *transfer* of utilities in our model, but there *is* interpersonal *comparison* of utilities. (The payoffs of the two players cannot be subject to arbitrarily different affine transformations as in the standard model of cardinal utilities, otherwise any notion of altruism or balance would become problematic.) We think that people can, and often do, put themselves in "the other's shoes" and have empathy (otherwise even communication by language would be difficult). Thus interpersonal comparison of utilities is not so unnatural.

2 Identical Types: the General Setting

A *game-form* is a map from strategy-pairs to outcomes[2]. For ease of exposition, we assume that the domain Γ of game-forms is finite and furthermore that, in each $\gamma \in \Gamma$, players have finitely many pure strategies (though the number of these strategies may well vary across Γ). Thus any game-form γ is given by a pair (U,V) of $n \times k$ matrices where player 1 has n pure strategies (the rows), 2 has k pure strategies (the columns)[3], and $U = (u_{ij}), V = (v_{ij})$ give respectively the monetary payoffs of $1,2$. In other words, if 1 picks row i and 2 picks column j, their payoffs in $\gamma = (U,V)$ are u_{ij} dollars and v_{ij} dollars respectively. (As was said, the integers n,k can depend on γ.)

There is a *probability distribution* π on Γ, indicating the likelihood of any game-form coming into play.

For any $\gamma = (U,V) \in \Gamma$, denote $\gamma^t = (V^t, U^t)$ where V^t (resp. U^t) means the transpose of the matrix V (resp. U). If the roles of the two players is interchanged, the game-form γ is transformed to $\gamma^t \in \Gamma$.

To rule out bias in favour of any player in our model, we postulate[4]

Axiom 1 ***(Unbiased Domain)*** *If $\gamma \in \Gamma$, then $\gamma^t \in \Gamma$; and, moreover, $\pi(\gamma) = \pi(\gamma^t)$.*

A game-form will yield a *game*, once we map outcomes into payoffs in accordance with one of the three standpoints S or A or B.

Actually it is just as easy to consider a continuum of standpoints or (player-) *types*, parametrized by $0 \leq c \leq 1$, where S or A or B will correspond to $c = 1$ or 0 or $1/2$ respectively.

For any $\gamma = (U,V) \in \Gamma$, denote $\gamma_c = (U_*, V_*)$ where U_* is obtained by replacing the ij^{th} entry u_{ij} of U by $cu_{ij} + (1-c)v_{ij}$; and, similarly, V_* is obtained by replacing the ij^{th} entry v_{ij} of V by $(1-c)u_{ij} + cv_{ij}$.

Then the game-form γ becomes the game γ_c, once we specify the type c of the players.

Let $\Sigma(\gamma) = \Sigma_1(\gamma) \times \Sigma_2(\gamma)$ denote the Cartesian product of the mixed-strategy sets $\Sigma_1(\gamma), \Sigma_2(\gamma)$ of players $1,2$ in the game-form $\gamma \in \Gamma$.

[2]Thus a game-form is one step short of a full-blown game. It will become such, once payoffs of all the players are defined at every outcome.

[3]*Throughout* this note, the row player will be 1 (Rowina) and the column player will be 2 (Colin).

[4]An equivalent way to state Axiom 1 is that nature picks a game form γ in Γ according to some common prior, and *independently* assigns the role of row or column player to $1,2$ with equal probability. (In this language, given any game form γ, Rowina is the row player and Colin the column player with probability $1/2$, and their roles are reversed in γ with the remaining probability.)

Definition 2 *A* strategy selection *is a map f on the domain Γ such that, for all $\gamma \in \Gamma$*

$$f(\gamma) \in \Sigma(\gamma)$$

and

$$f(\gamma) = (p,q) \Longrightarrow f(\gamma^t) = (q,p)$$

The second display reiterates the absence of bias, not merely in the domain (as in Axiom 1) but also in the strategy selection on the domain.

Next, for any $0 \leq c \leq 1$ and $\gamma \in \Gamma$, denote by $\Phi(\gamma_c) \subset \Sigma(\gamma)$ the set of all mixed strategy *Nash Equilibria* (NE) of the game γ_c.

Definition 3 *A c*-NE selection *is a strategy selection f which satisfies*

$$f(\gamma) \in \Phi(\gamma_c)$$

for all $\gamma \in \Gamma$.

By Axiom 1 and Nash's theorem, c-NE selections exist for all c.

It will also be useful to formalize the payoff of a c-type player at any strategy selection f.

Definition 4 *For any $0 \leq c \leq 1$ and any strategy selection f, the c*-payoff to player α at f *is*

$$Exp\,(f,c,\alpha) = \sum_{\gamma \in \Gamma} \pi(\gamma)\zeta\,(f\,(\gamma)\,,\gamma,c,\alpha)$$

where $\zeta\,(f\,(\gamma)\,,\gamma,c,\alpha)$ denotes the expected payoff to player α of type c in the game-form γ when the mixed strategies $f\,(\gamma)$ are played.[5]

Before we state our proposition, we need to also rule out domains which consist *exclusively* of games of coordination. To this end, let us say that the pair (i^*, j^*) of pure strategies in the game-form $\gamma = (U,V)$ is a *maximizer* if

$$u_{i^* j^*} + v_{i^* j^*} = \max_{i,j}\left(u_{ij} + v_{ij}\right) := \max \gamma$$

Axiom 5 ***(Non-trivial Domain)*** *There exist game-forms $\gamma \in \Gamma$, with $\pi(\gamma) > 0$, such that no maximizer of γ is a Nash Equilibrium of the (selfish) game γ_1.*

[5] i.e., when $\alpha = 1$, we have $\zeta\,(f\,(\gamma)\,,\gamma,c,\alpha) = \sum_{i,j} p_i q_j\,(cu_{ij} + (1-c)\,v_{ij})$ where $f\,(\gamma) = (p,q)$ and $\gamma = (U,V)$; and the payoff when $\alpha = 2$ is exactly the same, except that we must switch c and $1-c$.

We ready to state

Proposition 6 *Suppose Axiom 1 holds. There exists a* $(1/2)$*-NE selection* f^* *(i.e., NE for the balanced type) such that, for* $\alpha = 1,2$*:*

$$Exp\ (f^*, c, \alpha) = M := \frac{1}{2} \sum_{\gamma \in \Gamma} \pi(\gamma) \max \gamma, \textit{ for } \text{all } 0 \leq c \leq 1$$

Furthermore, for any *strategy selection* f*, and any player* $\alpha = 1,2$ *we have:*

$$M \geq Exp\ (f, c, \alpha), \textit{ for } \text{all } 0 \leq c \leq 1$$

Finally, if Axiom 5 is also satisfied, then at any (selfish) 1*-NE selection* f *and* $\alpha = 1,2$ *we have*

$$M > Exp\ (f, c, \alpha), \textit{ for } \text{all } 0 \leq c \leq 1;$$

(and the same strict inequality holds for (altruistic) 0*-NE selections* f*, provided we postulate the variant of Axiom 5 with* γ_0 *in place of* γ_1*.).*

Proof. (As the reader will have surely anticipated), define f^* by selecting a pure strategy maximizer (i^*, j^*) in each game γ, ensuring that (j^*, i^*) is then selected in γ^t. (This is feasible on account of Axiom 1.) Then f^* is clearly a $(1/2)$-NE-selection[6]. Now, for *any* $0 \leq c \leq 1$, the c-payoff of player 1 at f^* — across the two games γ and γ^t — is

$$\pi(\gamma)\left[cu_{i^*j^*} + (1-c)v_{i^*j^*}\right] + \pi(\gamma)\left[(1-c)u_{i^*j^*} + cv_{i^*j^*}\right] = (1/2)\left[\pi(\gamma)\max\gamma + \pi(\gamma^t)\max\gamma^t\right]$$

since $\pi(\gamma) = \pi(\gamma^t)$ and $u_{i^*j^*} + v_{i^*j^*} = \max\gamma = \max\gamma^t$. Player 2 gets the same by symmetry. Adding across all pairs γ, γ^t in the domain Γ establishes the first display of the proposition.

Next let f be any strategy selection which assigns mixed strategies $f(\gamma) = (p, q)$ to γ, and hence $f(\gamma^t) = (q, p)$ to γ^t. Then 1 will get the payoff $cu_{ij} + (1-c)v_{ij}$ in γ with probability $\pi(\gamma)p_i q_j$, and he will get $(1-c)u_{ij} + cv_{ij}$ in γ^t with the same probability. But $u_{ij} + v_{ij} \leq \max\gamma = \max\gamma^t$. Adding this inequality over all strategy pairs (prescribed by f) across all the games in Γ establishes the second display.

The third display follows by the same argument, in light of Axiom 5. ■

[6]Needless to say, there may be other $(1/2)$-NE-selections which give much lower expected payoffs to balanced types.

Remark 7 *We could have considered abstract strategy sets S_1 and S_2 and monetary payoffs on their Cartesian product, with enough convexity assumptions to guarantee existence of NE. The same results may then be established mutatis mutandis.*

Remark 8 *Also with n players, define c-type to be a player who puts weight c on his own reward and weight $1-c$ on the average of others' rewards. Then our analysis can be extended straightforwardly.*

Remark 9 *The second display implies in particular that f^*, generated by the balanced standpoint $1/2$, is as good as* any *c-NE selection (not just selections for the selfish $c=1$ or altruistic $c=0$), and that this is so* no matter which standpoint in the interval $[0,1]$ is used to compare the two.

Remark 10 *The third display of the proposition is silent on the "percentage drop" in efficiency of NE payoffs of S or A games (compared to the "benign" NE selection $f^*$of B-games, and always using the standpoint $1,0$ when talking of the S,A games respectively). This is a delicate question and will depend very much on the domain Γ. See Powers et al (2016) for a detailed analysis of various domains of 2×2 games.*

3 Mixed Types: Prisoner's Dilemma

So far, we have compared three starkly different standpoints S,A,B from which players may approach a game, and shown that the B-standpoint (that of the "middle path") will be most successful, achieving the highest average payoffs for its adherents. The question was not raised as to what might happen when adherents with different standpoints come face-to-face. Let us examine this possibility in the context of the well-known "prisoners' dilemma".

With κ for "cooperate" and δ for "defect", consider the game-form:

	κ	δ
κ	$-3,-3$	$-12,-1$
δ	$-1,-12$	$-5,-5$

(The only "non-ordinal" constraint on the monetary rewards, that we shall be invoking here, is that when player 1 unilaterally deviates from (δ,δ) to (κ,δ), what he loses is greater than what 2 gains; the remaining constraints are purely ordinal, as in the standard prisoners' dilemma[7].)

The best response map in the above 2×2 game-form for each of the three types S, A, B is given in the tables below, where the first row shows the opponent's strategy and the second row shows the corresponding best response:

***S*-type:**

κ	δ
δ	δ

***A*-type:**

κ	δ
κ	κ

***B*-type:**

κ	δ
κ	δ

Notice that both S and A have strictly dominant strategies, namely δ and κ respectively; whereas B plays "tit-for-tat".

Now let us think of a related game where each player chooses one of three possible programs for his computer or automaton, namely: "play according to S (or A, or B"). When any two automata meet, they converge to the unique NE of the ensuing 2×2 game. This gives rise to the overall 3×3 symmetric game below:

	S	A	B
S	$-5, -5$	$-1, -12$	$-5, -5$
A	$-12, -1$	$-3, -3$	$-3, -3$
B	$-5, -5$	$-3, -3$	$-3, -3$

Here (B, B) is an NE, and so is (S, S). However B dominates[8] A; and once we delete A, then B dominates S in the remaining subgame, hence (B, B) is — in this sense — the more "stable" of the two.

[7]To be explicit, replace $-3, -12, -1, -9$ by a, b, c, d. Then the standard prisoners' dilemma requires $c > a$ and $d > b$. We introduce the additional constraint: $d - b > c - d$. (Our analysis holds for any such a, b, c, d.)

[8]By "B dominates A" we mean that that B does better than A against *some* strategy of the opponent (in the game at hand, against S), and does no worse against any other strategy. (This is also called "weak domination" sometimes.)

4 References

M. R. Powers, M. Shubik, W. Wang . “Expected Worth for 2×2 Matrix Games with Variable Grid Sizes” (2016), Cowles Foundation Research Paper 2039 R