Effectivity and Noncooperative Solution Concepts

Paul Harrenstein[1] and Mehdi Dastani[2]

1. Institut für Informatik
Ludwig-Maximilians-Universität, Munich

2. Institute of Information and Computing Sciences
Utrecht University, Utrecht

Abstract. Game theory analyzes the strategic aspects of situations of social interaction by representing them as mathematical structures. Any such structure should at least account for the participants' powers and their preferences. The way powers and preferences are modelled makes a difference as to which strategic concepts are available for mathematical scrutiny. In *strategic games* the players' powers are modelled as possible courses of action. This provides sufficient structure to analyze interactive situations by means of the noncooperative solution concepts, in particular *Nash equilibrium*.
There are also different ways of representing the players' powers. One of them being by means of *effectivity sets*, i.e., sets of outcomes in which a player can guarantee a game to end by choosing a particular course of action. One could favor this way of representing interactive situations as, indeed, it has been fashionable in recent formalisms, such as ATL, ATEL and coalition logic.
This paper concerns the extent to which the notions available in strategic games are still applicable in case the players' powers are represented in terms of their effectivity. We address this issue in particular for Nash equilibrium and a related solution concept, which we call *outcome equilibrium*.

1 Introduction

In a situation of social interaction various active elements operate, each of which has its individual interests and capabilities. Game theory strives for a better mathematical understanding of such situations by formally modelling them as *games* and analyzing these by means of a range of solution concepts. (cf., e.g., von Neumann and Morgenstern, 1944; Luce and Raiffa, 1957; Osborne and Rubinstein, 1994). Any such formal representation should at least account for the active elements of a game—i.e., the players or agents—the preferences of the players over the possible outcomes of the game, and the various ways in which the players can manipulate the outcome of the game. Multiagent systems constitute a natural field of application for game theoretic concepts. Moreover, the issue of how to represent multiagent systems is an important issue when it comes to formally reasoning about them.

Formally, games can be described at different levels of abstraction. Perhaps most common in this respect is a *game in strategic or normal form*.[1] The manipulative powers of a player are then represented by a number of *strategies*. Each player of the game

[1] Canonical references that may be consulted for a more detailed exposition of the elementary concepts of game theory presented in this paragraph are, e.g., von Neumann and Morgenstern

has to decide on playing one of the strategies at his disposal. Each combination of strategies, or *strategy profile*, makes that a particular state of affairs, or *outcome*, comes about. Given the preferences of the players over the outcomes — which can be modelled in a number of ways, e.g., implicitly by utility vectors or by explicitly stipulating a preference relation over the outcomes — the formal structure of a strategic game constitutes the basis for the definition of the noncooperative game theoretic solution concepts, in particular that of a *Nash equilibrium* (Nash, 1950, 1951).

At this point we wish to emphasize a fundamental conceptual distinction between strategy profiles and outcomes. The way we employ the notions in this paper, strategies and strategy profiles pertain to the actions of the players, whereas outcomes are the entities with respect to which the preferences of the players are defined. A strategy profile is a combination of strategies, one for each player of the game, whereas an outcome is intuitively a state of affairs that comes about as the result of playing the strategy profile. Generally we assume that players are partial to the outcomes that may result from their actions, not so much to the actions themselves. Yet, if need be, outcomes may include a description of the acts through which they were brought about. Thus, the despicable or laudable character of certain acts can be accounted for in the desirability of their outcomes. Rather more frivolously, one could say that the strategies are the means and the outcomes the ends. Furthermore, it is important to note that, although each strategy profile has precisely one outcome, the same outcome may be brought about by different strategy profiles.

To illustrate the distinction between outcomes and strategies, consider the case of a burglar trying to open a safe that is secured by a unique seven digit code. If he succeeds in entering the right code the safe will open and the contents be his, otherwise an alarm will go off, forcing him to leave the building immediately without having another go at it. The burglar can enter any code he likes and the different codes he can enter are most naturally seen the strategies he can adopt. Yet his preferences are over the two possible outcomes of this situation, viz. whether the safe is open or remains closed after he has entered his code. By contrast, by means of which particular code the safe is secured is of no interest to him, e.g., he is quite indifferent as to whether the code is 4543360 or 3938785 as long as the code he enters opens the safe. Opening the safe cannot rightly be said to be a strategy available to the burglar, unaware as he is of the right code. Still, it could be said to be in the power of the burglar to open the safe, since entering the right code will do exactly that.

For many applications obscuring the distinction between outcomes and strategy profiles is quite harmless or even commendable. For the purposes of this paper, however, we cannot do enough to emphasize the distinction.

In some cases the preferred way to describe a situation of social interaction may be such that the powers of the players are formulated on a more abstract level in terms of outcomes rather than in terms of strategies. This may be the case, for example, when the relevant strategic information is lacking. This level of abstraction of representing the players' preferences, moreover, is in keeping with the employment of formal languages the syntax of which does not allow for explicit reference to strategic concepts

---

(1944), Luce and Raiffa (1957), Owen (1982), Myerson (1991) and Osborne and Rubinstein (1994) when reasoning about interactive situations. As examples of such frameworks should be mentioned ATL (e.g., Alur, Herzinger, and Kupferman, 2002), ATEL (e.g., van der Hoek and Wooldridge, 2003) and coalition logics (Pauly, 2001). Also in the field of *cooperative* game theory it is natural to abstract away from the strategic structure. The main concern in cooperative game theory is which outcomes a coalition can achieve by coordinating their strategies. By means of which strategies an outcome is achieved, however, is then quite irrelevant. The general issue this paper is concerned with is the extent to which the *noncooperative* solution concepts of game theory are still available when this higher level representation of the players' powers is adopted.

Nash equilibrium is a game theoretic solution concept the definition of which typically pertains to strategy profiles. A strategy profile is a Nash equilibrium whenever no player would benefit by unilaterally deviating from it. As such it gives content to a notion of strategic stability. In an equivalent fixed point formulation, a Nash equilibrium can be understood as combining strategies each of which is a best response with respect to the others.

The two aspects of a game that would thus seem to be relevant for singling out its Nash equilibria are the players' preferences and what each player can achieve by choosing a strategy or by deviating from a strategy profile. Usually, the preferences of the players are defined as a weak, i.e., transitive, reflexive and connected (complete), order over the possible outcomes of the game. The manipulative powers of the players in a game determine unequivocally a family of sets of subsets of outcomes. With each player is associated the set of outcomes he can guarantee the game to end in by choosing an appropriate strategy (cf., e.g., Moulin and Peleg, 1982; Abdou and Keiding, 1991). This set of sets of outcomes we will refer to as the player's *effectivity set*.[2]

We find, however, that it is not in general possible to find the outcomes that the Nash equilibria of a game give rise to, on the basis of the players' preferences and their individual, or even their coalitional, effectivity sets alone. This observation seems to reconfirm Nash equilibrium as an essentially strategic concept.

Finally, we suggest an alternative equilibrium concept, which we will call *outcome equilibrium*. The outcome equilibria can be singled out on the basis of the players' preferences and their coalitional effectivity sets alone. Yet no such result is possible if only the individual effectivity sets are taken into account.

---

[2] In the literature on coalitional NTU games, the effectivity of a coalition is also commonly defined as a set of payoff vectors rather than as sets of sets of outcomes (cf., e.g., Aumann, 1959, 1961a, 1967; Aumann and Peleg, 1960; Owen, 1982; Myerson, 1991). An important distinction is further between α- and β-effectivity. A coalition is said to be *α-effective* for a payoff profile v or a set of outcomes X, if its members have a joint strategy that guarantees each of them at least the payoffs specified in v or an outcome in X, no matter which strategies the players that are not part of the coalition may choose. By contrast, a coalition is β-effective for a payoff profile v or a set of outcomes X, if the players that are not part of the coalition cannot prevent the coalition to achieve a payoff profile as least as good as v or to have the game end in X. In this paper α-effectivity we assume α-effectivity, with the notable exception of Footnote 5, below.

2 Strategic Games and Effectivity

In this section we introduce some of the elementary concepts of game theory as well as Pauly's results concerning the characterization of effectivity sets.

In the introduction we emphasized the distinction between outcomes and strategy profiles. We make this distinction also formally explicit, and define a *strategic game* G as a tuple $(N, O, \{S_i\}_{i\in N}, \{\leq_i\}_{i\in N}, \omega)$, where N is a finite set of players, O is a finite set of outcomes,[3] and, for each player i, S_i is a finite set of strategies and $\leq_i$ is a reflexive, transitive and connected (complete) relation, or weak order, over O. These latter relations reflect the players' preferences with regard to the outcomes. We write $x <_i y$, if $x \leq_i y$ but not $y \leq_i x$. The set of *strategy profiles*, denoted by S, is given by the Cartesian product $\bigtimes_{i\in N} S_i$. In this paper, we do not consider mixed strategies, i.e., probability distributions over strategies. Finally, ω is a function mapping each strategy profile onto an outcome in O, i.e., $\omega\colon S \to O$. It be noted that formally the set of outcomes could be taken to be virtually anything. In particular the outcomes of a game could be understood to be utility vectors that assign to each player a numerical value or payoff.[4] Alternatively one could have the outcomes represent possible states of affairs, as they are known from possible world semantics for various logics for reasoning about multiagent systems. For technical convenience, we confine our attention to games with two or more players.

For i a player of a game G and X a subset of outcomes, we say that an outcome x in X is a *maximal* (*minimal*) *element* of X with respect to $\leq_i$ whenever there is no element y in X such that $x <_i y$ ($y <_i x$). Similarly, we have $\mathrm{Max}_i(X)$ ($\mathrm{Min}_i(X)$) denote the set of maximal (minimal) outcomes for i in X.

We also adopt the following notational convention. For C a subset of players and s and s' strategy profiles, we have (s_{-C}, s'_C) denote the strategy profile s'' such that, for all players i, s''_i equals s'_i, if i is in C, and s''_i equals s_i, otherwise. I.e., (s_{-C}, s'_C) is the strategy profile s'' such that for all players i:

$$s''_i \quad =_{df.} \quad \begin{cases} s'_i & \text{if } i \in C, \\ s_i & \text{otherwise.} \end{cases}$$

In case C is a singleton $\{i\}$, we generally omit the parentheses. Accordingly, (s_{-i}, s'_i) denotes the strategy profile $(s_1, \ldots, s_{m-1}, s'_m, s_{m+1}, \ldots, s_n)$, where s is the strategy profile given by $(s_1, \ldots, s_n)$ and i's strategy is the m-th element.

By a *strategic* or *noncooperative solution concept* we understand, quite generally and informally, any notion that singles out for each strategic game a, possibly empty, set of strategy profiles that are somehow significant from a game theoretic point of view. Perhaps the most important of solution concepts is *Nash equilibrium*. A strategy profile s is said to *contain a best response* for a player i whenever $\omega((s_{-i}, s'_i)) \leq_i \omega(s)$,

[3] The assumption that the set of outcomes be finite we use in proving Proposition 3, below. Mathematically more sophisticated and more liberal restrictions are quite possible, but do not really bear on the main tenets of this paper.

[4] Under the present definition the set of outcomes is assumed to be finite. Hence the set of outcomes cannot be taken to be the set of *all* real or integral utility vectors.

for all strategy profiles s'. A *Nash equilibrium* is then defined as a strategy profile that contains a best response for all players of the game.

Deviating but slightly from Pauly (2001), we represent the (individual) power of a player i in a game G by a set of sets of outcomes $\boldsymbol{E}_i^G$, also called the player's *individual effectivity* or *individual effectivity set*. Each set of outcomes in player i's effectivity set is such that i has at her disposal a strategy that guarantees the outcome of the game to be within that set. Which particular outcome within any such set actually comes about, however, may depend on the choices of the other players. Thus, for i a player of a game G, define $\boldsymbol{E}_i^G$ such that for each subset X of outcomes :

$$X \in \boldsymbol{E}_i^G \text{ iff there is an } s \text{ such that for all } s' : \omega((s'_{-i}, s_i)) \in X.$$

We say that a family $\{\boldsymbol{X}_i\}_{i \in N}$ of sets of subsets of outcomes *corresponds to the (individual) powers of the players in* G, if $\boldsymbol{X}_i$ equals $\boldsymbol{E}_i^G$, for each player i in N.

The concept of effectivity can straightforwardly and conservatively be extended as to apply to groups of players just as well as to individuals. Thus, each group of players C of a game G is associated with a set of subsets of outcomes, denoted by $\boldsymbol{E}_C^G$. For each group of players C, each element of $\boldsymbol{E}_C^G$ is a set of outcomes the players in C can force the game to terminate in by choosing suitable strategies independently of the choice of strategy made by the players outside the group. Thus, formally, for each set X of outcomes and each subset of players C, we have:

$$X \in \boldsymbol{E}_C^G \text{ iff there is an } s \text{ such that for all } s' : \ \omega((s'_{-C}, s_C)) \in X.$$

Having $-i$ denote the group of all players except i, then in particular:

$$X \in \boldsymbol{E}_{-i}^G \text{ iff there is an } s \text{ such that for all } s' : \ \omega((s_{-i}, s'_i)) \in X.$$

Also observe that, in general, $\boldsymbol{E}_{\{i\}}^G$ equals $\boldsymbol{E}_i^G$ and that each player's and each coalition's effectivity set is closed under supersets. We say that a family $\{\boldsymbol{X}_C\}_{C \subseteq N}$ of sets of subsets of outcomes *corresponds to the coalitional powers in* G, if $\boldsymbol{X}_C$ equals $\boldsymbol{E}_C^G$, for each group of players C of N. It be observed that it is not in general the case that the effectivity of a coalition is uniquely determined by the individual effectivity of its members.

We are now in a position to define the notions of an *individual effectivity game* and a *coalitional effectivity game*. We define an *individual effectivity game* as a tuple $(N, O, \{\boldsymbol{X}_i\}_{i \in N}, \{\leq_i\}_{i \in N})$, where N, O and $\{\leq_i\}_{i \in N}$ are as in a strategic game and $\{\boldsymbol{X}_i\}_{i \in N}$ is a set of subsets of outcomes that corresponds to the powers of the players in some strategic game with N, O and $\{\leq_i\}_{i \in N}$ as players, outcomes and preferences respectively. A *coalitional effectivity game* is defined analoguously as a tuple $(N, O, \{\boldsymbol{X}_C\}_{C \subseteq N}, \{\leq_i\}_{i \in N})$, such that $\{\boldsymbol{X}_C\}_{C \subseteq N}$ corresponds to the coalitional powers in some strategic game. We say an individual (coalitional) effectivity game E *corresponds* to a strategic game G, whenever both are defined on the same set of players and outcomes, the preferences over the outcomes of each player in E coincide with those in G and the individual (coalitional) effectivity sets in E correspond to the individual (coalitional) powers in G.

Pauly (2001) successfully addresses the issue of which set theoretic conditions a set of subsets of outcomes has to satisfy, if it is to constitute a family of effectivity sets for the players of a game. Moreover, he achieved a much similar result for an analogous notion of coalitional effectivity sets, i.e., the sets of subsets of outcomes a group of players can guarantee a game to end in by choosing appropriate strategies. Thus, to appraise whether a family of sets of subsets of outcomes corresponds to a family of effectivity sets, no *strategic* notions need to be invoked as such. This makes that effectivity games could also be defined without explicit reference to strategic games. Furthermore, we say that a coalitional effectivity game $(N, O, \{\boldsymbol{X}_C\}_{C\subseteq N}, \{\leq_i\}_{i\in N})$ *extends* an individual effectivity game $(N, O, \{\boldsymbol{Y}_i\}_{i\in N}, \{\leq_i\}_{i\in N})$ whenever $\boldsymbol{X}_{\{i\}} = \boldsymbol{Y}_i$, for each $i \in N$.

Representing a social situation as a strategic game it is possible to single out the strategy profiles, if any, that comply with a noncooperative solution concept. Via the outcome function ω of the strategic game, with this set of strategy profiles can be associated a set of outcomes. The issue we are concerned with is whether the same outcomes, viz. those associated with strategy profiles complying with a solution concept, can be singled out given the more abstract representation of the same situation as an effectivity game. In particular, we will address this issue for Nash equilibrium. For this purpose we introduce the notion of a solution concept being *defined* for an effectivity game, either individual or coalitional. A noncooperative solution concept κ is said to be *defined* for an individual (coalitional) effectivity game E, if the set of outcomes elicited by κ is the same in any strategic game to which E corresponds, i.e., if the set $\{\omega(s) : s \text{ complies with } \kappa \text{ in } G\}$ equals the set $\{\omega'(s) : s \text{ complies with } \kappa \text{ in } G'\}$, for any strategic games G and G' to which E corresponds. In general, we say that a solution concept κ is defined for individual (coalitional) effectivity games, if κ is defined for *all* individual (coalitional) effectivity games. An immediate consequence of this definition is that, if a solution concept is defined for an individual effectivity game, then also for any coalitional or individual effectivity game that extends it. In this paper we thus address the issue whether Nash equilibrium is in general defined for either individual or coalitional effectivity games.

3 Nash Equilibria in Effectivity Games

The set of Nash equilibria of a strategic game are those strategy profiles that contain a best response for each player. Nevertheless, it is perfectly possible for an outcome that it is not yielded by any Nash equilibrium of a game, even though for each player i, it is the outcome of some strategy profile containing a best response for i. I.e., it is not in general the case that the subset of outcomes of a strategic game G given by:

$$\bigcap_{i\in N} \{\omega(s) : \ s \text{ contains a best response for } i \text{ in } G\}$$

equals the subset of outcomes given by:

$$\{\omega(s) : \ s \text{ is a Nash equilibrium in } G\}.$$

This phenomenon is illustrated by the two-player game $G1$ depicted in Figure 1. The one player, *Row*, chooses rows, whereas the other player, *Col*, chooses columns, determining one of the entries of the matrix as the outcome of the game. Let us assume that both players prefer c to outcome b and value a least of all. The outcomes associated with the strategy profiles containing a best response for *Row*, as well as those containing a best response for *Col*, are then given by the set $\{b, c\}$. Hence, the intersection of the sets of outcomes yielded by a strategy profile containing a best response for one of the players coincides with $\{b, c\}$ as well. The only Nash equilibrium of this game is the strategy profile in which *Row* chooses the bottom row and *Col* opts for the right column. This strategy profile, however, yields c as the unique outcome.

$$\begin{pmatrix} a & b \\ b & \boldsymbol{c} \end{pmatrix}$$

Fig. 1. The two-player game $G1$ in which the intersection of the outcomes containing a best response does not coincide with the set of outcomes yielded by the Nash equilibria. Assume that for i either of the players, $a <_i b <_i c$. The only Nash equilibrium outcome, c, is in boldface.

This example, of course, merely shows that outcomes and strategy profiles belong to different categories. Two strategy profiles may have very different strategic properties and still give rise to the same outcome in a game. In a sense outcomes could thus be said to abstract from the strategic features of the underlying strategy profiles. The above example demonstrates, moreover, that by focussing on outcomes these strategic aspects of strategy profiles may be lost beyond recall.

$$\begin{pmatrix} a & \boldsymbol{b} & b \\ b & b & \boldsymbol{c} \\ a & a & b \end{pmatrix}$$

Fig. 2. The two-player game $G2$. Assuming the players' preferences as in $G1$, the two Nash equilibrium outcomes, c as well as b, are in boldface.

The effecitivity sets as they were introduced in the previous section are composed of sets of subsets of outcomes. The question we will be concerned with at this point is whether the structure of a game G that is preserved in a family $\{\boldsymbol{E}_i^G\}_{i \in N}$ together with the preferences of the players suffice to single out precisely the set of outcomes yielded by the Nash equilibria of G. Or, to phrase the issue slightly differently, consider a family $\{\boldsymbol{X}_i\}_{i \in N}$ which corresponds to the individual powers of a strategic game G. The point that is here at stake is whether it is possible to determine the set of Nash equilibrium outcomes of the game G? I.e., is Nash equilibrium defined for individual effectivity games, and if not, whether it is defined for coalitional effectivity games.

We find that the answer to both these questions is to be negative. To appreciate this, consider the two-person game $G2$ depicted in Figure 2, where the players and their preferences are as in the game $G1$ of the previous example, i.e., *Row* chooses rows, *Col* chooses columns and both players prefer c to both a and b, and b to a. Moreover, the individual effectivity sets of both players in $G2$ coincide with one another just as well as with those in $G1$, i.e., we have:

$$\boldsymbol{E}^{G1}_{Row} \;=\; \boldsymbol{E}^{G1}_{Col} \;=\; \boldsymbol{E}^{G2}_{Row} \;=\; \boldsymbol{E}^{G2}_{Col} \;=\; \{\,\{a,b\}\,,\{b,c\}\,,\{a,b,c\}\,\}.$$

Hence, the same families of effectivity sets correspond to the individual powers in both $G1$ and $G2$ and the players' preferences over the outcomes are also the same in both games. Observe, however, that the game $G2$ has two Nash equilibria, whereas $G1$ has only one. Moreover, one of these Nash equilibria—viz. the strategy profile in which *Row* chooses the top row and *Col* the middle column—gives rise to a different outcome (viz. outcome b) than the one Nash equilbrium in $G1$ (viz. outcome c). We may conclude that in this case it is impossible to single out the Nash equilibrium outcomes on the basis of the effectivity and preferences of the players alone, i.e., that Nash equilibrium is not defined for effectivity games.

Also observe that things would have been no better, if we had considered the collective instead of the individual effectivity sets. In both $G1$ and $G2$ the effectivity of the grand coalition of both players is given by $2^{\{a,b,c\}} - \{\emptyset\}$ and that of the empty coalition by $\{\,\{a,b,c\}\,\}$. Furthermore, apart from the singleton coalitions, there are no other coalitions possible in a two-player game. Hence, the same pair of games also presents a counter example against the claim that Nash equilibrium is defined for coalitional effectivity games. The following proposition summarizes these negative results.

Proposition 1. *Nash equilibrium is defined for neither individual nor coalitional effectivity games.*

In the above counterexample the Nash outcomes of the games $G1$ and $G2$ were distinct, yet the Nash outcomes of $G1$ constituted a subset of those of $G2$. This is not necessarily so for strategic games to which the same effectivity game corresponds. We close this section with an example of an effectivity game that corresponds to two strategic games with *disjoint* sets of Nash outcomes. Consider the two strategic games in Figure 3 and let the preferences of the two players be given by $a \leq_{Row} b \leq_{Row} x \leq_{Row} y \leq_{Row} c$ and $a \leq_{Col} c \leq_{Col} x \leq_{Col} y \leq_{Col} b$. It can be readily recognized that $G3$ and $G4$ correspond to the same effectivity game, both individual and coalitional. Nevertheless, x is the only Nash outcome in $G3$, whereas y is the only Nash outcome in $G4$.

4 Outcome Equilibrium

To counterbalance the negative results of the previous section, in this section we consider a solution concept for which the collective effectivity and preference sets do suffice

$$\begin{pmatrix} \mathbf{x} & a & a & c \\ a & x & a & b \\ a & b & y & a \\ a & c & a & y \end{pmatrix} \qquad \begin{pmatrix} x & a & c & a \\ a & x & b & a \\ b & a & y & a \\ c & a & a & \mathbf{y} \end{pmatrix}$$

Fig. 3. Two 4×4 games in strategic form, $G3$ and $G4$. The preferences of the players are given by $a \leq_{Row} b \leq_{Row} x \leq_{Row} y \leq_{Row} c$ and $a \leq_{Col} c \leq_{Col} x \leq_{Col} y \leq_{Col} b$. In the leftmost game the only Nash outcome is x, whereas in the rightmost game the only Nash outcome is y. Nevertheless, the effectivity of both players *Row* and *Col* are identical.

$$\begin{pmatrix} a & e & d \\ c & d & c \\ b & a & e \end{pmatrix}$$

Fig. 4. The two-player game *G*5.

to single out its instances. The fundamental idea is to replace the concept of a strategy containing a best response by a notion of a player's satisfaction with a particular outcome. Formally we say that an outcome *o* is *satisfactory for a player i* in a game *G*, whenever $\omega(s) \leq_i o$ for some strategy profile *s* that contains a best response for *i* in *G*.The intuition behind this concept is that a player should be able to reconcile himself with any outcome that is no worse than his best effort could have achieved in the least favorable circumstances, i.e., if his opponents had been conspiring against him and he still had managed to avert personal disaster. An *outcome equilibrium* of a game *G* we then define as a strategy profile that yields a satisfactory outcome to all players of *G*. Obviously, every Nash equilibrium is also an outcome equilibrium but not necessarily the other way round.

As an illustration of this concept, consider the two-player game *G*5 in Figure 4. Let in this case the preferences of the players, *Row* and *Col*, be antagonistic. In particular we assume $a <_{Row} b <_{Row} c <_{Row} d <_{Row} e$. Accordingly, for *Col* we have $e <_{Col} d <_{Col} c <_{Col} b <_{Col} a$. The strategy profiles containing a best response to *Row*, as can easily be checked, result in either outcome *c* or outcome *e*, of which he prefers *c* least. Hence, all of the outcomes at least as desirable as *c*—viz. the outcomes *c*, *d* and *e*—are satisfactory for *Row*. In a similar fashion it can be established that the outcomes satisfactory for *Col* are *a*, *b* and *c*. So, in this case, *c* is the only outcome elicited by an outcome equilibrium in *G*5. Observe in this context that, if *d* were eventually to emerge as the outcome of the game, *Row* can be certain not to have played a best response strategy. However, since he could have been worse off had he played a best response against *Col*'s choosing the left column, *d* still qualifies as a satisfactory outcome for *Row*.

$$\begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix} \quad \begin{pmatrix} b & c & a \\ a & b & c \\ c & a & b \end{pmatrix} \quad \begin{pmatrix} c & a & b \\ b & c & a \\ a & b & c \end{pmatrix}$$

Fig. 5. The three-player game *G*6, in which *c* is the best, *b* the second best and *a* the worst outcome for each of the players. Its only outcome equilibrium is then *c*.

But how much better than Nash equilibrium does outcome equilibrium fare with respect to its characterization in terms of players' preferences and effectivity? We find that the answer to this should be negative in case only individual effectivity is taken into account, whereas a characterization is possible, if the richer structure of coalitional effectivity may be invoked.

As to the first, negative, claim, consider the games *G*6 and *G*7 depicted in Figures 5 and 6, respectively. Both games involve three players, *Row*, *Col*, and *Mat*. The first two

$$\begin{pmatrix} c & b & a \\ a & b & c \\ c & b & a \end{pmatrix} \quad \begin{pmatrix} b & c & a \\ b & b & b \\ c & a & b \end{pmatrix} \quad \begin{pmatrix} c & a & b \\ a & b & c \\ a & b & c \end{pmatrix}$$

Fig. 6. Another three player game, $G7$. The players' preferences are as in $G6$. Here, apart from c, also b is an outcome equilibrium.

choose rows and columns, as before, and *Mat* chooses matrices. Observe that in both games the individual effectivity of all the players is identical:

$$\boldsymbol{E}^{G6}_{Row} = \boldsymbol{E}^{G6}_{Col} = \boldsymbol{E}^{G6}_{Mat} = \boldsymbol{E}^{G7}_{Row} = \boldsymbol{E}^{G7}_{Col} = \boldsymbol{E}^{G7}_{Mat} = \{\{a, b, c\}\}.$$

Let, moreover, the preferences of the players over the outcomes in both games coincide, each preferring c to b, with a being valued least of all. Then, c is the only outcome yielded by an outcome equilibrium in $G6$. However, in $G7$ not only c but also b qualifies as an outcome elicited by an outcome equilibrium. Hence the following proposition.

Proposition 2. *Outcome equilibrium are not in general defined for individual effectivity games.*

By contrast, the following proposition bears out the positive side of the issue.

Proposition 3. *Outcome equilibrium is defined for coalitional effectivity games.*

Proof (sketch). Let E be a coalitional effectivity game given by $(N, O, \{\boldsymbol{X}_C\}_{C\subseteq N}, \{\leq_i\}_{i\in N})$. Define $X^* =_{df.} \bigcap_{i\in N} \bigcup_{Y\in \boldsymbol{X}_{-i}} \mathrm{Max}_i(Y)$. It can then be demonstrated that X^* equals the set of outcomes yielded by the outcome equilibria of E.[5] □

[5] In connection with the concept of outcome equilibrium and the proof of Proposition 3, there is an interesting connection with the β-core as it was proposed in Aumann (1959, 1961a) in the setting of the theory of cooperative games without side payments. On the basis of the β-effectivity of coalitions (cf., Footnote 2, above) a dominance relation on the outcomes can be defined as follows. An outcome a is said to β-dominate another outcome b *via* a coalition C if (a) $b <_i a$ for all members i of C and (b) there is some set of outcomes X for which C is β-effective and $x \leq_i a$, for all $x \in X$. The β-core then consists precisely of those outcomes that are not β-dominated by any other outcome *via* any coalition. It can now be proved that an outcome equilibrium is an outcome that is not β-dominated *via* any *one-player* coalition. As the outcome sets for which coalitions are β-effective can be retrieved from coalitional effecivity sets $\boldsymbol{E}^G_C$ as they were defined in this paper, i.e., as α-effectivity, this observation provides an alternate proof of Proposition 3. It be noted in this context, that with outcome equilibrium we merely wished to illustrate how a solution concept akin to Nash equilibrium and *similarly defined in strategic terms*, can be characterized using coalitional effectivity, whereas Nash equilibrium cannot. We do not pretend to have introduced an entirely new solution concept in its own right or claim that it is somehow superior to either Nash equilibrium or the β-core. The authors are indebted to Moshe Tennenholtz for pointing out the connection between outcome equilibrium and the β-core.

5 Conclusion

Many existing game theoretic solution concepts, such as Nash equilibrium, pertain to strategy profiles rather than to outcomes. We have argued that the outcomes generated by the Nash equilibria of a game cannot in general be determined in terms of the outcomes the players can force the game to end in, i.e., their effectivity, and the player's preferences alone. Apparently, the solution concept relies too much on the strategic structure of games as given by the strategy profiles. Therefore, we have proposed an alternative game theoretic solution concept, called outcome equilibrium, which can be defined in terms of the players' coalitional effectivity together with the players' preferences.

However, outcome equilibrium shares a number of supposedly nocuous properties with Nash equilibrium. An outcome equilibrium is not guaranteed to exist and even if it does exist, it is not generally unique. Moreover, it is not in general the case that, for s and s' strategy profiles and i a player, if $\omega(s)$ and $\omega(s')$ are outcome equilibria, so are $\omega((s_{-i}, s'_i))$ and $\omega((s'_{-i}, s_i))$. Also, outcome equilibria need not be *Pareto efficient* in the sense that s may be an outcome equilibrium with there still being available another strategy profile s' such that each player of the game prefers $\omega(s')$ to $\omega(s)$. In any such case, however, s' will be an outcome equilibrium as well.

Studying the strategic aspects of an interactive situation by representing it as an effectivity game raises a number of questions. Obviously, the analysis can be extended to other noncooperative solution concepts, such as individually rational outcomes, securing players their security level, and weakly and strongly dominant strategies. It will then appear that individually rational outcomes are defined for both individual and coalitional effectivity games whereas neither weakly nor strongly dominant strategies are defined for either.

Thus we find that Nash equilibrium and strategic dominance require a more detailed representation of a situation of social interaction than do outcome equilibrium or the players' security level. This suggests a classification of game theoretic concepts, depending on the level specificity with which the strategic structure of a situation of social interaction should be represented for them to become available. The strategic form is a rather elaborate model of a situation of social interaction and often quite unattainable in practical settings. Thus, one could come to consider more austere models, which abstract from more structure. An urgent question could then be which game theoretic concepts would still be at one's disposal. The type of issue this paper addresses could of course be varied upon for different models of interaction and different game theoretical concepts.

Acknowledgements

The authors thank Jelle Gerbrandy for his useful comments on an earlier version of this paper. Parts of this work were supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/3-1.

References

Abdou, J. and Keiding, H. (1991), *Effectivity Functions in Social Choice*. Volume 8 of *Theory and Decision Library*. Kluwer Academic Publishers, Dordrecht, Boston, London.

Alur, R., Herzinger, T. A. and Kupferman, O. (2002), "Alternating-Time Temporal Logic." *Journal of the ACM* **49**, 672–713.

Aumann, R. J. (1959), "Acceptable Points in General *n*-Person Games." in: Luce, R. D. and Tucker, A. W. (eds.), *Contributions to the Theory of Games IV*, Annals of Mathematical Studies 40, pp. 287–324. Princeton U.P., Princeton.

——— (1961a), "The Core of a Cooperative Game without Side Payments." *Transactions of the American Mathematical Society* **98**, 539–552.

——— (1961b), "The Core of a Cooperative Game without Side Payments." *Transactions of the American Mathematical Society* **98**, 539–552.

——— (1967), "A Survey of Cooperative Games without Side Payments." in: Shubik, M. (ed.), *Essays in Mathematical Economics in Honor of Oskar Morgenstern*, pp. 3–27. Princeton University Press.

——— (1987), "Game Theory." in: Eatwell, J., Milgate, M. and Newman, P. (eds.), *Game Theory*, The New Palgrave, pp. 1–54. Macmillan, London and Basingstoke.

Aumann, R. J. and Peleg, B. (1960), "Von Neumann-Morgenstern Solutions to Cooperative Games without Side Payments." *Bulletin of the American Society* **66**, 173–179.

Benthem, J. F. A. K. van (2001), *Logic and Games. Lecture Notes*. ILLC, Amsterdam.

Hoek, W. van der and Wooldridge, M. (2003), "Cooperation, Knowledge and Time: Alternating-time Temporal Epistemic Logic and its Applications." *Studia Logica* **77** (3), 125–157.

Luce, R. D. and Raiffa, H. (1957), *Games and Decisions*. John Wiley & Sons, New York.

Moulin, H. and Peleg, B. (1982), "Cores of Effectivity Functions and Implementation Theory." *Journal of Mathematical Economics* **10**, 115–145.

Myerson, R. B. (1991), *Game Theory. Analysis of Conflict*. Harvard U.P., Cambridge, Mass.

Nash, J. (1950), "Equilibrium Points in *n*-Person Games." *PNAS* **36**, 48–49.

——— (1951), "Non-Cooperative Games." *Annals of Mathematics* **54**, 286–295.

Neumann, J. von and Morgenstern, O. (1944), *Theory of Games and Economic Behavior*. Princeton U.P., Princeton, N.J. (2nd edition 1947).

Osborne, M. J. and Rubinstein, A. (1994), *A Course in Game Theory*. MIT Press, Cambridge, Mass.

Owen, G. (1982), *Game Theory*. Academic Press, 2nd edition.

Pauly, M. (2001), *Logic for Social Software*. Ph.D. thesis, Institute for Logic, Language and Information, Amsterdam.

Peleg, B. (1984), *Game Theoretic Analysis of Voting in Committees*. Cambridge U.P.