Dependence Theory via Game Theory

Davide Grossi
Institute for Logic, Language and Computation; University of Amsterdam
Science Park 904, 1098 XH
Amsterdam, The Netherlands
d.grossi@uva.nl

Paolo Turrini
Department of Information and Computing Sciences; Utrecht University
Padualaan 14, 3584CH
Utrecht, The Netherlands
paolo@cs.uu.nl

ABSTRACT

In the multi-agent systems community, dependence theory and game theory are often presented as two alternative perspectives on the analysis of social interaction. Up till now no research has been done relating these two approaches. The unification presented provides dependence theory with the sort of mathematical foundations which still lacks, and shows how game theory can incorporate dependence-theoretic considerations in a fully formal manner.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems*

General Terms

Economics, Theory

Keywords

Dependence theory, Game theory

1. INTRODUCTION

The present paper brings together two independent research threads in the study of social interaction within multiagent systems (MAS): game theory [16] and dependence theory [4]. Game theory was born as a branch of economics but has recently become a well-established framework in MAS [14], as well as in computer science in general [9]. Dependence theory was born within the social sciences [5] and brought into MAS by [4]. Although not having uniform formulation, and still lacking thorough formal foundations, its core ideas made their way into several researches in MAS (e.g., [1,15]).

The paper moves from the authors' impression that, within the MAS community, dependence theory and game theory are erroneously considered to be alternative, when not incompatible, paradigms for the analysis of social interaction. An impression that has recently been reiterated during the AAMAS'2009 panel discussion "Theoretical Foundations for Agents and MAS: Is Game Theory Sufficient?". It is our conviction that the theory of games and that of dependence are highly compatible endeavours. By relating the two theories in a unified formal setting, the following research objectives are met: i) dependence theory can tap into the highly developed mathematical framework of game theory, obtaining the sort of mathematical foundations that it is still missing; ii) game theory can incorporate a dependence-theoretic perspective on the analysis of strategic interaction.

With respect to this latter point, the paper shows that dependence theory can play a precise role in games by modeling a way in which cooperation arises in strategic situations:

> "As soon as there is a possibility of choosing with whom to establish parallel interests, this becomes a case of choosing an ally. [...] One can also state it this way: A parallelism of interests makes a cooperation desirable, and therefore will probably lead to an agreement between the players involved." [16, p. 221]

Once this intuitive notion of "parallelism of interests" is taken to mean "mutual dependence" [4] or "dependence cycle" [15] the bridge is laid and the notion of agreement that stems from it can be fruitfully analyzed in dependence-theoretic terms.

The paper is structured as follows. Section 2 briefly introduces dependence theory and some notions and terminology of game theory. The section points also at the few related works to be found in the literature. Section 3 gives a game-theoretic semantics to statements of the type "i depends on j for achieving an outcome s." Most importantly, the section relates cycles and equilibrium points in games (Theorem 1) giving a formal game-theoretic argument for the centrality of cycles in dependence theory. Section 4 connects cycles to the possibility of agreements among players and introduces a way to order them. Maximal agreements in such ordering are, in Section 4.3, related to the core of corresponding coalitional games where, following the intuition of the above quote, coalitions form from dependence cycles via agreements (Theorem 2). Conclusions follow in Section 5.

2. PRELIMINARIES

The section is devoted to the introduction of the conceptual and technical apparatus that will be dealt with in the paper. Relevant and related literature is also pointed at.

2.1 Dependence theory in a nutshell

Dependence theory has, at the moment, several versions and no unified theory. There are mainly informal [3,4] and a few formal and computational [1,7,15] accounts. Yet, the aim

Cite as: Dependence Theory via Game Theory, Davide Grossi and Paolo Turrini, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.

of the theory is clear, and is best illustrated by the following quote:

> "One of the fundamental notions of social interaction is the *dependence* relation among agents. In our opinion, the terminology for describing interaction in a multi-agent world in necessarily based on an analytic description of this relation. Starting from such a terminology, it is possible to devise a calculus to obtain predictions and make choices that simulate human behaviour" [4, p. 2].

So, dependence theory boils down to two issues: i) the identification (and representation) of dependence relations among the agents in a system (to the nature of this relation we will come back in Section 3); ii) the use of such information as a means to obtain predictions about the behaviour of the system. While all contributions to dependence theory have focused on the first point, the second challenge, "devise a calculus to obtain predictions", has been mainly left to computer simulation [7]. When, in Section 1, we talked about providing formal foundations to dependence theory, we meant precisely this: find a suitable formal framework within which dependence theory can be given analytical predictive power. The paper shows in Section 4 that game theory can be used for such a purpose.

2.2 A game-theory toolkit

We sketch some of the game theoretic notions we will be dealing with in the paper. The reader is referred to [10] for a more extensive exposition.

Definition 1 (Game). *A (strategic form) game is a tuple $\mathcal{G} = (N, S, \Sigma_i, \succeq_i, o)$ where: N is a set of players; S is a set of outcomes; Σ_i is a set of strategies for player $i \in N$; $\succeq_i$ is a total preorder*[1] *over S (its irreflexive part is denoted $\succ_i$); $o : \bigtimes_{i \in N} \Sigma_i \rightarrow S$ is a bijective function from the set of strategy profiles to S.*[2]

Examples of games represented as payoff matrices are given in Figure 1. A strategy profile will be denoted σ, while strategies will be denoted as i-th projections of profiles, so σ_i denotes an element of Σ_i. Following the usual convention, $\sigma_C = (\sigma_i)_{i \in C}$ denotes the strategy $|C|$-tuple of the set of agents $C \subseteq N$. Given a strategy profile σ and an agent i, we call an i-variant of σ any profile which differs from σ at most for σ_i, i.e., any profile (σ'_i, σ_{-i}) with σ' possibly different from σ, where $-i = N \backslash \{i\}$. Similarly, a C-variant of σ is any profile $(\sigma'_C, \sigma_{\overline{C}})$ with σ' possibly different from σ where $\overline{C} = N \backslash C$. It is assumed that $(\sigma_N, \sigma_\emptyset) = (\sigma_N) = (\sigma_\emptyset, \sigma_N)$.

As to the solution concepts, we will work with Nash equilibrium, which we will refer to also as best response equilibrium (*BR*-equilibrium), and the dominant strategy equilibrium (*DS*-equilibrium).

Definition 2 (Equilibria). *Let $\mathcal{G}$ be a game. A strategy profile σ is: a BR-equilibrium (Nash equilibrium) iff $\forall i, \sigma' : o(\sigma) \succeq_i o(\sigma'_i, \sigma_{-i})$; it is a DS-equilibrium iff $\forall i, \sigma' : o(\sigma_i, \sigma'_{-i}) \succeq_i o(\sigma')$.*

[1] A total preorder is a transitive and total relation. Recall that a total relation is also reflexive.

[2] The definition could be dispensed with the outcome function. However, we chose for this presentation because it eases the formulation of the results presented in Section 3.4. If $o(\sigma) = s$, we will however use σ and s interchangeably.

So, a *BR*-equilibrium is a profile where all agents play a best response and a *DS*-equilibrium is a profile where all agents play a dominant strategy.

In addition to the games in strategic form (Definition 1) we will also work with coalitional games, i.e., cooperative games with non-transferable pay-offs abstractly represented by effectivity functions. In this case our main references are [8, 11].

Definition 3 (Coalitional game). *A coalitional game is a tuple $\mathcal{C} = (N, S, E, \succeq_i)$ where: N is a set of players; S is a set of outcomes; E is function $E : 2^N \rightarrow 2^{2^S}$; $\succeq_i$ is a total preorder on S.*

Function E—effectivity function—assigns to every coalition the sets of states that the coalition is able to enforce. In coalitional games, the standard solution concept is the core.

Definition 4 (The Core). *Let $\mathcal{C} = (N, S, E, \succeq_i)$ be a coalitional game. We say that a state s is dominated in $\mathcal{C}$ if for some C and $X \in E(C)$ it holds that $x \succ_i s$ for all $x \in X, i \in C$. The core of $\mathcal{C}$, in symbols $CORE(\mathcal{C})$ is the set of undominated states.*

Intuitively, the core is the set of those states in the game that are stable, i.e., for which there is no coalition that is at the same time able and interested to deviate from them.

2.3 Games and dependencies in the literature

To the best of our knowledge, almost no attention has been dedicated up till now to the relation between game theory and dependence theory, with two noteworthy recent exceptions: [2], and [13] which is the last contribution of a line of work starting with [1]. Although mainly motivated by the objective of easing the computational complexity of computing solution concepts in Boolean games [6], those works pursue a line of research that has much in common with ours. They study the type of dependence relations arising within Boolean games and they relate them to solution concepts such as Nash equilibrium and the core. In the present paper we proceed in a similar fashion, although our main concern is, instead of the application of dependence theory to the analysis of games, the use of game theory as a formal underpinning for dependence theory. As a consequence, our analysis needs to shift from Boolean games to standard games.

3. DEPENDENCIES IN GAMES

It is now time to move to the presentation of our results. The present section takes the primitive of dependence theory, i.e., the relation "i depends on j for achieving goal g", and defines it in game-theoretic terms.

3.1 Dependence as 'need for a favour'

Let us start off with a classic example [10].

Example 1 (dilemma). *Consider the Prisoner's dilemma payoff matrix in Figure 1. As is well-know the outcome (D, R) is a dominant strategy equilibrium (and thus a Nash equilibrium). To achieve this outcome in the game, it is fair to say that neither Row depends on Column, nor vice versa. Players just play their dominant strategies, they owe nothing to each other. What about outcome (D, L)? Here the situation is clearly asymmetric. If this outcome were to be achieved, Row would depend on Column in that, while Row plays its dominant strategy, Column has to play a dominated one which maximizes Row's welfare. Same analysis, with roles flipped, holds for (U, R). Asymmetry is broken in (U, L). If*

	L	R
U	2,2	0,3
D	3,0	1,1

Prisoner's dilemma

	L	R
U	3,3	2,2
D	2,2	1,1

Full Convergence

	L	R
U	1,1	0,0
D	0,0	1,1

Coordination

	L	R
U	3,3	2,2
D	2,5	1,1

Partial Convergence

Figure 1: Two players strategic games.

this was the outcome to be selected, as we might expect, Row would depend on Column and Column on Row since they both would have to play a dominated strategy which maximizes the opponent's welfare.

The literature on dependence theory features a number of different relations of dependency. Yet, in its most essential form, a dependence relation is a relation occurring between two agents i and j with respect to a certain state (or goal) which i wants to achieve but which it cannot achieve without some appropriate action of j.

> "x depends on y with regard to an act useful for realizing a state p when p is a goal of x's and x is unable to realize p while y is able to do so." [4, p.4]

By taking a game theoretic perspective, this informal relation acquires a precise and natural meaning: a player i depends on a player j for the realization of a state p, i.e., of the strategy profile σ such that $o(\sigma) = p$, when, in order for σ to occur, j has to favour i, that is, it has to play in i's interest. To put it otherwise, i depends on j for σ when, in order to achieve σ, j has to do a favour to i by playing σ_j (which is obviously not under i's control).[3] This intuition is made clear in the following definition.

Definition 5 (Best for someone else). *Assume a game $\mathcal{G} = (N, S, \Sigma_i, \succeq_i, o)$ and let $i, j \in N$. 1) Player j's strategy in σ is a* best response for *i iff $\forall \sigma'$, $o(\sigma) \succeq_i o(\sigma'_j, \sigma_{-j})$. 2) Player j's strategy in σ is a* dominant strategy for *i iff $\forall \sigma'$, $o(\sigma_j, \sigma'_{-j}) \succeq_i o(\sigma')$.*

Definition 5 generalizes the standard definitions of best response and dominant strategy by allowing the player holding the preference to be different from the player whose strategies are considered. By setting $i = j$ we obtain the usual definitions. We are now in the position to mathematically define the notion(s) of dependence as game theoretic notions.

Definition 6 (Dependence). *Let $\mathcal{G} = (N, S, \Sigma_i, \succeq_i, o)$ be a game and $i, j \in N$. 1) Player i* BR-depends *on j for strategy σ—in symbols, $iR^{BR}_{\sigma}j$—if and only if σ_j is a best response for i in σ. 2) Player i* DS-depends *on j for strategy σ—in symbols, $iR^{DS}_{\sigma}j$—if and only if σ_j is a dominant strategy for i.*

Intuitively, i depends on j for profile σ in a best response sense if, in σ, j plays a strategy which is a best response for i given the strategies in σ_{-j} (and hence given the choice of i itself). Similarly, i depends on j for profile σ in a dominant strategy

[3]It might be worth noticing that while the notion of dependence relation we use is a three-place one, in dependence theory, it often has higher arity, incorporating ingredients such as plans and actions (e.g. [15]).

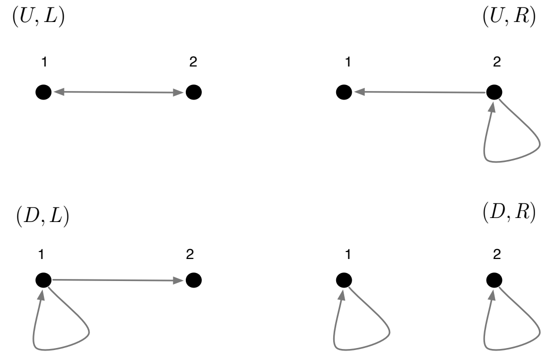

Figure 2: BR-dependences in the Prisoner's dilemma (1 denotes Row and 2 denotes Column).

sense if j's strategy in σ is the best j can do for maximizing i's welfare. We have thus obtained a notion of dependence with a clear game theoretic meaning. Note that if $iR^{DS}_{\sigma}j$ then $iR^{BR}_{\sigma}j$, but not vice versa, as a dominant strategy for i is always a best response for i.

Therefore, with any game $\mathcal{G}$ two dependence structures (N, R^{BR}_{σ}) and (N, R^{DS}_{σ})[4] can be associated where $\sigma \in \bigtimes_{i \in N} \Sigma_i$. Such structures, the one based on the notion of best response and the other on the notion of dominant strategy, are multigraphs on the finite carrier N and containing a finite number of binary relations, one for each possible profile of the game $\mathcal{G}$. They are a game theoretic formalization of the informal notion of dependence described above. Figure 2 depicts the BR-dependence structure of the game discussed in Example 1. So, for instance, the relation $R^{BR}_{(U,L)}$ depicted in the up-right corner of Figure 2 is such that Column depends on itself (it is a reflexive point), as it plays its own best response, and on Row, as Row does not play its own best response but a best response for Column.

In the Prisoner's dilemma the R^{BR}_{σ} and R^{DS}_{σ} relations coincide, that is, the BR- and DS-dependencies are equivalent. It should be clear that this is not always the case, as the reader can notice by considering the Coordination game in Figure 1. In such game no dominant strategy exists for any player with respect to any player.

3.2 Relational vs. game-theoretic properties

In general, relations R^{BR}_{σ} and R^{DS}_{σ} do not enjoy any particular structural property (e.g., reflexivity, completeness, transitivity, etc.). In the present section we study under what conditions they come to enjoy particular properties. We consider the two properties of reflexivity and universality,[5] which will be of use later in the paper.

Fact 1 (Properties of dependence relations). *Let $\mathcal{G}$ be a game and (N, R^{x}_{σ}) be its dependence structure with $x \in \{BR, DS\}$. For any profile σ it holds that:*

1. *R^{BR}_{σ} is reflexive iff σ is a BR-equilibrium;*
2. *R^{DS}_{σ} is reflexive iff σ is a DS-equilibrium;*

[4]We use this lighter notation instead of the heavier $(N, \{R^{BR}_{\sigma}\}_{\sigma \in \bigtimes_{i \in N} \Sigma_i})$ and $(N, \{R^{DS}_{\sigma}\}_{\sigma \in \bigtimes_{i \in N} \Sigma_i})$.
[5]We recall that a relation R on domain N is reflexive if $\forall x \in N$, xRx. It is universal if R is the Cartesian square of the domain, i.e., $R = N^2$. Obviously, if R is universal, it is also reflexive.

	g	$\neg g$
g	3,3,3	2,4,2
$\neg g$	4,2,2	1,1,0

g

	g	$\neg g$
g	2,2,4	0,1,1
$\neg g$	1,0,1	1,1,1

$\neg g$

Figure 3: A three person Prisoner's dilemma.

3. R^{BR}_{σ} *is universal iff* $\forall i, j \in N$ σ_i *is a best response for* j;

4. R^{DS}_{σ} *is universal iff* $\forall i, j \in N$ σ_i *is a dominant strategy for* j.

Proof. [First claim:] From left to right, we assume that $\forall i \in N, iR^{BR}_{\sigma}i$. From Definition 6, it follows that $\forall i \in N, \forall \sigma'$: $o(\sigma) \succeq_i o(\sigma'_i, \sigma_{-i})$, that is, σ is a Nash equilibrium. From left to right, we assume that σ is a Nash equilibrium. From this it follows that $\forall i \in N$ σ_i is a best response for i, from which the reflexivity of R^{BR}_{σ} follows by Definition 6. [Second claim:] It can be proven in a similar way. [Third and fourth claims:] The proof is by direct application of Definition 6. □

Figure 1 provides neat examples of the claims listed in Fact 1. For instance, the first two claims are illustrated by profile (D, R) in the Prisoner's dilemma, which is both a Nash as well as a dominant strategy equilibrium. The profile (U, L) of the game Full Convergence is an instance of claim 3 whose outcome is most preferred by all players. Finally, profiles (U, L) and (D, R) in the Coordination game instantiate claim 4 by being the best profiles among all profiles that could be obtained by a unilateral deviation of any of the players. With respect to this, it is worth noticing that while (U, L) and (D, R) in the Coordination game correspond to a universal BR-dependence, the profile (D, R) in the Prisoner's dilemma only corresponds to a reflexive one.

3.3 Cycles

In the informal discussion of Example 1 we pointed at an essential difference between the pair of profiles (D, R) and (U, L) and the pair of profiles (D, L) and (U, R), namely the asymmetry in the dependence structure that the last two exhibit (see also Figure 2). For instance, in (D, L) Row BR-depends on Column while Column does not BR-depend on Row.

The game-theoretic instability of the profiles can be viewed, in dependence theoretic terms, as a lack of balance or reciprocity in the corresponding dependence structure. According to [4], a dependence is reciprocal when it allows for the occurrence of "social exchange", or exchange of favours, between the agents involved. This happens in the presence of cycles in the dependence relation [15]. In a cycle, the first player of the cycle could be prone to do what the last player asks since it can obtain something from the second player who, in turn, can obtain something from the third and so on.

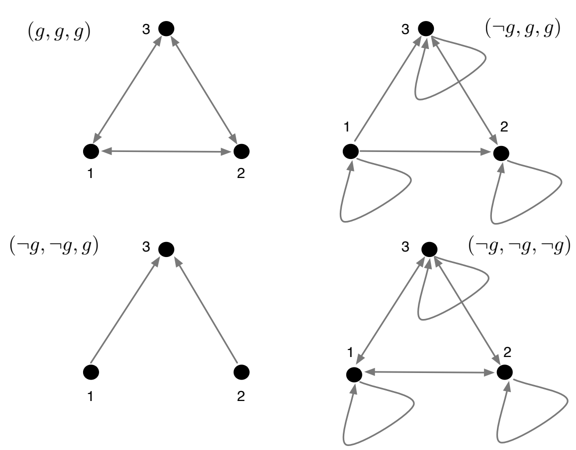

Figure 4: Some BR-dependences of Example 2.

Definition 7 (Dependence cycles). *Let* $\mathcal{G} = (N, S, \Sigma_i, \succeq_i, o)$ *be a game,* (N, R^x_{σ}) *be its dependence structure for profile* σ *with* $x \in \{BR, DS\}$*, and let* $i, j \in N$*. An* R^x_{σ}*-dependence cycle* $\mathfrak{c}$ *of length* $k-1$ *in* $\mathcal{G}$ *is a tuple* $(a_1, \ldots, a_k)$ *such that:* $a_1, \ldots, a_k \in N$; $a_1 = a_k$; $\forall a_i, a_j$ *with* $1 \leq i \neq j < k$, $a_i \neq a_j$; $a_1 R^x_{\sigma} a_2 R^x_{\sigma} \ldots R^x_{\sigma} a_{k-1} R^x_{\sigma} a_k$*. Given a cycle* $\mathfrak{c} = (a_1, \ldots, a_k)$*, its orbit* $O(\mathfrak{c}) = \{a_1, \ldots, a_{k-1}\}$ *denotes the set of its elements.*

In other words, cycles are sequences of pairwise different agents, except for the first and the last which are equal, such that all agents are linked by a dependence relation. Cycles become of particular interest in games with more than two players, so let us illustrate the definition by the following example.

Example 2 (Cycles in three person games.). *Consider the following three-person variant of the Prisoner's dilemma. A committee of three juries has to decide whether to declare a defendant in a trial guilty or not. All the three juries want the defendant to be found guilty, however, all three prefer that the others declare her guilty while she declares her innocent. Also, they do not want to be the only ones declaring her guilty if the other two declare her innocent. They all know each other's preferences. Figure 3 gives a payoff matrix for such game. Figure 4 depicts some cyclic BR-dependencies inherent in the game presented. Player 1 is Row, player 2 Column, and player 3 picks the right or left table. Among the ones depicted, the reciprocal profiles are clearly* (g, g, g), $(\neg g, \neg g, \neg g)$ *(which is also universal) and* $(\neg g, g, g)$*, only the last two of which are Nash equilibria (reflexive). Looking at the cycles present in these BR-reciprocal profiles, we notice that* (g, g, g) *contains the* 2×3 *cycles of length 3, all yielding the partition* $\{\{1, 2, 3\}\}$ *of the set of agents* $\{1, 2, 3\}$*. Profile* $(\neg g, g, g)$*, instead, yields two partitions:* $\{\{1\}, \{2\}, \{3\}\}$ *and* $\{\{1\}, \{2, 3\}\}$*. The latter is determined by the cycles* $(1, 1)$ *and* $(2, 3, 2)$ *or* $(1, 1)$ *and* $(3, 2, 3)$*. Finally, profile* $(\neg g, \neg g, g)$ *is such that both 1 and 2 depend on 3. Yet, neither of them plays a best response strategy.*

Reciprocity obtains a formal definition as follows.

Definition 8 (Reciprocal profiles). *Let* $\mathcal{G}$ *be a game and* (N, R^x_{σ}) *be its dependence structure with* $x \in \{BR, DS\}$ *and* σ *be a profile. A profile* σ *is* reciprocal *if and only if there exists a partition* $P(N)$ *of* N *such that each element* p *of the partition is the orbit of some* R^x_{σ}*-cycle.*

So, a profile is reciprocal when the corresponding dependence relation, be it a *BR*- or *DS*-dependence, clusters the agents into non-overlapping groups whose members are all part of some cycle of dependencies. Notice that the definition covers 'degenerate' cases such as the case of *trivially reciprocal* profiles where cycles are of the type (i, i), i.e., whose orbit is a singleton (cf. Fact 1). This is the case, for instance, in the (D, R) profile in the Prisoner's dilemma. Also, notice that several different cycles can coexist. This is for instance the case in universal dependence relations (cf. Fact 1). A profile yielding a cycle

	L	R
U	0,0	1,0
D	0,1	1,0

$\mathcal{G}$

	L	R
	0,0	0,1
	1,0	1,0

$\mathcal{G}^{\mu}$

Figure 5: The two horsemen game and its permutation.

with orbit N, e.g. (U, L) in the Prisoner's dilemma, is called *fully reciprocal*.

The literature on dependence theory stresses the existence of cycles as the essential characteristic for a social situation to give rise to some kind of cooperation. To say it with [1], cycles formalize the possibility of social interaction between agents of a *do-ut-des* (give-to-get) type. However, the fact that cycles are prerequisite for cooperation is, in dependence theory, normally taken for granted. The next section shows how the importance of dependence cycles can be well understood once game theory is taken into the picture.

3.4 Cycles as equilibria somewhere else

Consider a game $\mathcal{G} = (N, S, \Sigma_i, \succeq_i, o)$ and a bijection $\mu : N \mapsto N$. The μ-permutation of game $\mathcal{G}$ is the game $\mathcal{G}^{\mu} = (N, S, \Sigma^{\mu}, \succeq_i, o^{\mu})$ where for all $i \in N$, $\Sigma^{\mu}_i = \Sigma_{\mu(i)}$ and the outcome function $o_{\mu} : \bigtimes_{i \in N} \Sigma_{\mu(i)} \rightarrow S$ is such that $o_{\mu}(\mu(\sigma)) = o(\sigma)$, with $\mu(\sigma)$ denoting the permutation of σ according to μ. Intuitively, a permuted game $\mathcal{G}^{\mu}$ is therefore a game where the strategies of each player are redistributed according to μ in the sense that i's strategies become $\mu(i)$'s strategies, where agents keep the same preferences over outcomes, and where the outcome function assigns the same outcomes to the same profiles.

Example 3 (Two horsemen [12]). *"Two horsemen are on a forest path chatting about something. A passerby M, the mischief maker, comes along and having plenty of time and a desire for amusement, suggests that they race against each other to a tree a short distance away and he will give a prize of \$100. However, there is an interesting twist. He will give the \$100 to the owner of the slower horse. Let us call the two horsemen Bill and Joe. Joe's horse can go at 35 miles per hour, whereas Bill's horse can only go 30 miles per hour. Since Bill has the slower horse, he should get the \$100. The two horsemen start, but soon realize that there is a problem. Each one is trying to go slower than the other and it is obvious that the race is not going to finish. [. . .] Thus they end up [. . .] with both horses going at 0 miles per hour. [. . .] However, along comes another passerby, let us call her S , the problem solver, and the situation is explained to her. She turns out to have a clever solution. She advises the two men to switch horses. Now each man has an incentive to go fast, because by making his competitor's horse go faster, he is helping his own horse to win!" [12, p. 195-196].*

Once we depict the game of the example as the left-hand side game in Figure 5, we can view the second passerby's solution as a bijection μ which changes the game to the right-hand side version. Now Row can play Column's moves and Column can play Row's moves. The result is a swap of (D, L) with (U, R), since (D, L) in $\mathcal{G}^{\mu}$ corresponds to (U, R) in $\mathcal{G}$ and vice versa. On the other hand, (U, L) and (D, R) stay the same, as the exchange of strategies do not affect them. As a consequence, profile (D, R), in which both horsemen engage in the race, becomes a dominant strategy equilibrium.

On the ground of these intuitions, we can obtain a simple characterization of reciprocal profiles as equilibria in appropriately permuted games.

Theorem 1 (Reciprocity in equilibrium). *Let $\mathcal{G}$ be a game and (N, R^x_{σ}) be its dependence structure with $x \in \{BR, DS\}$ and σ be a profile. It holds that σ is x-reciprocal iff there exists a bijection $\mu : N \mapsto N$ s.t. σ is a x-equilibrium in $\mathcal{G}^{\mu}$.*

Proof. The theorem states two claims: one for $x = BR$ and one for $x = DS$. [First claim.] From left to right, assume that σ is BR-reciprocal and prove the claim by constructing the desired μ. By Definition 8 it follows that there exists a partition P of N such that each element p of the partition is the orbit of some R^{BR}_{σ}-cycle. Given P, observe that any agent i belongs to at most one member p of P. Now build μ so that $\mu(i)$ outputs the successor j (which is unique) of i in the cycle whose orbit is the p to which i belongs. Since each j has at most one predecessor in a cycle, μ is an injection and since domain and codomain coincide μ is also a surjection. Now it follows that for all i, j, $iR^{BR}_{\sigma}j$ implies, by Definition 6, that $\sigma_{\mu(i)}$ is a best response for i in σ. But in $\mathcal{G}^{\mu}$ it holds that $\sigma_{\mu(i)} \in \Sigma_i$ and since σ is reciprocal, by Definition 8, we have that for all i σ_i is a best response in $\mathcal{G}^{\mu}$, and hence it is a Nash equilibrium. From right to left, assume μ to be the bijection at issue. It suffices to build the desired partition P from μ by an inverse construction of the one used in the left to right part of the claim. We set $iR^{BR}_{\sigma}j$ iff $\mu(i) = j$. The definition is sound w.r.t. Definition 6 because σ being a Nash equilibrium we have that $iR^{BR}_{\sigma}j$ iff j plays a best response for i in σ. Since μ is a bijection, it follows that R^{BR}_{σ} contains cycles whose orbits are disjoint and cover N. Therefore, by Definition 8, we can conclude that σ is BR-reciprocal. [Second claim.] The proof is similar to the proof of the first claim. □

A profile is reciprocal if and only if it is an equilibrium—under a given solution concept—in the game yielded by a reallocation of the players' strategy spaces, where the players keep the same preferences of the original game. In a nutshell, a reciprocal profile is nothing but an equilibrium after strategy permutation. The following corollary is a direct consequence of Theorem 1.

Corollary 1 (Trivially and fully reciprocal profiles). *Let $\mathcal{G}$ be a game and (N, R^x_{σ}) be its dependence structure with $x \in \{BR, DS\}$ and σ be a profile. It holds that:*

1. *σ is trivially x-reciprocal iff σ is an x-equilibrium in $\mathcal{G}^{\mu}$ where μ is the identity;*

2. *σ is fully x-reciprocal iff σ is an x-equilibrium in $\mathcal{G}^{\mu}$ where μ is such that $\mu^{|N|}$ is the identity,[6] and for no $n < |N|$ μ^n is the identity.*

The corollary makes explicit that: trivially reciprocal profiles are the equilibria arising from the maintenance of the 'status quo', so to say; and that fully reciprocal profiles are the equilibria arising after the widest reallocation of strategies possible.

From the foregoing results, it follows that permutations can be fruitfully viewed as ways of *implementing* a reciprocal profile, where implementation has to be understood as a way of transforming a game in such a way that the desirable outcomes, in the transformed game, are brought about at an equilibrium point.[7]

[6]Function $\mu^{|N|}$ is the $|N|$th iteration of μ.

[7]This is another way of looking at *constrained mechanism design* as described in [14, Ch. 10.7] or, indeed, social software [12].

	L	R
U	2,2	0,3
D	3,0	1,1

$\mathcal{G}^{\mu}$

L	R
2,2	3,0
0,3	1,1

$\mathcal{G}^{\nu}$

Figure 6: Agreements between the prisoners

4. DEPENDENCY SOLVED: AGREEMENTS

We have seen in Section 1 that one of the original, and yet unmet, objectives of dependence theory was to "devise a calculus to obtain predictions" [4] of agents' behaviour. The present section shows how, by looking at dependence theory the way we suggest, dependence structures lend themselves to analytical predictions of the behaviour of the system to which they pertain.

4.1 Agreements

Let us start with the definition of an agreement.

Definition 9 (Agreements). *Let $\mathcal{G}$ be a game, (N, R^x_σ) be its dependence structure in σ with $x \in \{BR, DS\}$, and let $i, j \in N$. A pair (σ, μ) is an x-agreement for $\mathcal{G}$ if σ is an x-reciprocal profile, and μ a bijection which x-implements σ in $\mathcal{G}$. The set of x-agreements of a game $\mathcal{G}$ is denoted x-$AGR(\mathcal{G})$.*

Intuitively, an agreement, of BR or DS type, can be seen as the result of agents' coordination selecting a desirable outcome and realizing it by an appropriate exchange of strategies. In fact, a bijection μ formalizes a precise idea of social exchange (see Section 2) in a game-theoretic setting. Notice that, from Definition 8 it follows that the bijection μ of an agreement (σ, μ) for a game $\mathcal{G}$ partitions N according to the cyclical dependence structure (N, R^x_σ). Such partition, which we call the partition of N induced by μ, is denoted $P_\mu(N)$. Definition 9 is applied in the following example.

Example 4 (Permuting Prisoners). *In the game Prisoner's Dilemma (Example 1) we observe two DS-agreements, whose permutations give rise to the games depicted in Figure 6. Agreement $((D, R), \mu)$ with $\mu(i) = i$ for all players, is the standard DS-equilibrium of the strategic game. But there is another possible agreement, where the players swap their strategies: it is $((U, L), \nu)$, for which $\nu(i) = N \backslash \{i\}$. Here Row plays cooperatively for Column and Column plays cooperatively for Row. Of the same kind is the agreement arising in Example 3. Notice that in such example, the agreement is the result of coordination mediated by a third party (the second passerby). Analogous considerations can also be done about Example 2 where, for instance, $((g, g, g), \mu)$ with $\mu(1) = 2, \mu(2) = 3, \mu(3) = 1$ is a BR- agreement.*

Definition 9 has defined agreements with respect to both notions of best response (*BR*-agreement) and dominant strategy (*DS*-agreement). In the remaining of the paper we will focus, for the sake of exposition, only on *DS*-agreements. So, by 'agreement' we will mean *DS*-agreement.

4.2 Ordering agreements

As there can be several possible agreements in a game, the natural issue arises of how to order them and, in particular, of how to order them so that the maximal agreements, with respect to such ordering, can be considered to be the agreements that will actually be realized in the game.

A first candidate for such ordering could be Pareto dominance. Such ordering on the agreements (σ, μ) would simply be the same as the one defined on the profiles σ of the agreements.[8] However, we are interested in an ordering which takes into consideration the fact that agreements can be strategically taken by coalitions. To define the desired ordering that takes strategic behaviour into account, we need the following notion of C-candidate and C-variant agreements.

Definition 10 (C-candidates and C-variants). *Let $\mathcal{G} = (N, S, \Sigma_i, \succeq_i, o)$ be a game and C a non-empty subset of N. An agreement (σ, μ) for $\mathcal{G}$ is a* C-candidate *if C is the union of some members of the partition induced by μ, that is: $C = \bigcup X$ where $X \subseteq P_\mu(N)$. An agreement (σ, μ) for $\mathcal{G}$ is a* C-variant *of an agreement (σ', μ') if $\sigma_C = \sigma'_C$ and $\mu_C = \mu'_C$, where μ_C and μ'_C are the restrictions of μ to C. As a convention we take the set of $\emptyset$-candidate agreements to be empty and an agreement (σ, ν) to be the only $\emptyset$-variant of itself.*

In other words, an agreement (σ, μ) is a C-candidate if $\{C, \overline{C}\}$ is a bipartition of $P_\mu(N)$, and it is a C-variant of (σ', μ') if it differs from this latter at most in its C-part. It is instructive to notice the following: if an agreement is a C-candidate, it is also a $\overline{C}$-candidate; all agreements are N-candidate; and if an agreement is a C-variant of a C-candidate, it is also a C-candidate. Finally, notice also that if (σ, μ) and (σ', μ') are two C-candidate agreements, then $((\sigma_C, \sigma'_{\overline{C}}), (\mu_C, \mu'_{\overline{C}}))$ is also an agreement. We can now define the following notion of dominance between agreements.

Definition 11 (Dominance). *Let $\mathcal{G} = (N, S, \Sigma_i, \succeq_i, o)$ be a game. An agreement (σ, μ) is dominated if for some coalition C there exists a C-candidate agreement (σ', μ') for $\mathcal{G}$ such that for all agreements (ρ, ν) which are $\overline{C}$-variants of (σ', μ'), $o(\rho, \nu) \succ_i o(\sigma)$ for all $i \in C$. The set of undominated agreements of $\mathcal{G}$ is denoted $DEP(\mathcal{G})$.*

Intuitively, an agreement is dominated when a coalition C can force all possible agreements to yield outcomes which are better for all the members of the coalition, regardless of what the rest of the players do, that is, regardless of the $\overline{C}$-variants of their agreements. To put it yet otherwise, the members of coalition C can exploit the dependencies among them creating a partial agreement (σ_C, μ_C) which suffices to force the outcomes of the game to be all better than the ones obtained via any other agreement.

Example 5 (Agreements with three players). *In the three person Prisoner's dilemma drawn in Figure 3 and analyzed in Example 2, the strategy profile $(\neg g, \neg g, \neg g)$ is a DS-equilibrium and an agreement under the identity permutation μ while $(\neg g, g, g)$ is an agreement under any permutation ν that induces a partition $\{\{1\}, \{2, 3\}\}$ on N. As can be checked, $((\neg g, \neg g, \neg g), \mu)$ is dominated, since there exist a coalition $C = \{2, 3\}$ and a C-candidate agreement, namely $((\neg g, g, g), \nu)$, that has itself as only $\{1\}$-variant, with $o(\neg g, g, g) \succ_i o(\neg g, \neg g, \neg g)$ for all $i \in C$.*

4.3 Dependencies and coalitions

The previous sections have shown how to identify dependences within games in strategic form. If such games are

[8] For completeness, here is the definition: Let $\mathcal{G} = (N, S, \Sigma_i, \succeq_i, o)$ be a game. An agreement (σ, μ) for $\mathcal{G}$ is Pareto optimal if for no agreement (σ', μ'), $o(\sigma')$ Pareto dominates $o(\sigma)$, i.e., $\neg\exists\sigma'$ s.t. $\forall i \in N : o(\sigma') \succ_i o(\sigma)$.

studied from the point of view of cooperative game theory—along the lines of the quote from [16] given in Section 1—what kind of cooperation is the one that arises based on reciprocal dependences and agreements? In other words, what is the place of dependencies within cooperative game theory?

We proceed as follows. First, starting from a game $\mathcal{G}$, we consider its representation $C^{\mathcal{G}}$ as a coalitional game (Definition 12). Then we refine such representation, obtaining a coalitional game $C^{\mathcal{G}}_{DEP}$ which captures the intuition that coalitions form only by means of agreements as they have been defined in Definition 9. At this point we show that the core of $C^{\mathcal{G}}_{DEP}$ coincides with the set of undominated agreements of $\mathcal{G}$ (Theorem 2), thereby obtaining a game-theoretical characterization of Definition 11.

So let us start with the definition of a cooperative game obtained from a strategic form one (cf. [11]).

Definition 12 (Coalitional games from strategic ones). *Let $\mathcal{G} = (N, S, \Sigma_i, \succeq_i, o)$ be a game. The coalitional game $C^{\mathcal{G}} = (N, S, E^{\mathcal{G}}, \succeq_i)$ of $\mathcal{G}$ is a coalitional game where the effectivity function $E^{\mathcal{G}}$ is defined as follows:*

$$X \in E^{\mathcal{G}}(C) \Leftrightarrow \exists \sigma_C \forall \sigma_{\overline{C}}\ o(\sigma_C, \sigma_{\overline{C}}) \in X.$$

Roughly, the effectivity function of $C^{\mathcal{G}}$ contains those sets in which a coalition C can force the game to end up, no matter what $\overline{C}$ does.

Definition 12 abstracts from dependence-theoretic considerations. So, in order to be able to study what outcomes are stable once we give to the players the power to form coalitions via agreements, we have to define the effectivity function so that the states for which a coalition is effective depend on the agreements it can force.

Definition 13 (Dependence games from strategic ones). *Let $\mathcal{G} = (N, S, \Sigma_i, \succeq_i, o)$ be a game. The dependence (coalitional) game $C^{\mathcal{G}}_{DEP} = (N, S, E^{\mathcal{G}}_{DEP}, \succeq_i)$ of $\mathcal{G}$ is a coalitional game where the effectivity function $E^{\mathcal{G}}_{DEP}$ is defined as follows:*

$$\begin{aligned} X \in E^{\mathcal{G}}_{DEP}(C) \quad \Leftrightarrow \quad & \exists \sigma_C, \mu_C \text{ s.t.} \\ & \exists \sigma_{\overline{C}}, \mu_{\overline{C}} : [((\sigma_C, \sigma_{\overline{C}}), (\mu_C, \mu_{\overline{C}})) \in AGR(\mathcal{G})] \\ & \text{and } [\forall \sigma_{\overline{C}}, \mu_{\overline{C}} : [((\sigma_C, \sigma_{\overline{C}}), (\mu_C, \mu_{\overline{C}})) \in AGR(\mathcal{G}) \\ & \text{implies } o(\sigma_C, \sigma_{\overline{C}}) \in X]]. \end{aligned}$$

where $\mu : N \rightarrow N$ is a bijection.

This somewhat intricate formulation states nothing but that the effectivity function $E^{\mathcal{G}}_{DEP}(C)$ associates with each coalition C the states which are outcomes of agreements (and hence of reciprocal profiles), and which C can force via partial agreements (σ_C, μ_C) regardless of the partial agreements $(\sigma_{\overline{C}}, \mu_{\overline{C}})$ of $\overline{C}$. Whether agreements exist at all depends, obviously, on the underlying game $\mathcal{G}$. The following fact compares Definitions 12 and 13.

Fact 2 ($E^{\mathcal{G}}$ vs. $E^{\mathcal{G}}_{DEP}$). *It does not hold that for all $\mathcal{G}$: $E^{\mathcal{G}}_{DEP} \subseteq E^{\mathcal{G}}$; nor it holds that for all $\mathcal{G}$: $E^{\mathcal{G}} \subseteq E^{\mathcal{G}}_{DEP}$.*

Proof. For the first inclusion consider a game $\mathcal{G}$ with tree players $1, 2, 3$ and two actions $\{a, b\}$ for each of them. Suppose the only possible agreement is the identity permutation $\mu(i) = i$ and (a, a, a) is a DS-equilibrium. We have that $\{o(a, a, a)\} \in E^{\mathcal{G}}_{DEP}(\{1\})$ while $\{o(a, a, a)\} \notin E^{\mathcal{G}}(\{1\})$. For the second inclusion take $\mathcal{G}$ the Prisoner's dilemma game in which $\{(U, L)\} \in E^{\mathcal{G}}(\{Column, Row\})$ but $\{(U, L)\} \notin E^{\mathcal{G}}_{DEP}(\{Column, Row\})$. □

The fact shows that dependence games are not just a refinement of coalitional ones. Dependence-based effectivity functions considerably change the powers of coalitions.

At this point two natural questions arise. First, given a game $\mathcal{G}$, what is the relation between the set $DEP(\mathcal{G})$ (Definition 11) and the set $CORE(C^{\mathcal{G}}_{DEP})$, i.e., the core of the dependence game built from $\mathcal{G}$? Second, along the lines of Fact 2, what is the relation between the set $CORE(C^{\mathcal{G}})$ and the set $CORE(C^{\mathcal{G}}_{DEP})$, that is, what is the relation between the cores of the coalitional and dependence games built on $\mathcal{G}$? These questions are answered by the two results below, but let us first point at some differences between coalitional games and dependence games by means of an example.

Example 6 (core in coalitional vs. dependence games). *The core of the coalitional game and of the dependence game built on the Prisoner's dilemma game (Figure 1) coincide and are $\{o(U, L)\}$, i.e., the cooperative outcome. Take $\mathcal{G}$ to be the Prisoner's dilemma game. That $CORE(C^{\mathcal{G}}) = \{o(U, L)\}$ is clear, as there exists no coalition C that can force a set of outcomes which are all better for all members of C. That $CORE(C^{\mathcal{G}}_{DEP}) = \{o(U, L)\}$ is a little subtler. As shown in Example 4, there are two possible agreements in the Prisoner's dilemma: the cooperative one leading to the fully DS-reciprocal profile, and the identity one, leading to the trivially DS-reciprocal profile. These profiles are the only ones that can be forced by the coalitions, and clearly $o(U, L)$ dominates $o(D, R)$.*

Theorem 2 (DEP vs. $CORE$). *Let $\mathcal{G} = (N, S, \Sigma_i, \succeq_i, o)$ be a game. It holds that, for all agreements (σ, μ):*

$$(\sigma, \mu) \in DEP(\mathcal{G}) \ \Leftrightarrow\ o(\sigma) \in CORE(C^{\mathcal{G}}_{DEP}).$$

where $\mu : N \rightarrow N$ is a bijection.

Proof. [Left to right:] By contraposition, assume $o(\sigma) \notin CORE(C^{\mathcal{G}}_{DEP})$. By Definition 4 this means that $\exists C \subseteq N, X \in E^{\mathcal{G}}_{DEP}(C)$ s.t. $x \succ_i o(\sigma)$ for all $i \in C, x \in X$. Applying Definition 13 we obtain that there exists an agreement $((\sigma'_C, \sigma'_{\overline{C}}), (\mu'_C, \mu'_{\overline{C}}))$ s.t. $\forall \sigma'_{\overline{C}}, \mu'_{\overline{C}}, o(\sigma'_C, \sigma'_{\overline{C}}) \in X$ and s.t. $x \succ_i o(\sigma)$ for all $i \in C, x \in X$. Now, $((\sigma'_C, \sigma'_{\overline{C}}), (\mu'_C, \mu'_{\overline{C}}))$ is obviously C-candidate, and all its C-variants yield better outcomes for C than σ. Hence, by Definition 11,$(\sigma, \mu) \notin DEP(\mathcal{G})$. [Right to left:] Notice that the set up of Definition 11 implies that, if (σ, μ) is dominated, then any other agreement for σ would also be dominated. So, by contraposition, assume $(\sigma, \mu) \notin DEP(\mathcal{G})$. By Definition 11, we obtain that there exists a C-candidate agreement (σ', μ') for $\mathcal{G}$ such that for all agreements (ρ, ν) which are $\overline{C}$-variants of (σ', μ'), $o(\rho, \nu) \succ_i o(\sigma)$ for all $i \in C$. But this means, by Definition 13, that $\exists C, X$ such that $X \in E^{\mathcal{G}}_{DEP}(C)$ and $x \succ_i o(\sigma)$ for all $x \in C$. Hence, by Definition 4, we obtain $\sigma \notin CORE(C^{\mathcal{G}}_{DEP})$. □

Intuitively, an agreement for a game $\mathcal{G}$ is undominated if and only if the outcome of its profile lies in the core of the dependence game built on $\mathcal{G}$.[9] To put it yet otherwise, Theorem 2 states that, given a game $\mathcal{G}$ and agreement (σ, μ), the permuted game $\mathcal{G}^{\mu}$ where σ is a DS-equilibrium (cf. Theorem 1) lying in $DEP(\mathcal{G})$, if and only if σ is in the core of the dependence game of $\mathcal{G}$.

[9]Notice also that, as a consequence of Theorem 2, the existence of undominated agreements guarantees the nonemptyness of the core. Notice that the converse need not hold as there could be undominated outcomes which are not the result of agreements.

Fact 3 ($CORE(C^{\mathcal{G}})$ vs. $CORE(C^{\mathcal{G}}_{DEP})$). *It does not hold that for all $\mathcal{G}$: $CORE(C^{\mathcal{G}}) \subseteq CORE(C^{\mathcal{G}}_{DEP})$; nor it holds that for all $\mathcal{G}$: $CORE(C^{\mathcal{G}}_{DEP}) \subseteq CORE(C^{\mathcal{G}})$.*

Proof. For the first inclusion, take $\mathcal{G}$ to be the Partial Convergence Game in Figure 1. While the outcome $o(D, L)$ belongs to $CORE(C^{\mathcal{G}})$, it does not belong to $CORE(C^{\mathcal{G}}_{DEP})$. The strategy profile (U, L) is an agreement under the identity permutation and it is the only Column-variant of itself. Even though $o(D, L) \succ_{Column} o(U, L)$, playing U is dominant strategy for Row and playing L is dominant strategy for Column. The only possible agreement under the identity permutation is not an optimal outcome for Column. As can be observed, Row can reach an agreement with the identity permutation, successfully deviating from (D, L). For the converse inclusion, take $\mathcal{G}$ to be the Coordination game (Figure 1). Then $CORE(C^{\mathcal{G}}) = \{o(U, L), o(D, R)\}$ while $CORE(C^{\mathcal{G}}_{DEP})$ contains all the four possible outcomes, simply because there is no *DS*-agreement in such game, hence no possible coalition. □

With this result at hand we can observe that dependence structures carry out a different—and alternative—notion of stability from the one usually studied in cooperative games.

4.4 Discussion

The results presented in the previous section hinge on the notion of agreement intended as *DS*-agreement. Variants of the above results can be obtained also for *BR*-agreements.

It is moreover worth stressing that Definition 11, and its corresponding game-theoretic notion of core in dependence games, is just one possible candidate for the formalization of a notion of 'stability' of agreements. Other notions of dominance among agreements can be isolated and related to solution concepts in coalitional games just like we did in Theorem 2. To give an example, an arguably natural candidate for a notion of 'stability' for agreements is the following one, which we call *Nash agreement*. An agreement (σ, μ) for $\mathcal{G}$ is Nash if and only if there exists no $C \subseteq N$, and agreement (σ', μ') such that (σ', μ') is a C-variant of (σ, μ) and, for all $i \in C$, $o(\sigma') \succ_i o(\sigma)$. The question is then: can we build a (coalitional) game, based on $\mathcal{G}$, and single out an appropriate solution concept such that the solution of that game corresponds to the set of Nash agreements of $\mathcal{G}$?

Finally, alternatives to the definition of agreements are also possible, and perhaps desirable. Agreements have been formalized here using permutations onto the set N. An alternative path to take is the use of partial agreements as only admissible coalitional strategies, allowing for permutations onto subsets of N. This should guarantee the core of the coalitional game to be included in the core of the dependence game so constructed (cf. Fact 3).

These particular research issues are left for future work. What we want to stress, however, is that a series of results of this kind could set the boundaries of dependence theory within the theory of games, thereby giving it a well-defined, and mathematically solid, 'own' place.

5. CONCLUSIONS

The contribution of the paper is two-fold. On the one hand it has been shown that central dependence-theoretic notions such as the notion of cycle are amenable to a game-theoretic characterization (Theorem 1). On the other hand dependence theory has been demonstrated to give rise to types of cooperative games where solution concepts such as the core can be applied to obtain the sort of 'analytical predictive power' that dependence theory unsuccessfully looked for since its beginnings [4] (Theorem 2).

Acknowledgments.

The authors would like to thank the anonymous reviewers of AAMAS 2010 for their helpful remarks. Davide Grossi is supported by the *Nederlandse Organisatie voor Wetenschappelijk Onderzoek* (NWO) under the VENI grant 639.021.816.

6. REFERENCES

[1] G. Boella, L. Sauro, and L. van der Torre. Strengthening admissible coalitions. In *Proceeding of the 2006 conference on ECAI 2006: 17th European Conference on Artificial Intelligence*, pages 195–199. ACM, 2006.

[2] E. Bonzon, M. Lagasquie-Schiex, and J. Lang. Dependencies between players in boolean games. *International Journal of Approximate Reasoning*, 50:899–914, 2009.

[3] C. Castelfranchi. Modelling social action for AI agents. *Artificial Intelligence*, 103:157–182, 1998.

[4] C. Castelfranchi, A. Cesta, and M. Miceli. Dependence relations among autonomous agents. In E. Werner and Y. Demazeau, editors, *Decentralized A.I.3*. Elsevier, 1992.

[5] J. Coleman. *Foundations of Social Theory*. Belknap Harvard, 1990.

[6] P. Harrenstein, W. van der Hoek, J. Meyer, and C. Witteveen. Boolean games. In J. van Benthem, editor, *Proceedings of TARK'01*, pages 287–298. Morgan Kaufmann, 2001.

[7] J.Sichman. Depint: Dependence-based coalition formation in an open multi-agent scenario. *Journal of Artificial Societies and Social Simulation*, 1(2), 1998.

[8] H. Moulin and B. Peleg. Cores of effectivity functions and implementation theory. *Journal of Mathematical Economics*, 10:115–145, 1982.

[9] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[10] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[11] M. Pauly. *Logic for Social Software*. ILLC Dissertation Series, 2001.

[12] R.Parikh. Social software. *Synthese*, 132(3):187–211, 2002.

[13] L. Sauro, L. van der Torre, and S. Villata. Dependency in cooperative boolean games. In A. Håkansson, N. Nguyen, R. Hartung, R. Howlett, and L. Jain, editors, *Proceedings of KES-AMSTA 2009*, volume 5559 of *LNAI*, pages 1–10. Springer, 2009.

[14] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic and Logical Foundations*. Cambridge University Press, 2008.

[15] J. Sichman and R. Conte. Multi-agent dependence by dependence graphs. In *Proceedings of AAMAS 2002*, ACM, pages 483–490, 2002.

[16] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.