

Logical Consequence and the Theory of Games

*Paul Harrenstein**
Utrecht University

Résumé : Les notions logiques de conséquence sont fréquemment reliées à des concepts de solution de la théorie des jeux. Dans ce contexte domine la correspondance entre une formule classiquement valide et l'existence d'une stratégie gagnante pour un joueur dans un jeu à deux joueurs. Nous proposons une extension conservatrice de la notion classique de conséquence basée sur une généralisation du concept de solution de jeu d'équilibre de Nash.

Abstract: Logical notions of consequence have frequently been related to game-theoretical solution concepts. The correspondence between a formula being classically valid and the existence of a winning strategy for a player in a related two-person game, has been most prominent in this context. We propose a conservative extension of the classical notion of consequence that is based on a generalization of the game-theoretical solution concept of Nash equilibrium.

*This research is partly but generously supported by the CABS (Collective Agent Based Systems) project of Delft University of Technology.

Introduction

In their *Theory of Games and Economic Behavior* von Neumann and Morgenstern argued that situations of conflicting interests present a problem that had been “nowhere dealt with in classical mathematics” [Neumann & Morgenstern 1944, p. 11]. They maintained that, due to its interactive nature, a conflict situation could not be analyzed as a traditional optimization or decision problem. Rather, it is a “peculiar and disconcerting mixture of several maximum problems” (*ibid.*, page 11). An optimization or decision problem for an individual can be represented formally as a function $f(\hat{x}_0, \dots, \hat{x}_n)$. The individual’s predicament is then to choose values for the variables x_0, \dots, x_n so as to maximize the value of $f(\hat{x}_0, \dots, \hat{x}_n)$. The variables on which the function depends are regarded as decision variables that are in the control of the individual. Pursuing this conceptualization, a situation of conflict could in similar terms be understood as a *collection* of functions $g_i(\hat{x}_0, \dots, \hat{x}_n)$, each one of which one of the participants tries to maximize by choosing suitable values for the variables in a way that furthers his idiosyncratic interests. Moreover, the variables on which these functions depend may overlap and the parties involved may have control over only some of the relevant variables. This makes that the optimal choices for an individual’s variables, from his perspective, may be dependent on the very choices the other participants make in their effort to maximize their functions from their respective points of view. Thus the issue may evoke a sense of impending circularity.

Traditional notions of optimality were thought to be no longer adequate for such problems and new mathematical notions – viz., game-theoretical solution concepts – had to be developed to take over their role (*ibid.*, page 39). In non-cooperative settings Nash equilibrium is archetypical in this respect. Informally, an assignment of values for the joint set of variables (henceforth a *strategy profile*) is a Nash equilibrium, if none of the individuals can improve on it by unilaterally choosing different values for the variables in his control.

Having distinguished optimization problems and game-theoretical problems thus, the satisfiability problem for Classical Propositional Logic (CPC) could be classified as an optimization problem with respect to truth. A formula is thought of as a function in its propositional variables. The issue is then to choose values for the propositional variables so as to satisfy the formula in question. Classical logical consequence can be understood in similar deliberative terms: a formula φ follows from a collection of premisses Γ if and only if, each choice for the truth values

of the propositional variables (henceforth a *valuation*) that succeeds in satisfying all formulas in Γ , is a choice that makes φ hold as well.

As in this formulation there is present a definite element of choice with respect to the possible truth-assignments, we come to think of propositional variables as binary decision variables that are in the control of a decision maker. The accompanying image of a logical possibility is that of a situation which obtains as the result of the decisions of an individual, rather than that of an unalterable state of affairs.¹

In line with this, it also becomes natural to consider the case in which control over the propositional variables is distributed over multiple agents. Logical space then assumes a game-theoretical aspect, with the valuations as strategy profiles.

In analogy with the relation between optimization and game-theoretical problems, these considerations give rise to the following issue, which can be regarded as the game-theoretical counterpart of the classical problem of logical consequence. *Which conclusions is one to draw from a family of theories, given that, for each of these theories, there is a player who controls a (disjoint) set of propositional variables and who seeks to satisfy his theory as well as he can by choosing appropriate values for the variables in his control?* This is a logical question, at the basis of which there is a game-theoretical problem. For its resolution we take recourse to the game-theoretical notion of a *maximum equilibrium*, which we will introduce as a generalization of Nash equilibrium.

We argue that any particular distribution of the propositional variables and any particular family of theories define a unique strategic game, which we refer to as a *distributed evaluation game*. We propose to consider as the consequences of a family of theories and a distribution of the propositional variables, those formulas that are satisfied in the maximum equilibria of the accompanying game. This defines a game-theoretical concept of consequence.

Example 1. Consider a propositional language with only two variables, a and b . Suppose that one player, *Row*, has control over a and wishes to satisfy the formula $a \wedge \neg b$. Let there further be another player, *Column*, who has control over b and aims at the satisfaction of the formula $\neg(a \vee b)$. The situation is summarized in the matrix below, in

¹This way of viewing propositional variables as controllable by individuals has its precursors in the field of Artificial Intelligence. A good example is Boutilier's distinction between controllable and uncontrollable propositions in [Boutilier 1994]. The image also underlies recent studies in distributed constraint satisfaction problems such as [Yokoo et al. 1998, Walshetal 2001].

which *Row* chooses rows – i.e., setting a to false (\emptyset) or to true ($\{a\}$) – and *Column* choosing columns in a similar fashion:²

	\emptyset	$\{b\}$
\emptyset	1	0
$\{a\}$	0	0
	1	0

The figures in boldface indicate the Nash equilibria, which in this particular game coincide with the maximum equilibria. Since a satisfied by both equilibria, viz., $\{a\}$ and $\{a, b\}$, it is considered a game-theoretical consequence of the theories $\{a \wedge \neg b\}$ and $\{\neg(a \vee b)\}$ given the distribution of a to *Row* and b to *Column*. However, b does not follow thus, as it not satisfied by the valuation $\{a\}$, although the latter is an equilibrium.

A player in a distributed evaluation game is thought of as preferring valuations that satisfy his theory to those that do not. We will argue, however, that theories can be interpreted as defining more gradated preferences over the valuations and that, by doing so, more justice is done to the interactive nature of the issue at hand. Thus a more comprehensive class of games is brought within the scope of propositional logic.

The two main ideas on which this paper pivots – distributing control over the propositional variables and interpreting theories as preference relations — may seem to indicate a rash departure from the traditional canons of propositional logic. Yet, we find that the concept of game-theoretical consequence can be regarded as a conservative extension of the classical notion in the following sense. In the special case in which the control over the propositional variables is concentrated in one player, a formula φ is a game-theoretical consequence of a (singleton!) family of theories $\{\Gamma_A\}$ whenever φ is a classical consequence of Γ .

1 Games and Maximum Equilibria

In this section we review some of the elementary concepts from game theory. The notion of a maximum equilibrium is also introduced.

²Each cell of the matrix represents a strategy profile. The bottom-left entries indicate the *ordinal* preferences of the player choosing rows, the upper-right entries those of the player choosing columns.

We define a *strategic game* as a tuple $(N, \{S_i\}_{i \in N}, \{\rho_i\}_{i \in N})$, where N is a countable and possibly infinite set of players and, for each player i in N , S_i is the set of strategies available to i . Accordingly, the generalized Cartesian product over S , i.e., $\times_{i \in N} S_i$, is the set of *strategy profiles* of the game, which we also denote by S . For each $i \in N$, ρ_i is a reflexive and transitive, but not necessarily connected relation on the strategy profiles S ; we also allow ρ_i to be the empty relation. Moreover, \leq_i is used as the infix notion of ρ_i . Thus, S could be considered to be an $\|N\|$ -dimensional space with for each strategy profile s and each player i in N , the strategy s_i its i -th coordinate. We will adopt the notation (s_{-i}, s'_i) for the point that is like s except for the i -th coordinate, which is identical with the i -th coordinate of s' . Intuitively, each (s_{-i}, s'_i) denotes a strategy profile that player i can reach from s by unilaterally deviating and playing s'_i .

By a *partial preorder* we understand a reflexive and transitive relation. If a partial preorder is moreover connected, we refer to it as a *total preorder*. The notion of a *Nash equilibrium* in pure strategies is usually defined on games in which the preferences of the players are total preorders over the strategy profiles. Then, for s a strategy profile in a game $(N, \{S_i\}_{i \in N}, \{\rho_i\}_{i \in N})$:

s is a *Nash equilibrium* iff for all $i \in N$, for all $s' \in S$: $(s_{-i}, s'_i) \leq_i s$.

We say strategy profile s is a *best response for a player i* if $(s_{-i}, s'_i) \leq_i s$, for all strategy profiles s' in S . Obviously, the set of strategy profiles that contain a best response for each player coincides with the set of Nash equilibria.

Our investigations, however, concern games in which the players' preference relations are also allowed to be *partial* preorders and even to be empty. We are now confronted with at least two obvious conservative extensions of the notion of a Nash equilibrium. On total preorders the notions of a maximal element (no other element is greater) and a maximum element (greater than any other element) coincide, but on partial preorders or the empty relation they may diverge. Accordingly, one could define, for i a player and s a strategy profile in a game G :

s is a *maximal response for i* iff for all $s' \in S$: $s \not\prec_i (s_{-i}, s'_i)$,

s is a *maximum response for i* iff for all $s' \in S$: $(s_{-i}, s'_i) \leq_i s$.

Lacking connectivity, the set of maximal responses for a player i may contain elements s, s' that are incomparable for i (on the i -th coordinate) but which are such that $s_j = s'_j$, for each $j \neq i$. This possibility

is excluded for maximum response strategies. Accordingly, we define a strategy profile s of a game G to be *maximal equilibrium* in G if for all players s is a maximal response. Similarly, a strategy profile s is a *maximum equilibrium* in G if s is a maximum response for all players. Observe that a strategy profile being a maximum equilibrium implies it being a maximal equilibrium, but not in general the other way round. By refining the preference orders of the players — i.e., if the preference relations become smaller — the number of maximal equilibria may increase; this is impossible with maximum equilibria. Hence, the following monotonicity property holds for maximum equilibria only.

Proposition 1. *Let $(N, \{S_i\}_{i \in N}, \{\rho_i\}_{i \in N})$ and $(N, \{S_i\}_{i \in N}, \{\rho'_i\}_{i \in N})$ be strategic games denoted by G and G' , respectively. Let, further, for each player $i \in N$, $\rho'_i \subseteq \rho_i$. Then the maximum equilibria of G' are included in the maximum equilibria of G .*

Proof. For the contrapositive, consider an arbitrary $s \in S$ that is *not* a maximum equilibrium in G . Then, for some player i and for some $s' \in S$, $((s_{-i}, s'_i), s) \notin \rho_i$. Since $\rho'_i \subseteq \rho_i$, $((s_{-i}, s'_i), s) \notin \rho'_i$. Hence, s is not a maximum equilibrium in G' . \square

2 Logical Consequence

Game-theoretical consequence will be defined for propositional languages and it will rely on an interpretation of formulas and theories as relations over the valuations. Some remarks with respect to classical propositional logic are in order.

A propositional logic is defined as a pair $(L(A), \vdash)$, where $L(A)$ is a propositional language over a set A of propositional variables and \vdash a binary relation between theories and formulas of $L(A)$. We will assume the formulas of a propositional language to be given by a minimal set that contains a non-empty but countable number of propositional variables A and that is closed under negation \neg and conjunction \wedge . The connectives falsum \perp , verum \top , disjunction \vee and implication \rightarrow are defined as usual.

For classical propositional logic $\Gamma \vdash \varphi$ informally reads “if all formulas in Γ are satisfied, then so is φ .” This notion can be given a formal semantics in terms of its valuations, i.e., (characteristic functions of) subsets of its propositional variables. Classical Tarskian semantics provides us with an inductive way to associate each formula φ with a set

of valuations – i.e., its *extension* $\llbracket \varphi \rrbracket$ — that is in accordance with the truth-functional readings of the connectives.³ We say a formula holds in a valuation s , if s is an element of the extension of φ , i.e., if $s \in \llbracket \varphi \rrbracket$. Assuming classical consequence being given independently of a specific semantical or proof theoretical characterization, the following soundness and completeness result is obtained:

$$\Gamma \vdash_{\text{CPC}} \varphi \quad \text{iff} \quad \bigcap_{\gamma \in \Gamma} \llbracket \gamma \rrbracket \subseteq \llbracket \varphi \rrbracket.$$

In the sequel we will write $\llbracket \Gamma \rrbracket$ for $\bigcap_{\gamma \in \Gamma} \llbracket \gamma \rrbracket$.

Alternatively, each formula can be interpreted as a binary relation on the set of valuations and define a semantics for logical consequence in terms of the maximum elements of these relations.⁴ Define for each formula φ a relation $\rho(\varphi)$ as follows, where s and s' range over valuations:

$$\rho(\varphi) =_{df.} \begin{cases} \{(s, s') : s \in \llbracket \varphi \rrbracket \text{ implies } s' \in \llbracket \varphi \rrbracket\} & \text{if } \llbracket \varphi \rrbracket \neq \emptyset, \\ \emptyset & \text{otherwise.} \end{cases}$$

Let further $\rho(\Gamma)$ be defined as $\bigcap_{\gamma \in \Gamma} \rho(\gamma)$. Then, obviously, $\Gamma \subseteq \Gamma'$ implies $\rho(\Gamma') \subseteq \rho(\Gamma)$, i.e., the larger the theory, the finer the relation it defines on the valuations. It can easily be checked that both $\rho(\varphi)$ and $\rho(\Gamma)$ are either the empty relation or a partial preorder, i.e., a reflexive and transitive relation, over the valuations. We have the following proposition and corollary.

Proposition 2. *For Γ a theory of a propositional language $L(A)$:*

$\llbracket \Gamma \rrbracket$ coincides with set of maximum elements of $\rho(\Gamma)$.

Proof. First assume $\llbracket \Gamma \rrbracket$ to be empty. Assume further for a reductio ad absurdum that s is a maximum element of $\rho(\Gamma)$ and consider an arbitrary $\gamma \in \Gamma$. Then, $(s', s) \in \rho(\gamma)$, for all valuations s' . So, in

³The resulting semantics is, of course, as follows:

$$\begin{aligned} \llbracket a \rrbracket &=_{df.} \{s \in 2^A : a \in s\} && \text{for } a \in A \\ \llbracket \neg \varphi \rrbracket &=_{df.} \overline{\llbracket \varphi \rrbracket} \\ \llbracket \varphi \wedge \psi \rrbracket &=_{df.} \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket. \end{aligned}$$

Here $\overline{\llbracket \varphi \rrbracket}$ denotes the complement of $\llbracket \varphi \rrbracket$ in 2^A , i.e., $2^A - \llbracket \varphi \rrbracket$.

⁴A *maximum* or *maximum element* of a binary relation ρ on a set S , is an element of S that is at least as great as any other in S with respect to ρ . The *maximal elements* of ρ in S are those elements than which none is greater in S with respect to ρ .

particular, $(s, s) \in \rho(\gamma)$ and from the definition of $\rho(\gamma)$ then follows that $\llbracket \gamma \rrbracket \neq \emptyset$. Hence, $s^* \in \llbracket \gamma \rrbracket$, for some s^* . Then also $(s^*, s) \in \rho(\gamma)$ and consequently $s \in \llbracket \gamma \rrbracket$ as well. With γ having been chosen as an arbitrary element of Γ , we have that $s \in \llbracket \Gamma \rrbracket$, which is at variance with the assumption that $\llbracket \Gamma \rrbracket$ be empty.

So, for the remainder of the proof we will assume $\llbracket \Gamma \rrbracket$ to be not empty. Consider an arbitrary valuation s . First assume that $s \notin \llbracket \Gamma \rrbracket$. Then $s \notin \llbracket \gamma \rrbracket$, for some $\gamma \in \Gamma$. With $\llbracket \Gamma \rrbracket$ not empty, we may assume there is some $s' \in \llbracket \gamma \rrbracket$. Then, however, $(s', s) \notin \rho(\gamma)$ and $(s', s) \notin \rho(\Gamma)$. Hence, s is no maximum element of $\rho(\Gamma)$. Finally, assume $s \in \llbracket \Gamma \rrbracket$. Now consider an arbitrary valuation s' along with an arbitrary $\gamma \in \Gamma$. Then, $s \in \llbracket \gamma \rrbracket$ and so $(s', s) \in \rho(\gamma)$. With γ having been chosen arbitrarily, also $(s', s) \in \rho(\Gamma)$ and, consequently, s is a maximum element of $\rho(\Gamma)$. \square

Corollary 3. *For Γ be theory and φ a formula of $L(A)$:*

$\Gamma \vdash_{\text{CPC}} \varphi$ iff the maximum elements of $\rho(\Gamma)$ are included in $\llbracket \varphi \rrbracket$.

Proof. By Proposition 2 and the definition of classical consequence. \square

Tarskian semantics for classical propositional logic disregards much of the ordinal structure a theory imposes on the valuations. This, of course, can be no censure of Tarskian semantics as a semantics for classical propositional logic. Its very soundness and completeness would belie such a claim. In the field of artificial intelligence and philosophical logic, however, semantical investigations into non-standard reasoning mechanisms have frequently found researchers in need of a richer structure on logical space. Formal analyses of default reasoning (e.g., [Veltman 1996]) and studies in non-monotonic consequence relations (cf. e.g., [Shoham 1988], [Krausetal 1990] and [Makinson 1994]) come under this heading. In this context, also qualitative decision theory (e.g., [Boutilier 1994]) and belief revision (e.g., [Gärdenfors 1988]) should be mentioned. In each of these cases the valuations that are, in a specified sense, *optimal* with respect to the additional structure play a crucial role in the definition of the key semantical concepts.

Our proposal for a game-theoretical notion of consequence is in line with these researches, be it that the structure imposed on logical space is that of a strategic game and that the notion of optimality is understood in terms of compliance with a game-theoretical solution concept.

In the formulation of the question as to the game-theoretical consequences of a family of theories, the theories are thought of as inducing




Figure 1: The extensions of three formulas φ , ψ and χ in logical space are depicted on the left. The relation $\rho(\{\varphi, \psi, \chi\})$ is indicated by the shades of grey: the darker the area, the higher ranked are the valuations in it, on the understanding that the valuations in $[[\varphi] - [[\psi]$ and those in $[[\psi] - [[\varphi]$ are incomparable. On the right, each block indicates a set of valuations still possible as outcomes given particular choices for the variables not controlled by the player with $\rho(\{\varphi, \psi, \chi\})$ as preference relation. The dark areas represent his maximum responses.

the players' preferences over the valuations. We will be particularly interested in the maximum elements of any such relation within certain subsets of the valuations, viz., the maximum elements in those subsets that contain the valuations that are still possible outcomes given a particular choice of strategy for all but one player. Whether a player is able to achieve an outcome he prefers above all others, may well depend on the decisions of his opponents. Moreover, the best an individual can achieve relative to some set of fixed values for the other players' variables may be inferior to what he can achieve relative to some other set of values for those variables. Any such locally optimal, but globally lesser optimal outcome constitutes a significant game-theoretical datum. In view of Corollary 3, however, the extension of a theory merely contains a player's most preferred outcomes, independently of her powers or the others players' preferences. By contrast, the relation $\rho(\Gamma)$ enables one to single out the maximum valuations within *any* subset of valuations, even if this set is disjoint from the extension $[[\Gamma]$ of Γ . In particular, it enables us to identify for each particular choice of strategy by the opponents, which are a player's maximum responses. Figure 1 illustrates this point graphically.

3 Game-Theoretical Consequence

The fundamental idea underlying the concept of game-theoretical consequence, to be introduced presently, is that, given a distribution of control over the propositional variables, each family of theories defines a strategic game, with the valuations as strategy profiles. This game is then amenable to game-theoretical analysis. To make this concept formally precise, let π be a partition of a set A of propositional variables and $\{T_i\}_{i \in \pi}$ a family of theories of $L(A)$. We define a *distributed evaluation game* for the propositional language $L(A)$ as a strategic game $(\pi, \{S_i\}_{i \in \pi}, \{\rho(T_i)\}_{i \in \pi})$, with for each i in π , S_i the set of i 's strategies given by 2^i , the powerset of the set of propositional variables in i . We will denote this game by $G(\{T_i\}_{i \in \pi})$.

We are now in a position to define formally the central notion of this paper: our concept of game-theoretical consequence.

Definition 1. Let $L(A)$ be a propositional language and π a partition of the set A of propositional variables. Let further $\{T_i\}_{i \in \pi}$ be a family of theories of $L(A)$ indexed by π . Then for each formula φ of $L(A)$:

$$\{T_i\}_{i \in \pi} \models \varphi \text{ iff } \varphi \text{ holds in all maximum equilibria of } G(\{T_i\}_{i \in \pi}).$$

The concept of unilateral deviation by a player from a strategy profile, on which the concept of maximum equilibrium relies, can be represented in neat set-theoretical terms. For s, s' and s'' strategy profiles (i.e., valuations) of a distributed evaluation game $G(\{T_i\}_{i \in \pi})$ and for i a player, c.q., a block, in π :

$$(s_{-i}, s'_i) = (s \cap \bar{i}) \cup (s' \cap i).$$

We now have the following proposition, which establishes game-theoretical consequence as a conservative extension of the consequence relation of classical propositional logic. Intuitively, it says that the game-theoretical problem of consequence reduces to that of classical consequence if there is only one player who wields control over all propositional variables.

Proposition 4. Let Γ be a theory and φ a formula in a propositional language in the propositional variables A . Then:

$$\{\Gamma_A\} \models \varphi \text{ iff } \Gamma \vdash_{\text{CPC}} \varphi.$$

Proof. Since in $G(\{\Gamma_A\})$ there is only one player, the set of maximum responses of A is identical with the set of maximum equilibria. Hence, in virtue of Proposition 2, it suffices to prove that the set of maximum responses of A in $G(\{\Gamma_A\})$ coincides with the set of maximum elements of $\rho(\Gamma_A)$. First observe that $(s_{-A}, s'_A) = (s \cap \bar{A}) \cup (s' \cap A) = s'$, for any two valuations s and s' . Hence, for any valuation s :

- s is a maximum element of $\rho(\Gamma_A)$
- iff for all s' : $(s', s) \in \rho(\Gamma_A)$
- iff for all s' : $((s_{-A}, s'_A), s) \in \rho(\Gamma_A)$
- iff s is a maximum response for A in $G(\{\Gamma_A\})$.

This concludes the proof. □

The next proposition also connects game-theoretical consequence and classical propositional logic. It guarantees the extrapolation of negative facts about the former to the latter. E.g., as a consequence of Proposition 5 we find that game-theoretical consequence is consistent, i.e., in general, $\{\Gamma_i\}_{i \in \pi} \not\models \perp$, if $\Gamma_i = \emptyset$, for all $i \in \pi$.

Proposition 5. *Let π a partition of A and $\{\Gamma_i\}_{i \in \pi}$ a family of theories in $L(A)$. Then, for all formulas φ :*

$$\{\Gamma_i\}_{i \in \pi} \models \varphi \text{ implies } \bigcup_{i \in \pi} \Gamma_i \vdash_{\text{CPC}} \varphi.$$

Proof. It suffices to show that $\bigcap_{i \in \pi} \llbracket \Gamma_i \rrbracket$ is contained in the set of maximum equilibria of $G(\{\Gamma_i\}_{i \in \pi})$. If $\bigcap_{i \in \pi} \llbracket \Gamma_i \rrbracket$ is empty, the proof is trivial, so assume $\bigcap_{i \in \pi} \llbracket \Gamma_i \rrbracket$ to be none empty. Consider an arbitrary valuation s and assume $s \in \bigcap_{i \in \pi} \llbracket \Gamma_i \rrbracket$. Consider an arbitrary $i \in \pi$ and an arbitrary $\gamma \in \Gamma_i$. Then, $s \in \llbracket \gamma \rrbracket$. Hence, $\llbracket \gamma \rrbracket \neq \emptyset$, and for all valuations s' , $(s', s) \in \rho(\gamma)$. With γ and i having been chosen arbitrarily it follows, subsequently, that s is a maximum response for i and that s is a maximum equilibrium in $G(\{\Gamma_i\}_{i \in \pi})$ as well. □

Game-theoretical consequence is *monotonic* and *finite* (or *compact*). The former is established by Proposition 6; the latter we state as a fact without giving a detailed proof.⁵ The general idea is, however, that for

⁵In another paper [Harrenstein 2002] the author gives a detailed proof for a notion very similar to that of game-theoretical consequence as it is presented here.

each formula γ and each block i in a partition π of propositional variables there is another formula γ^i , such that in general:

$$\{\Gamma_i\}_{i \in \pi} \models \varphi \text{ iff } \bigcup_{i \in \pi} \{\gamma^i : \gamma \in \Gamma_i\} \vdash_{\text{CPC}} \varphi.$$

Game-theoretical consequence then inherits finiteness from CPC.

Proposition 6 (Monotony). *Let π be a partition of the set A of propositional variables in $L(A)$ and let $\{\Gamma_i\}_{i \in \pi}$ and $\{\Gamma'_i\}_{i \in \pi}$ be families of theories such that $\Gamma_i \subseteq \Gamma'_i$, for all $i \in \pi$. Then for all formulas φ :*

$$\{\Gamma_i\}_{i \in \pi} \models \varphi \text{ implies } \{\{\Gamma'_i\}_{i \in \pi} \models \varphi\}.$$

Proof. Observe that $\rho(\Gamma'_i) \subseteq \rho(\Gamma_i)$, for each $i \in \pi$. Then the claim follows from Proposition 1 and the definition of game-theoretical consequence. \square

Fact 7 (Finiteness). *Let π a partition of the propositional variables of $L(A)$ and let $\{\Gamma_i\}_{i \in \pi}$ be a family of theories and φ a formula of $L(A)$. Then $\{\Gamma_i\}_{i \in \pi} \models \varphi$ implies that, for each $i \in \pi$, there is a finite $\Gamma'_i \subseteq \Gamma_i$ such that $\{\Gamma'_i\}_{i \in \pi} \models \varphi$.*

Another important property of classical consequence is that of *reflexivity*, i.e., $\varphi \vdash \varphi$, for all formulas φ . However, for game-theoretical consequence it is not in general the case that $\{\psi_i\}_{i \in \pi} \models \varphi$, even if ψ_i is syntactically identical to φ , for each $i \in \pi$. Consider the following example.

Example 2. Consider again the game of Example 1, but now assume that both players seek to satisfy $\{a \wedge b\}$. The corresponding game can now be summarized in the following matrix:

	\emptyset	$\{b\}$
\emptyset	$\mathbf{0}$ $\mathbf{0}$	$\mathbf{0}$ $\mathbf{0}$
$\{a\}$	$\mathbf{0}$ $\mathbf{1}$	$\mathbf{0}$ $\mathbf{1}$

The valuation \emptyset is a maximum equilibrium in this game. However, $\emptyset \notin \llbracket a \wedge b \rrbracket$. Accordingly, $\{\{a \wedge b\}_{\{a\}}, \{a \wedge b\}_{\{b\}}\} \not\models a \wedge b$.

4 Varying Partitions and Coalitions

A distinguishing feature of game-theoretical consequence is the distribution of control over the propositional variables. The families of theories game-theoretical consequence concerns, may be indexed by different partitions of the propositional variables. A formal topic that suggests itself is how the set of game-theoretical consequences of a family of theories indexed by a partition relates to the set of game-theoretical consequences of a family of theories indexed by another partition. In order to assay this issue with some success, we need a firm grip on the ways theories can systematically be combined into one theory and how a theory can be distributed over various theories. We argue that social choice theory furnishes us with some of the concepts needed.

In this context recall that the partitions over a set can be ordered as a complete lattice with respect to their coarseness as follows. For π and π' partitions over some set define:

$$\pi \leq \pi' \text{ iff for all } x \in \pi \text{ there is a } y \in \pi' \text{ such that } x \subseteq y.$$

Intuitively, $\pi \leq \pi'$ denotes that π is at least as fine as π' .

So, for any partitions π and π' such that $\pi \leq \pi'$, any element j of π' is the union of elements of π . As we identified partitions of propositional variables and sets of players, each element of π' could be conceived of as a *coalition* in which players in π have joined forces. But, if coalitions are to be regarded as fully fledged participants in strategic situations, they should also be ascribed preferences over the possible outcomes. Thus, the question remains what the collective preferences are like and how they depend on the preferences of the members of the coalition.

One of the fundamental concepts in this context is that of the (strong) *Pareto property* (cf., [Arrow 1963], [Kelly 1987]). The collective preference relation satisfies this property if and only if the coalition as a whole prefers x to y , if every member of the coalition does. An admittedly rather blunt way of combining preferences that satisfies the Pareto property is by taking the latter as both a sufficient and a necessary condition for coalitional preference relations, i.e., by simply intersecting the coalition members' preference relations.⁶ This procedure results in at most

⁶Be it noted that the intersection of a set of relations is not in general a total preorder, not even if all the relations in the set are preorders themselves. This would disqualify intersection as defining a feasible social choice function as they are commonly understood. In our analysis, however, we merely required the preference orders to be reflexive and transitive and the intersection of any set of reflexive and transitive relations is again a reflexive and transitive relation.

a reduction of the number of the maximum equilibria of a distributed evaluation game.

Proposition 8. *Let π and π' be partitions of some set A such that $\pi \leq \pi'$. Let G and G' be the games $(\pi, \{S_i\}_{i \in \pi}, \{\rho_i\}_{i \in \pi})$ and $(\pi', \{S_j\}_{j \in \pi'}, \{\rho_j\}_{j \in \pi'})$, where for each k in either π or π' , $S_k = 2^k$. Let further for each $j \in \pi'$ the preference relation be defined as the intersection of the preference relations of its constituent members, i.e., $\rho_j =_{df.} \bigcap_{\substack{i \in \pi \\ i \subseteq j}} \rho_i$. Then, the maximum equilibria in G' are also maximum equilibria in G .*

Proof. By contraposition. Assume for some strategy profile s that it be no maximum equilibrium in G . Then there is some $i \in \pi$ and some strategy profile s' such that $((s_{-i}, s'_i), s) \notin \rho_i$. Consider this i and s' . Then, $((s_{-i}, s'_i), s) \notin \bigcap_{\substack{i \in \pi \\ i \subseteq j}} \rho_i$. Now consider the unique $j \in \pi'$ such that $i \subseteq j$. Observe that $(s_{-i}, s'_i) = (s_{-j}, (s_{-i}, s'_i)_j)$. Hence, $((s_{-j}, (s_{-i}, s'_i)_j), s) \notin \bigcap_{\substack{i \in \pi \\ i \subseteq j}} \rho_i$. That is $(s_{-j}, (s_{-i}, s'_i)_j) \not\leq_j s$. We may conclude that s is not a maximum equilibrium in G' either. \square

This proposition concerns games and their equilibria but it has a logical counterpart in the following corollary.

Corollary 9. *Let π and π' be partitions of a set of propositional variables A such that $\pi \leq \pi'$. For $\{\Gamma_i\}_{i \in \pi}$ a family of theories, define for each $j \in \pi'$, $\Gamma_j^* =_{df.} \bigcup_{\substack{i \in \pi \\ i \subseteq j}} \Gamma_i$. Then, for all formulas φ :*

$$\{\Gamma_i\}_{i \in \pi} \models \varphi \text{ implies } \{\Gamma_j^*\}_{j \in \pi'} \models_{\pi'} \varphi.$$

Proof. Observe that the relation associated with each Γ_j^* is identical to $\bigcap_{\substack{i \in \pi \\ i \subseteq j}} \rho(\Gamma_i)$. The proposition then follows from Proposition 8. \square

The converses of neither Proposition 8 nor Corollary 9 hold in general. Due to its greater strategic power, a coalition can reach a greater number of strategy profiles by unilaterally deviating than any of its members can. Among this greater number there may be strategy profiles that may render deviation from a particular strategy profile s attractive, whereas no such strategy profiles was accessible from s by any of the coalition's members alone. This phenomenon even occurs if all members of the coalition have the same preferences over the outcomes.

5 Conclusion

In this paper we proposed a concept of logical consequence based on the game-theoretical notion of maximum equilibrium. Classical consequence was proved to be a special case of game-theoretical consequence. From this perspective, it stands to reason to investigate game-theoretical consequence using the standard logical techniques and concepts. The issue of sound and complete formal and axiomatic systems for it is still very much open in this respect.

Game-theoretical consequence, however, also raises some issues of its own, for the proper treatment of which it would seem that concepts from other sciences should be employed. We have already mentioned social choice theory as a possible conceptual source to get a firm grasp of how to combine and distribute theories, if the latter are looked upon as representing preference orders.

In distributed evaluation games the players were identified with the variables they control. The emphasis has so far been on the set maximum equilibria given different theories defining the preferences of the players. We could also invert this image, and take the preferences of players as fixed and investigate the sets of maximum equilibria by varying assignments of the variables to the players. Game theory may here provide the apposite concepts.

Another issue is that of the existence of maximum equilibria in distributed evaluation games. This is the game-theoretical counterpart of the issue of satisfiability in classical logic. Maximum equilibria in pure strategies do not in general exist, and only pure strategies we considered. Lattice theoretic restrictions may be imposed on the strategies and preferences of players so that the existence of equilibria is guaranteed (cf. [Topkis 1998], [Fudenberg & Tirole 1991]). An example is the lattice-theoretical concept of (quasi-)supermodularity, which is closely related to economic notion of complementarity. These reflections, however, raise the question what these concepts correspond to on a logical level.

Game-theoretical consequence provides a generalization of classical logic, in the study of which we argued concepts from game-theory, economics and social choice theory become relevant and apposite.