

Local Equilibria in Logic-Based Multi-Player Games

Julian Gutierrez, Paul Harrenstein, Thomas Steeples, Michael Wooldridge
University of Oxford
Oxford, UK

ABSTRACT

Game theory provides a well-established framework for the analysis and verification of concurrent and multi-agent systems. Typically, the analysis of a multi-agent system involves computing the set of equilibria in the associated multi-player game representing the behaviour of the system. As systems grow larger, it becomes increasingly harder to find equilibria in the game – which represent the rationally stable behaviours of the multi-agent system (the solutions of the game). To address this issue, in this paper, we study the concept of *local equilibria*, which are defined with respect to (maximal) stable coalitions of agents for which an equilibrium can be found. We focus on the solutions given by the Nash equilibria of Boolean games and iterated Boolean games, two logic-based models for multi-agent systems, in which the players’ goals are given by formulae of propositional logic and LTL, respectively.

KEYWORDS

Iterated Boolean Games; Nash Equilibria; Formal Verification

ACM Reference Format:

Julian Gutierrez, Paul Harrenstein, Thomas Steeples, Michael Wooldridge. 2018. Local Equilibria in Logic-Based Multi-Player Games. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, Stockholm, Sweden, July 10–15, 2018, IFAAMAS, 8 pages.

1 INTRODUCTION

Game theory [15] provides an important framework for the analysis of concurrent and multi-agent systems. Within this framework, a concurrent and multi-agent system is viewed as a game, where agents correspond to players, system executions to plays, and individual agent behaviours are modelled as player strategies, which are used to resolve the possible choices available to each player. Since agents are assumed to be acting strategically, game theory proposes to analyse the behaviour of these systems using solution concepts based on the idea of *equilibria*, amongst which Nash equilibrium is the best known and most widely used [15].

This is certainly a very powerful and interesting framework, with several applications in computer science, artificial intelligence, and multi-agent systems research [14, 18, 19]. Typically, analysing a concurrent and multi-agent system in this game-theoretic setting boils down to computing the set of (Nash) equilibria in the associated multi-player game. However, as systems grow large it becomes increasingly harder to compute – or even ensure the existence of – a single equilibrium representing a possible stable behaviour of the multi-agent system at hand.

Consider, for instance, of a multi-agent system consisting of $k + n$ agents, with k much larger than n , where a stable behaviour could

be found but only “locally” with respect to k players in the system. To make this idea more concrete, let us illustrate the situation with an example. Suppose you have a game with player set $\{1, \dots, k, k + 1, k + 2\}$, where each agent i has only two strategies, namely either $f_i = a$ or $f_i = b$, and a utility function utility u_i defined as follows:

$$u_i(f_1, \dots, f_k, f_{k+1}, f_{k+2}) = \begin{cases} 1 & \text{if } f_i = f_j \text{ for all } j \leq k, \text{ and } i \leq k \\ 1 & \text{if } f_i = f_{k+2} \text{ and } i = k + 1 \\ 1 & \text{if } f_i \neq f_{k+1} \text{ and } i = k + 2 \\ 0 & \text{otherwise.} \end{cases}$$

Informally, in this game the first k players are playing a coordination game (they all wish to play the same action, either a or b) whereas the last two players are playing a ‘matching pennies’ game against each other (while player $k + 1$ wishes to match player $k + 2$ ’s choice, player $k + 2$ desires both choices to be different). Since in pure strategies there is no Nash equilibrium for the matching pennies (sub)game, the whole game, including the $k + 2$ players, does not have a Nash equilibrium. Nevertheless, it should be easy to see that there is a ‘local’ equilibrium between the first k players in the game: simply set $f_i = f_j$, for all $1 \leq i, j \leq k$. No matter which actions players $k + 1$ and $k + 2$ choose, the other players in the game (players 1 to k) have no incentive to deviate. However, in pure strategies, a solution (e.g., a Nash equilibrium) for the whole system does not exist due to the irreconcilably antagonistic behaviour of players $k + 1$ and $k + 2$.

To address this issue, we introduce and study the concept of *local equilibrium*, which, essentially, formalises the intuitive idea that an equilibrium can be defined/found ‘locally’ with respect to a subset of the players in a game. We investigate two complementary notions of local equilibrium, an *existential* one and a *universal* one. In the existential case, we consider a local equilibrium for a given set of players with respect to *some* behaviour of the remaining players in the game. In the universal case, by contrast, we ask whether, for given a set of players, it is possible to find a local equilibrium for *every* behaviour of the remaining players in the game. As we are interested in local equilibria that are as inclusive as possible, we also define a very general notion of *maximality* that allows us to compare different local equilibria and select the most appropriate ones (e.g., local equilibria with the biggest number of players in equilibrium).

In particular, we study the complexity of computing a local equilibrium (formal definitions will be given later) in the logic-based game-theoretic settings provided by Boolean games (BGs [9, 10]) and by iterated Boolean games (iBGs [7]), two models for concurrent and multi-agent system where players’ goals are given, respectively, by propositional logic and by Linear Temporal Logic (LTL [16]) formulae. To give a comprehensive overview of the problem at hand, we define a number of decision problems pertaining to our notion of local equilibrium and establish their computational complexity.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Our main complexity results range from problems that are coNP-complete to ones that can be solved in 3EXPTIME. They show that, in general, reasoning about local equilibria (*i.e.*, equilibria with respect to a subset of the players in the game) may be harder than reasoning about the ‘global’ notion of equilibrium. Table 1 summarises our main complexity results and can be found at the end of the paper. Some additional complexity results, along with concluding remarks, related work, and ideas for further research can also be found at the end of the paper.

2 PRELIMINARIES

In this section we introduce the main technical concepts and models used in this paper.

Valuations and Runs. Let Φ be a finite set of Boolean variables. A *valuation* for propositional logic is a set $v \in 2^\Phi$, with the interpretation that $p \in v$ means that p is true under valuation v , while $p \notin v$ means that p is false under v . By $p\bar{q}\bar{r}$ we denote the valuation in which variable p is set to true and variables q and r are set to false, and similarly for other valuations. A *run* is an infinite sequence $\rho = v_0, v_1, v_2, \dots$ of valuations. Using square brackets around parameters referring to time points, we let $\rho[t]$ denote the valuation v_t assigned to time point t by run ρ .

Linear Temporal Logic. In this paper we use *Linear Temporal Logic* (LTL [16]), which extends propositional logic with two operators, X (‘next’) and U (‘until’) so as to be able to express properties of runs. The syntax of LTL is defined with respect to a set Φ of variables as follows:

$$\phi ::= \top \mid p \mid \neg\phi \mid \phi \vee \psi \mid X\phi \mid \phi U \psi$$

where $p \in \Phi$. The remaining classical logical operators are defined as usual. Also, we write $F\phi = \top U \phi$ and $G\phi = \neg F\neg\phi$, for ‘eventually’ and ‘always’, respectively. We interpret formulae of LTL with respect to pairs (ρ, t) , where $\rho \in (2^\Phi)^\omega$ is run and $t \in \mathbb{N}$ is a temporal index into ρ , as follows:

$$\begin{aligned} (\rho, t) &\models \top \\ (\rho, t) &\models p && \text{iff } p \in \rho[t] \\ (\rho, t) &\models \neg\phi && \text{iff it is not the case that } (\rho, t) \models \phi \\ (\rho, t) &\models \phi \vee \psi && \text{iff } (\rho, t) \models \phi \text{ or } (\rho, t) \models \psi \\ (\rho, t) &\models X\phi && \text{iff } (\rho, t+1) \models \phi \\ (\rho, t) &\models \phi U \psi && \text{iff for some } t' \geq t : ((\rho, t') \models \psi \text{ and} \\ &&& \text{for all } t \leq t'' < t' : (\rho, t'') \models \phi). \end{aligned}$$

If $(\rho, 0) \models \phi$, we write $\rho \models \phi$ and say that ρ *satisfies* ϕ . An LTL formula ϕ is *satisfiable* if there is a run satisfying ϕ .

Iterated Boolean Games. An *iterated Boolean game* (iBG [7]) is a tuple

$$G = (N, \Phi, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n),$$

where $N = \{1, \dots, n\}$ is a set of agents, Φ is a finite and non-empty set of Boolean variables, and for each agent $i \in N$, Φ_i is the set of Boolean variables uniquely controlled by i . The *choices* available to agent i are then given by the different ways i can choose truth values for the variables under her control, *i.e.*, by the valuations $v_i \subseteq 2^{\Phi_i}$. We require that Φ_1, \dots, Φ_n forms a partition of Φ . Finally, γ_i is the LTL goal that agent i aims to see satisfied by choosing her strategy.

A strategy σ_i for agent i is a complete plan how to make choices for the variables under her control over time depending on the choices made by the other agents previously. Thus, games are played by each agent selecting a strategy. Formally a *strategy* σ_i for agent i is a finite state machine $(Q_i, q_i^0, \delta_i, \tau_i)$ with output (a transducer), where Q_i is a finite and non-empty set of *states*, $q_i^0 \in Q_i$ is the *initial state*, $\delta_i : Q_i \times 2^\Phi \rightarrow Q_i$ is a deterministic *transition function*, and $\tau_i : Q_i \rightarrow 2^{\Phi_i}$ is an *output function*. Let Σ_i be the set of strategies for agent i . Once every agent i has selected a strategy σ_i , a *strategy profile* $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$ results and the game has an *outcome*, which we will denote by $\rho(\vec{\sigma})$. Because strategies are deterministic, $\rho(\vec{\sigma})$ is the unique run over Φ induced by $\vec{\sigma}$, *i.e.*, the infinite run v_0, v_1, v_2, \dots such that

$$\begin{aligned} v_0 &= \tau_1(q_1^0) \cup \dots \cup \tau_n(q_n^0) \\ v_{k+1} &= \tau_1(q_1^{k+1}) \cup \dots \cup \tau_n(q_n^{k+1}), \end{aligned}$$

where $q_i^{k+1} = \delta_i(q_i^k, v_k)$, for every agent i .

Nash equilibrium. Since the outcome of a game determines whether each goal γ_i is or is not satisfied, we can now define a preference relation \geq_i over outcomes for each agent i with goal γ_i . For strategy profiles $\vec{\sigma}$ and $\vec{\sigma}'$, we have

$$\rho(\vec{\sigma}) \geq_i \rho(\vec{\sigma}') \quad \text{if and only if } \rho(\vec{\sigma}') \models \gamma_i \text{ implies } \rho(\vec{\sigma}) \models \gamma_i.$$

On this basis, we define the concept of Nash equilibrium [15]: a strategy profile $\vec{\sigma}$ is a *Nash equilibrium* of G if and only if for every agent i and every strategy σ'_i , we have

$$\rho(\vec{\sigma}) \geq_i \rho((\vec{\sigma}_{-i}, \sigma'_i)),$$

where $(\vec{\sigma}_{-i}, \sigma'_i)$ denotes $(\sigma_1, \dots, \sigma_{i-1}, \sigma'_i, \sigma_{i+1}, \dots, \sigma_n)$, the strategy profile where the strategy of agent i in $\vec{\sigma}$ is replaced by σ'_i . Let $NE(G)$ denote the set of Nash equilibria of G .

Boolean games. The source of inspiration behind the iBGs model was the simpler model of strategic interaction provided by Boolean games (BGs [9, 10]). In Boolean games, the agents’ goals are given by propositional logic formulae and plays have only one round. In this more basic setting, a strategy σ_i for agent i is represented as a valuation over Φ_i , *i.e.*, $\sigma_i \in 2^{\Phi_i}$, and an outcome $\rho(\vec{\sigma})$ of a BG is just a valuation over Φ , *i.e.*, $\rho(\vec{\sigma}) \in 2^\Phi$. We will generally not notationally distinguish between profiles and valuations. On this basis, the concept of Nash equilibrium can be defined just as for iBGs by letting ‘ \models ’ be the satisfaction relation for propositional logic formulae.

3 LOCAL EQUILIBRIA AND \exists/\forall -STABILITY

The main concern of this paper is with local equilibria in games, be they Boolean games or iterated Boolean games. We formalise locality as a property of *coalitions*, *i.e.*, of groups of agents. For technical convenience we also allow the empty set \emptyset of agents as a coalition that controls the empty set \emptyset of propositional variables, and consequently has one joint strategy at their disposal, *viz.*, \emptyset .

Local Equilibria. A strategy profile $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$ is an equilibrium local to a coalition C whenever no agent within C can profit by deviating from $\vec{\sigma}$. Formally, we say that a strategy profile $\vec{\sigma}$ is a *C-equilibrium* if $\vec{\sigma} \geq_i (\vec{\sigma}_{-i}, \sigma'_i)$, for every agent i in C and every strategy σ'_i for i . Clearly, a profile is a Nash equilibrium if and only if it

is an N -equilibrium. The formal connection between C -equilibrium and Nash equilibrium can also be made in a slightly different way. For a game $G = (N, \Phi, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n)$ and coalition C , we can define the game $G_C = (N, \Phi, \Phi_1, \dots, \Phi_n, \gamma'_1, \dots, \gamma'_n)$, where

$$\gamma'_i = \begin{cases} \gamma_i, & \text{if } i \in C, \\ \top, & \text{otherwise.} \end{cases}$$

Thus, we see that in G_C the agents who do not belong to C are fully indifferent. We now have the following lemma:

LEMMA 3.1. *A strategy profile $\vec{\sigma}$ is a C -equilibrium of game G if and only if it is a Nash Equilibrium of G_C .*

PROOF. Let $\vec{\sigma}$ be a C -Nash equilibrium of G , and G_C as defined above. For all $i \in C$ and every strategy σ'_i for i , if $\vec{\sigma} \not\models \gamma_i$, then $(\vec{\sigma}_{-i}, \sigma'_i) \not\models \gamma_i$. For these agents $\gamma_i = \gamma'_i$, and thus the Nash equilibrium condition is satisfied for them. Furthermore, trivially, $\vec{\sigma} \models \top$ and, hence, $\vec{\sigma} \models \gamma'_i$ for all $i \notin C$. Thus, for all agents $i \in N$, if $\vec{\sigma} \not\models \gamma'_i$, then $(\vec{\sigma}_{-i}, \sigma'_i) \not\models \gamma'_i$, for every strategy σ'_i for agent i . It follows that $\vec{\sigma}$ is a Nash equilibrium of G_C .

Conversely, suppose $\vec{\sigma}$ is a Nash equilibrium of G_C . Then for all agents $i \in N$, if $\vec{\sigma} \not\models \gamma'_i$, then $(\vec{\sigma}_{-i}, \sigma'_i) \not\models \gamma'_i$, for every strategy σ'_i for agent i . Since $C \subseteq N$, and the goals of those agents in C are the same in both G and G_C , it immediately follows that $\vec{\sigma}$ is a C -Nash equilibrium of G . \square

\exists - and \forall -Stable Coalitions. Observe that whether a profile is a C -Nash equilibrium is, by definition, independent of the preferences of the agents that are not in C . Yet, as we will see, one profile $\vec{\sigma}$ may be a C -Nash equilibrium and another $\vec{\sigma}'$ not, even if every agent in C plays the same strategy in $\vec{\sigma}$ and $\vec{\sigma}'$. We therefore distinguish between \forall -stable and \exists -stable coalitions. A coalition C is \exists -stable if there is some profile $\vec{\sigma}_{-C}$ for the players outside C and some profile $\vec{\sigma}'_C$ for the players in C such that $(\vec{\sigma}_{-C}, \vec{\sigma}'_C)$ is a C -equilibrium. Similarly, coalition C is \forall -stable if for every profile $\vec{\sigma}_{-C}$ for the players outside C there is some profile $\vec{\sigma}'_C$ for the players in C such that $(\vec{\sigma}_{-C}, \vec{\sigma}'_C)$ is a C -equilibrium.

For an example, consider the Boolean game in Figure 1, where the strategy profile pqr is a $\{1, 2\}$ -Nash equilibrium: player 1 and 2 both get their goal satisfied and do not want to deviate. However, player 3 would like to deviate and play \bar{r} and, hence, pqr is not a Nash equilibrium. Furthermore, coalition $\{1, 2\}$ is \exists -stable but not \forall -stable. To see the former, observe that if player 3 sets r to true, players 1 and 2 could set their variables, p and q , respectively, to true as well. Having seen that pqr is a $\{1, 2\}$ -equilibrium, it follows that $\{1, 2\}$ is \exists -stable. On the other hand, $\{1, 2\}$ fails to be \forall -stable: if player 3 sets r to false, players 1 and 2 are caught in a ‘matching pennies’ type of situation, for which no $\{1, 2\}$ -Nash equilibrium exists.

It is worth observing that singleton coalitions are always both \exists -stable and \forall -stable, as a single agent always has a strategy at her disposal that maximises her payoff for each given profile of strategies the other players may be playing. As Nash equilibria, and therewith N -equilibria, are not guaranteed to exist, it follows that neither \exists -stability nor \forall -stability are *upward monotonic* in the sense that if a coalition C is \exists/\forall -stable, then so is every coalition $D \supseteq C$. The example above, moreover, shows that \forall -stability is not *downward monotonic* either in the sense that C being \forall -stable

	q	\bar{q}	
p	1, 1, 0	0, 0, 1	
\bar{p}	0, 0, 1	1, 1, 0	
	r		

	q	\bar{q}	
	0, 1, 1	1, 0, 0	
	1, 0, 0	0, 1, 0	
	\bar{r}		

Figure 1: A three-player Boolean game, where player 1 controls p and chooses rows, player 2 controls q and chooses columns, and player 3 controls r and chooses matrices.

implies every coalition $D \subseteq C$ to be \forall -monotonic: since $\bar{p}\bar{q}r$ is an N -equilibrium, it follows that N is \forall -stable; however, coalition $\{1, 2\}$ is not. On the other hand, we do have the following lemma with respect to \exists -stable coalitions.

LEMMA 3.2. *\exists -stability is downward monotonic.*

PROOF. Assume that coalition C is \exists -stable and $D \subseteq C$. Then, some C -equilibrium $\vec{\sigma}$ exists. That is, no player in C would like to unilaterally deviate from $\vec{\sigma}$. But then, in particular, no player in D would like to deviate from $\vec{\sigma}$ either. Hence, $\vec{\sigma}$ is a D -equilibrium and D is \exists -stable. \square

Maximal Coalitions. Among other things, Lemma 3.2 shows that a game may have many local equilibria. Then, a natural question arises: given a game with multiple local equilibria, which ones should be considered as more desirable than others, and why? Indeed, observe also that any game with n players has at least $n + 1$ \exists -stable Nash equilibria (one for each singleton coalition plus one for the empty coalition). In order to address this issue, we now define a cost function, $\kappa : N \rightarrow \mathbb{R}$, which indicates how ‘valuable’ a player is to the system designer. This cost function will, in turn, allow us to measure in a formal and mathematically concrete way how valuable a (local) equilibrium is to us. Every cost function κ over players is straightforwardly extended to a cost function over coalitions in such that, for every coalition C ,

$$\kappa(C) = \sum_{i \in C} \kappa(i).$$

An \exists -stable coalition C in a game G is then called κ -maximal if there is no \exists -stable coalition D in G with $\kappa(D) > \kappa(C)$. Similarly, an \forall -stable coalition C is said to be κ -maximal if there is no \forall -stable coalition D in G with $\kappa(D) > \kappa(C)$.

One may now wonder why these should be useful notions to study. A first point to consider is that the existence of Nash equilibria is not guaranteed in (iterated) Boolean games. There may be ‘stubborn’ players who sabotage everyone else’s goal if they can not achieve their goal, or it might just be the case that two players’ goals are incompatible. If we can coerce some players not to care about their strategies or to have them play them in a particular way, this may relax the game enough to induce the existence of a Nash equilibrium. However, we may also want to maximise the value of \exists/\forall -stable coalitions. Thus, in an important sense, once κ is defined, κ -maximal stable coalitions can be thought of as representing ‘local optima’ in the game.

3.1 Complexity Problems

Both Boolean games and iBGs can be used to model the behaviour of concurrent and multi-agent systems. Once a multi-agent system has been modelled as a multi-player game, from a game-theoretic point of view, a number of queries about the correctness and verification of the system naturally arise. Among these, the most basic one is whether the system, seen as a game, has a Nash equilibrium (cf., [20]). We will study similar questions in the context of local equilibria. In particular, we will be interested in whether some coalition of players is κ -maximal, for a given cost function κ . More specifically, we will study the following two decision problems:

(\exists/\forall)-STABLE-MAXIMAL

Given: Game G , a set $C \subseteq N$, and cost function κ .

Problem: Is C a κ -maximal (\exists/\forall)-stable coalition?

Some particular cases will be more relevant than others. For instance if $C = N$ and κ is the constant function to zero, \exists -STABLE-MAXIMAL is the NON-EMPTYNESS problem for Boolean games, *i.e.*, whether a Boolean game has a Nash equilibrium. In the context of local equilibria, another problem will also be of interest, namely, whether a set of players is or is not (\exists/\forall)-stable. This problem, which naturally arises when κ is the constant function to zero, is defined as follows:

(\exists/\forall)-STABILITY

Given: Game G and set of players $C \subseteq N$.

Problem: Is C an (\exists/\forall)-stable coalition?

Finally, another basic problem we will be interested in is deciding whether a given strategy profile is a C -equilibrium, *i.e.*, if such a profile is a local (Nash) equilibrium with respect to C . Formally, this decision problem is stated as follows:

C -MEMBERSHIP

Given: Game G , set of players C , strategy profile $\vec{\sigma}$.

Problem: Is $\vec{\sigma}$ a C -Nash equilibrium?

4 COMPUTATIONAL RESULTS

In this section we investigate the concept of local equilibria through a systematic study of the decision problems formally stated in the previous section. In order to do so, we begin with Boolean games, that is, the case where players' goals are given by propositional logic formulae.

4.1 Local Equilibria in Boolean Games

We will conduct our study of Boolean games moving from the simplest problems to the hardest, starting with C -MEMBERSHIP.

PROPOSITION 4.1. C -MEMBERSHIP is coNP-complete.

PROOF. Given a Boolean game G , the game G_C can be constructed in polynomial time. Because of Lemma 3.1, we can then decide whether $\vec{\sigma}$ is a C -equilibrium by checking if $\vec{\sigma}$ is a Nash equilibrium of G_C , which is a coNP-complete problem [3]; we will call this problem MEMBERSHIP. For hardness, we reduce from MEMBERSHIP by noting that, due to Lemma 3.1, $\langle G, \vec{\sigma} \rangle \in$ MEMBERSHIP if and only if $\langle G, N, \vec{\sigma} \rangle \in C$ -MEMBERSHIP. \square

We now consider \exists -STABILITY and show that it is as hard as checking for the existence of a Nash equilibrium in a Boolean game, a problem we will denote by NON-EMPTYNESS. As one may expect,

this result will be useful to establish the computational complexity of \exists -STABLE-MAXIMAL.

PROPOSITION 4.2. \exists -STABILITY is Σ_2^P -complete.

PROOF. Firstly, recall that coalition C is \exists -stable if and only if a C -equilibrium exists (Lemma 3.1). As $\Sigma_2^P = \text{NP}^{\text{NP}} = \text{NP}^{\text{coNP}}$, for membership in Σ_2^P it suffices to observe that a certificate for \exists -STABILITY is a strategy profile $\vec{\sigma}$, which, by Proposition 4.1, can be checked in polynomial time to be a 'yes'-instance (by conferring with a coNP-oracle whether $\vec{\sigma}$ is a C -equilibrium). For Σ_2^P -hardness, we reduce from NON-EMPTYNESS by noting that, due to Lemma 3.1, $\langle G \rangle \in$ NON-EMPTYNESS if and only if $\langle G, N \rangle \in \exists$ -STABILITY. \square

We now address the complementary problem: \forall -STABILITY. As the next result shows, for a given Boolean game G , checking $\langle G, C \rangle \in \forall$ -STABILITY is considerably different from checking whether $\langle G, C \rangle \in \exists$ -STABILITY. For technical reasons and to simplify notations, we will use Ψ and Θ as metavariables over propositional formulas with free variables. Moreover, the following construction will be useful. For a given Boolean game $G = (N, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n)$, coalition C , and strategy profile $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$, let

$$G_{\vec{\sigma}_C} = (N, \Phi_1, \dots, \Phi_n, \gamma'_1, \dots, \gamma'_n),$$

where

$$\gamma'_i = \begin{cases} \bigwedge_{p \in \sigma_i} p \wedge \bigwedge_{p \in \Phi_i \setminus \sigma_i} \neg p & \text{if } i \in C, \\ \gamma_i & \text{otherwise.} \end{cases}$$

Intuitively, the game $G_{\vec{\sigma}_C}$ differs from G only in that the members in C solely wish that their part of a particular strategy profile $\vec{\sigma}$ be played. Moreover, each member i of coalition C can enforce by itself that its goal is satisfied by playing σ_i , and accordingly will do so in every Nash equilibrium. Then, we have the following result.

PROPOSITION 4.3. \forall -STABILITY is Π_3^P -complete.

PROOF. For membership, recall that $\Pi_3^P = \text{coNP}^{\Sigma_2^P}$. We can now proceed as follows. Given Boolean game G and coalition C , guess a profile $\vec{\sigma}_{-C}$, and construct $G_{\vec{\sigma}_{-C}}$. The latter can be achieved in polynomial time. At this point, we can confer with an oracle for NON-EMPTYNESS with respect to $G_{\vec{\sigma}_{-C}}$. Then, the pair $\langle C, \vec{\sigma}_{-C} \rangle$ is a counterexample to C being \forall -stable if and only if the oracle returns "no". As NON-EMPTYNESS is Σ_2^P -complete, we are done.

For hardness, we reduce QBF $_{\forall,3}$. Let $Q = \forall X \exists Y \forall Z \Psi$ be a quantified Boolean formula. We construct a three-player Boolean game G_Q with two additional variables y and z that are not included in $X \cup Y \cup Z$. Now let:

$$\begin{aligned} \Phi_1 &= X & \gamma_1 &= \top \\ \Phi_2 &= Y \cup \{y\} & \gamma_2 &= \Psi \vee (y \leftrightarrow z) \\ \Phi_3 &= Z \cup \{z\} & \gamma_3 &= \neg Y \wedge (\neg y \leftrightarrow z) \end{aligned}$$

It can then be shown that $\{2, 3\}$ is \forall -stable in G_Q if and only if $Q = \forall X \exists Y \forall Z \Psi$ evaluates to true.

First assume that Q evaluates to true, and consider an arbitrary strategy σ_1 for player 1. Then, there is a valuation σ_2 to Y such that for all valuations σ_3 to Z , the formula Ψ is satisfied at $\sigma_1 \cup \sigma_2 \cup \sigma_3$. It is then easy to see that the strategy profile $(\sigma_1, \sigma_2, \sigma_3)$ is a $\{2, 3\}$ -equilibrium: player 2 has its goal achieved whereas player 3 will not get γ_3 satisfied by playing any other strategy.

For the opposite direction, assume that Q evaluates to false. Then, there is some assignment σ_1 to X such that for all assignments to Y , there is an assignment to Z that renders Ψ false. Consider arbitrary strategies σ_2 and σ_3 for players 2 and 3, respectively. If $(\sigma_1, \sigma_2, \sigma_3)$ does not satisfy γ_2 , player 2 can and would like to deviate to a strategy σ'_2 such that $(\sigma_1, \sigma'_2, \sigma_3) \models \Psi$ and, moreover, $y \in \sigma'_2$ if and only if $z \in \sigma_3$. If, on the other hand, strategy profile $(\sigma_1, \sigma_2, \sigma_3)$ does satisfy γ_2 , then player 3 does not satisfy its goal. In that case however, the latter can deviate by playing an appropriate strategy σ'_3 such that Φ is false under strategy profile $(\sigma_1, \sigma_2, \sigma'_3)$ and $z \in \sigma'_3$ if and only if $y \notin \sigma_2$. \square

We are now in a position to study the maximality problems for Boolean games. In order to do so, let us first consider the following auxiliary decision problem and associated lemma.

SEMI- \exists -STABLE-MAXIMAL

Given: Game G , a set of players C , and cost function κ .

Problem: Is $\kappa(C) \geq \kappa(D)$ for all \exists -stable coalitions $D \subseteq N$?

LEMMA 4.4. SEMI- \exists -STABLE-MAXIMAL is in Π_2^P .

PROOF. For membership in Π_2^P , observe that a “no” instance of SEMI- \exists -STABLE-MAXIMAL is certified by a pair $\langle D, \vec{\sigma} \rangle$, where D is a coalition with $\kappa(D) > \kappa(C)$ and profile $\vec{\sigma}$ is a D -equilibrium. Since $\Pi_2^P = \text{coNP}^{\text{NP}} = \text{coNP}^{\text{coNP}}$, given a pair $\langle D, \vec{\sigma} \rangle$, we can verify in polynomial time whether $\kappa(D) > \kappa(C)$ and, because of Lemma 4.1, also whether $\vec{\sigma}$ is a D -equilibrium by conferring with the coNP-oracle. \square

Using both \exists -STABILITY and SEMI- \exists -STABLE-MAXIMAL, we can now determine the computational complexity of the maximality problem for \exists -stable coalitions in Boolean games.

PROPOSITION 4.5. \exists -STABLE-MAXIMAL is D_2^P -complete.

PROOF. For membership in D_2^P , note that \exists -STABLE-MAXIMAL can be written as the intersection of SEMI- \exists -STABLE-MAXIMAL and \exists -STABILITY. And, by Lemma 4.4 and Proposition 4.2, we know that these problems are in Π_2^P and Σ_2^P , respectively. Then, it follows that \exists -STABLE-MAXIMAL is in D_2^P .

For hardness, we reduce from $\text{QBF}_{2, \exists}\text{-QBF}_{2, \forall}$. Firstly, let $Q = \langle Q_1, Q_2 \rangle$, with $Q_1 = \exists X_1 \forall X_2 \Theta(X_1, X_2)$, and $Q_2 = \forall Y_1 \exists Y_2 \Psi(Y_1, Y_2)$. At this point, we may assume that $X_1 \cup X_2$ and $Y_1 \cup Y_2$ are disjoint. Also, define a game, G , with four players, 1, 2, 3, and 4, and with four fresh auxiliary variables, p, q, r , and s . The players control the variables as follows,

$$\begin{aligned} \Phi_1 &= X_1 \cup \{p\} & \Phi_2 &= X_2 \cup \{q\} \\ \Phi_3 &= Y_1 \cup \{r\} & \Phi_4 &= Y_2 \cup \{s\} \end{aligned}$$

and the players' goals are given by,

$$\begin{aligned} \gamma_1 &= \Theta \vee (p \leftrightarrow q) & \gamma_2 &= \neg\Theta \wedge \neg(p \leftrightarrow q) \\ \gamma_3 &= \neg\Psi \vee \neg(r \leftrightarrow s) & \gamma_4 &= \Psi \wedge (r \leftrightarrow s). \end{aligned}$$

Obviously, this game can be constructed in polynomial time. Let, moreover, $\kappa(i) = 1$ for all players i . Now consider coalition $C = \{1, 2, 4\}$. We claim that Q_1 and Q_2 are both valid if and only if C is a κ -maximal \exists -stable coalition.

First assume that both Q_1 and Q_2 are valid. Then, there exists a valuation $v'_1 \subseteq X_1$ such that for all valuations $v_2 \subseteq X_2$, we have

$(v'_1, v_2) \models \Theta$. Let v'_2 be an arbitrary valuation of X_2 . So naturally, $(v'_1, v'_2) \models \Theta$. Now, let $v'_3 \subseteq Y_1$ be an arbitrary valuation. Since Q_2 is valid, there exists some valuation, $v'_4 \subseteq Y_2$ such that we have $(v'_3, v'_4) \models \Psi$.

Let $\vec{\sigma}^* = (\sigma_1^*, \sigma_2^*, \sigma_3^*, \sigma_4^*)$ be such that

$$\begin{aligned} \sigma_1^* &= v'_1 \cup \{p\} & \sigma_2^* &= v'_2 \cup \{q\} \\ \sigma_3^* &= v'_3 \cup \{r\} & \sigma_4^* &= v'_4 \cup \{s\}. \end{aligned}$$

We claim that $\vec{\sigma}^*$ is a C -equilibrium of G .

As $(v'_1, v'_2) \models \Theta$, player 1 has their goal satisfied at $\vec{\sigma}^*$ and would prefer not to deviate. Additionally, since for all $v_2 \subseteq X_2$, we have $(v'_1, v_2) \models \Theta$, it follows that player 2 has no deviations available to them that can force $\neg\Theta$ to be true. Therefore, player 2 would prefer not to deviate as well. Finally, since $(v'_3, v'_4) \models \Psi$, and since r and s are both set to true, player 4 has their goal satisfied and would also prefer not to deviate. It then follows that $\vec{\sigma}^*$ is a C -equilibrium.

Moreover, note that $D = N = \{1, 2, 3, 4\}$ is the only subset of N with $v_\kappa(G_C) < v_\kappa(G_D)$, where $v_\kappa(G_C) = \kappa(C)$ and $v_\kappa(G_D) = \kappa(D)$. However, D is not \exists -stable. To see this, consider an arbitrary profile $\vec{\sigma}'$. As $\forall Y_1 \exists Y_2 \Psi(Y_1, Y_2)$ is valid, player 4 can always deviate by choosing values for Y_2 and s so as to satisfy γ_4 , if $\vec{\sigma}' \not\models \gamma_4$. On the other hand, if $\vec{\sigma}' \models \gamma_4$, player 3 can deviate by choosing the opposite value for r and have γ_3 satisfied. It follows that C is κ -maximal \exists -stable.

For the opposite direction, assume that Q_1 or Q_2 is not valid. If the former, for all valuations $v_1 \subseteq X_1$, there exists some valuation $v_2 \subseteq X_2$ such that $(v_1, v_2) \models \neg\Theta$. By a similar line of reasoning as in the previous paragraph, we find that either player 1 or player 2 would like to deviate from any given strategy profile $\vec{\sigma}$. As a consequence, coalition C is not \exists -stable, let alone κ -maximal \exists -stable.

The analysis here is almost identical to the previous paragraph, so we omit the details. Since we have $\gamma_1 = \neg\gamma_2$, and since player 2 can always force their goal to be true, but player 1 can always force their goal to be true by deviating after “seeing” how player 2 has won, this implies that G_C has no Nash equilibria, and so is not a κ -maximal game.

The other possibility is that Q_2 is not valid, with Q_1 still valid. In this case, we find that the grand coalition $N = \{1, 2, 3, 4\}$ is \exists -stable and, as $\kappa(N) > \kappa(C)$, also that coalition C is not κ -maximal \exists -stable. To see that coalition D is \exists -stable, observe that, as Q_1 is valid, there exists some valuation $v'_1 \subseteq X_1$ such that for all valuations $v_2 \subseteq X_2$, we have $(v'_1, v_2) \models \Theta$. Similarly, since Q_2 is not valid, there exists some valuation $v'_3 \subseteq Y_1$, such that for all valuations $v_4 \subseteq Y_2$, we have $(v'_3, v_4) \models \Psi$. So, let $v'_2 \subseteq X_2$, and $v'_4 \subseteq Y_2$ be arbitrary valuations. Again, define the strategy profile $\vec{\sigma}^* = (\sigma_1^*, \sigma_2^*, \sigma_3^*, \sigma_4^*)$ to be such that

$$\begin{aligned} \sigma_1^* &= v_1 \cup \{p\} & \sigma_2^* &= v_2 \cup \{q\} \\ \sigma_3^* &= v_3 \cup \{r\} & \sigma_4^* &= v_4 \cup \{s\}. \end{aligned}$$

To conclude the proof we show that $\vec{\sigma}^*$ is a C -equilibrium of G . To see this, observe that both player 1 and player 3 have their goals satisfied, so they will not choose to deviate. Now, players 2 and 4, no matter how they deviate, cannot make $\neg\Theta$ and Ψ (respectively) true, and so, cannot satisfy their goals. Thus, they have no incentive to deviate either. \square

Remark: Note that cost functions are a natural extension of the concept of the cardinality of a coalition, which is captured by the special case in which the cost of each player is 1. It is this natural case that we exploit in the proof of Theorem 4.5.

To end this subsection, we now investigate the complexity of \forall -STABLE-MAXIMAL. Formally, we have the following result.

PROPOSITION 4.6. \forall -STABLE-MAXIMAL is in Π_4^P .

PROOF. A counter-example against \forall -STABLE-MAXIMAL for coalition C would either be coalition C itself if it is not \forall -stable or a coalition D with $\kappa(D) > \kappa(C)$ that is \forall -stable. Thus, membership in $\Pi_4^P = \text{coNP}^{\Sigma_3^P} = \text{coNP}^{\Pi_3^P}$ can be established simply by guessing a coalition D .

If $D = C$, then we can query the oracle whether C is \forall -stable. We have found a counter-instance if and only if the answer is “no”. If, on the other hand, $D \neq C$, we can check in polynomial time whether $\kappa(D) > \kappa(C)$ and query the oracle as to whether coalition D is \forall -stable. We have found a counter-instance if and only if both answers are positive. \square

A lower bound in Π_3^P follows from \forall -STABILITY. However, whether the problem is Π_4^P -complete or not remains unknown.

4.2 Local Equilibria in iBGs

We now study the notion of local equilibria in the more general model of iterated Boolean Games (iBGs). As in the previous subsection, we first consider the *membership problem* and then the others. However, contrarily to the case for Boolean games, in this section we will study the *maximality problems* for iBGs first and then understand the stability problems as special cases.

The first result pertains to membership, that is, checking whether a given strategy profile is a local (Nash) equilibrium of a given iBG. This problem, as well as the conventional membership problem for iBGs, can be solved in PSPACE. The reasoning in the proof is analogous to that for Boolean games, *i.e.*, as in Proposition 4.1 (since Lemma 3.1 applies to both Boolean games and iBGs uniformly), and leveraging that MEMBERSHIP for iBGs is PSPACE-hard [7].

PROPOSITION 4.7. C -MEMBERSHIP is PSPACE-complete.

This result confirms what was shown in the previous section, namely, that checking whether a given strategy profile is a local (Nash) equilibrium is as hard as checking whether a strategy profile is a Nash equilibrium. In sharp contrast to this fact, we have the following result. Whereas for Boolean games, the \exists -STABLE-MAXIMAL problem was harder than the NON-EMPTYNESS problem, unless the polynomial hierarchy is not strict, in the case of iBGs both problems are equally hard, and can be solved in 2EXPTIME. Formally, we have:

PROPOSITION 4.8. \exists -STABLE-MAXIMAL is 2EXPTIME-complete.

PROOF. For membership in 2EXPTIME, first assume we want to check whether $\langle G, C, \kappa \rangle \in \exists$ -STABLE-MAXIMAL, for some iBG G , set of players C , and cost function κ . Since the NON-EMPTYNESS problem for iBGs is 2EXPTIME-complete [7], there is some algorithm, \mathcal{A} , and some fixed, positive integer k , that given a game G , determines

whether or not $NE(G)$ is non-empty in time $T(|G|)$, where $|G|$ is the size of G , with $T(n) = O(2^{2^{n^k}})$. We use the following algorithm to determine whether G_C is a κ -maximal subgame of G . First, we use \mathcal{A} to check if $NE(G_C)$ is non-empty. If it is empty, we reject G_C . Otherwise, for every $D \subseteq N$, look at G_D , and calculate its value. If $v_\kappa(G_C) \geq v_\kappa(G_D)$, we move onto the next subset. Otherwise, we use \mathcal{A} to determine whether or not $NE_D(G)$ is non-empty. If it is, then we reject G_C . If we have done this for every subset without yet rejecting G_C , then we accept it.

Now, for a given subset $D \subseteq N$, we can use algorithm \mathcal{A} to determine whether $NE_D(G)$ is non-empty in time $T(|G_D|) \leq T(|G|)$. Additionally, determining the value of a subgame, G_D , can be done in time $O(|D|) \leq O(n)$. Finally, there are 2^n subsets of N , so this algorithm is time bounded by,

$$\begin{aligned} \sum_{D \subseteq N} O(|D|) + T(|G_D|) &\leq \sum_{D \subseteq N} O(n) + T(|G|) \\ &= O(n2^n) + 2^n T(|G|) \\ &\leq O(|G| 2^{|G|}) + 2^{|G|} T(|G|) \\ &= O\left(|G| 2^{|G|} + 2^{|G|} 2^{2^{|G|^k}}\right) \\ &= O\left(2^{2^{|G|^k}}\right). \end{aligned}$$

Here, we have also used the fact that $n < |G|$. It therefore follows that \exists -STABLE-MAXIMAL is a member of 2EXPTIME.

For hardness, we reduce from NON-EMPTYNESS for iBGs. Let $G = \langle N, \Phi, (\Phi_i)_{i \in N}, (Y_i)_{i \in N} \rangle$ be a game. Now, define a trivial cost function, $\kappa : N \rightarrow \mathbb{R}$, that is defined as $\kappa(i) = 0$, for each player i in the game, that is, the constant function to zero. We now claim that

$$G \in \text{NON-EMPTYNESS} \text{ iff } \langle G, N, \kappa \rangle \in \exists\text{-STABLE-MAXIMAL}.$$

Suppose that $G \in \text{NON-EMPTYNESS}$. So $NE(G) \neq \emptyset$. To show that $\langle G, N, \kappa \rangle \in \exists$ -STABLE-MAXIMAL, we need to demonstrate that $G_N \equiv G$ is a κ -maximal subgame of G . We already know that $NE_N(G) = NE(G) \neq \emptyset$. Let $D \subseteq N$. Then,

$$v_\kappa(G_D) = \sum_{i \in D} \kappa(i) = 0,$$

and so, $v_\kappa(G_D) = 0$, for all $D \subseteq N$. In particular, for all $D \subseteq N$ with $NE_D(G) \neq \emptyset$, we have $v_\kappa(G_N) \geq v_\kappa(G_D)$. Thus, G_N is a κ -maximal subgame of G , and therefore we obtain that $\langle G, N, \kappa \rangle \in \exists$ -STABLE-MAXIMAL.

Conversely, suppose that $\langle G, N, \kappa \rangle \in \exists$ -STABLE-MAXIMAL. By definition, $NE(G) = NE_N(G) \neq \emptyset$, and so $G \in \text{NON-EMPTYNESS}$. Moreover, this construction can evidently be done in polynomial time. Thus, \exists -STABLE-MAXIMAL is 2EXPTIME-complete. \square

As an immediate consequence of Proposition 4.8 we now obtain the following result about \exists -stable coalitions in iBGs.

COROLLARY 4.9. \exists -STABILITY is 2EXPTIME-complete.

PROOF. For membership we first let κ be the constant (cost) function to zero, *i.e.*, $\kappa(i) = 0$ for each player i in the game. Then, we observe that

$$\langle G, C \rangle \in \exists\text{-STABILITY} \text{ iff } \langle G, C, \kappa \rangle \in \exists\text{-STABLE-MAXIMAL}.$$

For the hardness part now observe that

$$G \in \text{NON-EMPTYNESS} \text{ iff } \langle G, N \rangle \in \exists\text{-STABILITY},$$

which concludes the proof. \square

The above proof shows that, for $\exists\text{-STABLE-MAXIMAL}$, the overall complexity is driven by the (sub)procedure to decide whether a coalition is \exists -stable, instead of the (sub)procedure to check whether a stable coalition is κ -maximal. We observe the same phenomenon for $\forall\text{-STABLE-MAXIMAL}$. However, in this case, the problem can be solved in 3EXPTIME.

We will first show that $\forall\text{-STABILITY}$ is in 3EXPTIME. In order to do so, we use Strategy Logic (SL [13]), in particular, over the model of iBGs. SL is a logic to reason about strategic behaviour in multi-player games. SL also extends LTL, in this case with three new operators: an existential strategy quantifier $\langle\langle x \rangle\rangle$, a universal strategy quantifier $[[x]]$, and an agent binding operator (i, x) . These three new operators can be read as “there is a strategy x ”, “for every strategy x ”, and “let agent i use the strategy associated with x ”, respectively.

Formally, SL formulae are inductively built from a set of atomic propositions Φ , variables Var , and players N , using the following grammar, where $p \in \Phi$, $x \in \text{Var}$, and $i \in N$:

$$\phi ::= \top \mid p \mid \neg\phi \mid \phi \vee \phi \mid \mathbf{X}\phi \mid \phi \mathbf{U} \phi \mid \langle\langle x \rangle\rangle\phi \mid (i, x)\phi$$

We define $[[x]]\phi \equiv \neg\langle\langle x \rangle\rangle\neg\phi$, and use all other usual classical and temporal logic abbreviations. We can now present the semantics of SL, where $\Sigma = \bigcup_{i \in N} \Sigma_i$ denotes the set of all strategies of all players $i \in N$ in an iBG. Given an iBG G , for all SL formulae ϕ , states $v \in 2^\Phi$ of G , and assignments $\chi \in \text{Asg} = (\text{Var} \cup \text{Ag}) \rightarrow \Sigma$, mapping variables and players to strategies, the relation $G, \chi, v \models \phi$ is defined as follows:

- (1) For the Boolean and temporal cases, the semantics is as for LTL formulae on iBGs (see, e.g., [7]);
- (2) For all formulae ϕ and variables $x \in \text{Var}$ we have:
 - (a) $G, \chi, v \models \langle\langle x \rangle\rangle\phi$ if there is a strategy $\sigma \in \Sigma$ such that $G, \chi[x \mapsto \sigma], v \models \phi$;
 - (b) $G, \chi, v \models [[x]]\phi$ if for all strategies $\sigma \in \Sigma$ we have that $G, \chi[x \mapsto \sigma], v \models \phi$.
- (3) For all $i \in N$ and $x \in \text{Var}$, we have $G, \chi, v \models (i, x)\phi$ if $G, \chi[i \mapsto \chi(x)], v \models \phi$.

For a sentence ϕ , we say that G at v satisfies ϕ , and write $G, v \models \phi$, if $G, \emptyset, v \models \phi$, where \emptyset is the empty assignment. Using SL we can describe \forall -stability for an instance $\langle G, C \rangle$ as follows:

$$\phi_{\forall, C} = [[-C]]\phi_{NE}$$

where $[[-C]]$ stands for $[[x]] \dots [[y]](c_1, x) \dots (c_k, y)$, with respect to the set of players $-C = \{c_1, \dots, c_k\} = N \setminus C$, and ϕ_{NE} is the usual SL formula for Nash equilibrium, namely,

$$\phi_{NE} = \langle\langle C \rangle\rangle \bigwedge_{i \in C} \neg\gamma_i \rightarrow [[z']](i, z')\neg\gamma_i$$

in which $\langle\langle C \rangle\rangle$ stands for $\langle\langle x' \rangle\rangle \dots \langle\langle y' \rangle\rangle(c'_1, x') \dots (c'_j, y')$, with $C = \{c'_1, \dots, c'_j\}$. Since formula $\phi_{\forall, C}$ has alternation depth 2, using [13], it can be (model) checked in 3EXPTIME over concurrent game structures [2], although the procedure is only polynomial in the size of the underlying concurrent game structure. Since we

can translate any iBG G into a concurrent game structure of exponential size, using the model checking procedure for SL we can model check iBGs with the same combined complexity, but with a procedure that is exponential in the size of the iBG G . Formally, we have:

LEMMA 4.10. *For iBGs, model checking an SL formula ψ with alternation depth k can be done in $(k + 1)\text{EXPTIME}$.*

We are now in a position to establish the complexity of the $\forall\text{-STABILITY}$ problem for iBGs.

PROPOSITION 4.11. *$\forall\text{-STABILITY}$ is in 3EXPTIME.*

PROOF. Lemma 4.10 yields an upper bound for $\forall\text{-STABILITY}$ via model checking of SL formulae over iBGs. Note that, for any iBG G and set of players C , we have:

$$\langle G, C \rangle \in \forall\text{-STABILITY} \text{ iff } G \text{ is a model of } \phi_{\forall, C},$$

from which membership in 3EXPTIME follows. \square

At this point, we should note that the proof of Proposition 4.8 has three parts: one to check that a coalition is \exists -stable, one to reason about κ -maximality, and the hardness part. We can see that from these the second part (to reason about maximality) also applies to \forall -stable coalitions since we can go over all coalitions $D \subseteq N$ to check if C is κ -maximal. Based on this observation, we have the following result.

PROPOSITION 4.12. *$\forall\text{-STABLE-MAXIMAL}$ is in 3EXPTIME.*

PROOF. For membership in 3EXPTIME, we can check if C is \forall -stable using Proposition 4.11 and if it is κ -maximal by comparing it with all other \forall -stable coalitions $D \subseteq N$. \square

It also follows that both $\forall\text{-STABILITY}$ and $\forall\text{-STABLE-MAXIMAL}$ are 2EXPTIME-hard. To see this, observe that $\langle G, C \rangle \in \forall\text{-STABILITY}$ if and only if $\langle G, C, \kappa \rangle \in \forall\text{-STABLE-MAXIMAL}$, e.g., when κ is the constant function to zero, and note that NON-EMPTYNESS can be reduced to $\forall\text{-STABILITY}$ when $C = N$ and κ as above.

Remark. Readers acquainted with SL may have noticed that a 2EXPTIME upper bound for $\exists\text{-STABILITY}$ can also be achieved via SL, essentially, following an approach similar to the one just used to analyse \forall -stable coalitions. In addition, such a result could be leveraged to obtain a 2EXPTIME upper bound for $\exists\text{-STABLE-MAXIMAL}$ as well. On the one hand, this solution would allow us to avoid invoking Lemma 3.1 and, hence, also the construction of the game G_C . On the other hand, it would require reasoning on the semantics of SL, which we have shown can be avoided when considering local equilibria with respect to \exists -stable coalitions. Indeed, all we require to reason about \exists -stable coalitions is to have a 2EXPTIME procedure (any!) to check for the existence of Nash equilibria in multi-player games with LTL goals.

5 CONCLUSIONS & FUTURE WORK

In this paper we have studied the new concept of *local equilibrium*, in particular in the context of Boolean games and iterated Boolean games. We have obtained a diverse landscape of complexity results, ranging from problems that are coNP-complete to problems solvable in 3EXPTIME. In particular, our results show that, in general,

	Boolean Games	iBGs
C-MEMBERSHIP	coNP	PSPACE
\exists -STABILITY	Σ_2^P	2EXPTIME
\forall -STABILITY	Π_3^P	in 3EXPTIME
\exists -STABLE-MAXIMAL	D_2^P	2EXPTIME
\forall -STABLE-MAXIMAL	in Π_4^P	in 3EXPTIME

Table 1: Overview of main complexity results.

reasoning about local (Nash) equilibria may be harder than reasoning about the “global” concept of Nash equilibrium. A summary of our main complexity results is in Table 1.

Additional complexity results. An inspection of many of the proofs underlying our results shows that minor variations can be applied so that the same complexities can be obtained for (some of) the problems we studied with respect to other solution concepts and models of games. For instance, with respect to the former, we know that the same complexities hold for the maximality problems if we consider subgame perfect Nash equilibrium or dominant strategies. On the other hand, with respect to the latter, also the same complexities hold for the maximality problems if we consider SRML games [8].

Local equilibria. To the best of our knowledge the concept of *local equilibria* has not been previously studied in the way we have presented it in this paper. However, there are games where the concept of locality has been investigated in other forms. For instance, in graphical games [5]—where games are given as graphs with each node representing a player and edges representing players’ choices—a player may be required to best respond only to the behaviour of its neighbours. This is clearly a different notion of locality, but it also embodies the idea that a player’s behaviour may be analysed only with respect to a subset of the overall set of players in the game. The concept of local (Nash) equilibrium has also been studied in the context of economic games [1], continuous games [17], and social networks [21]. These are all different notions of locality, both between them and with respect to ours. However, these other notions of locality share one common feature: they try to characterise the fact that if the space of strategies is reduced then a global optimum may not be achieved, only a local one (as constrained by the reduced space of strategies under consideration), which may be easier to compute or model more faithfully the fact that players in a game may have only bounded rationality.

Strategy Logic and Formal Verification. Our results also relate to recent work on rational synthesis [6, 11] and rational verification [7, 8, 20]. In these papers the concept of locality is not present, neither the concept of cost functions in the way we use them here. However, the solutions in such papers make use of Strategy Logic (SL), which at one point we also use here. Although SL can easily provide some upper bounds, as shown in this paper, this does not immediately mean that they are optimal (cf., \forall -STABILITY). In addition, SL does

not yield optimal upper bounds for BGs and other related settings. For these we would need another logic.

Future work. From a practical point of view, it should be possible to have an implementation of local equilibria for BGs or for games with imperfect recall using MCMAS [12], since MCMAS supports SL with memoryless strategies [4]. From a theoretical point of view, ways to lower the overall complexity of the problems we considered should be investigated. Two promising directions are to consider games with simpler types of goals, or games with simpler strategy models.

Acknowledgment.

The authors acknowledge with gratitude the support of ERC grant 291528 (“RACE”) at Oxford.

REFERENCES

- [1] Carlos Alos-Ferrer and Ana B. Ania. 2001. Local Equilibria in Economic Games. *Economic Letters* 70, 2 (2001), 165–173.
- [2] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *Journal of the ACM* 49, 5 (2002), 672–713.
- [3] Elise Bonzon, Marie-Christine Lagasque-Schiech, Jérôme Lang, and Bruno Zanuttini. 2006. Boolean Games Revisited. In *ECAI (Frontiers in Artificial Intelligence and Applications)*, Vol. 141. IOS Press, Trentino, Italy, 265–269.
- [4] Petr Cermák, Alessio Lomuscio, Fabio Mogavero, and Aniello Murano. 2014. MCMAS-SLK: A Model Checker for the Verification of Strategy Logic Specifications. In *CAV (LNCS)*, Vol. 8559. Springer, Vienna, Austria, 525–532.
- [5] Steven N. Durlauf and Lawrence E. Blume. 2010. Graphical Games. In *Game Theory*. Palgrave Macmillan Publishers, London, UK, 153–157.
- [6] Dana Fisman, Orna Kupferman, and Yoad Lustig. 2010. Rational Synthesis. In *TACAS (LNCS)*, Vol. 6015. Springer, Paphos, Cyprus, 190–204.
- [7] Julian Gutierrez, Paul Harrenstein, and Michael Wooldridge. 2015. Iterated Boolean games. *Information and Computation* 242 (2015), 53–79.
- [8] Julian Gutierrez, Paul Harrenstein, and Michael Wooldridge. 2017. From model checking to equilibrium checking: Reactive modules for rational verification. *Artificial Intelligence* 248 (2017), 123–157.
- [9] Paul Harrenstein. 2004. *Logic in Conflict*. Ph.D. Dissertation. Utrecht University, Utrecht, Netherlands.
- [10] Paul Harrenstein, Wiebe van der Hoek, John-Jules Ch. Meyer, and Cees Witteveen. 2001. Boolean Games. In *TARK*. Morgan Kaufmann Publishers, Siena, Italy, 287–298.
- [11] Orna Kupferman, Giuseppe Perelli, and Moshe Y. Vardi. 2016. Synthesis with rational environments. *Annals of Mathematics and Artificial Intelligence* 78, 1 (2016), 3–20.
- [12] Alessio Lomuscio, Hongyang Qu, and Franco Raimondi. 2017. MCMAS: An open-source model checker for the verification of multi-agent systems. *STTT* 19, 1 (2017), 9–30.
- [13] Fabio Mogavero, Aniello Murano, Giuseppe Perelli, and Moshe Y. Vardi. 2014. Reasoning About Strategies: On the Model-Checking Problem. *ACM Transactions on Computational Logic* 15, 4 (2014), 34:1–34:47.
- [14] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani (Eds.). 2007. *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK.
- [15] Martin J. Osborne and Ariel Rubinstein. 1994. *A Course in Game Theory*. MIT Press, Cambridge, MA, USA.
- [16] Amir Pnueli. 1977. The Temporal Logic of Programs. In *FOCS*. IEEE, Rhode Island, USA, 46–57.
- [17] Lillian J. Ratliff, Samuel A. Burden, and S. Shankar Sastry. 2016. On the Characterization of Local Nash Equilibria in Continuous Games. *IEEE Trans. Automat. Control* 61, 8 (2016), 2301–2307.
- [18] Yoav Shoham and Kevin Leyton-Brown. 2008. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, Cambridge, UK.
- [19] Igor Walukiewicz. 2004. A Landscape with Games in the Background. In *LICS*. IEEE, Turku, Finland, 356–366.
- [20] Michael Wooldridge, Julian Gutierrez, Paul Harrenstein, Enrico Marchioni, Giuseppe Perelli, and Alexis Toumi. 2016. Rational Verification: From Model Checking to Equilibrium Checking. In *AAAI*. AAAI Press, Phoenix, USA, 4184–4191.
- [21] Yichao Zhang, M. A. Aziz-Alaoui, Cyrille Bertelle, and Jihong Guan. 2014. Local Nash Equilibrium in Social Networks. *Scientific Reports – Nature* 4, 6224 (2014), 1–8.