

Original citation:

Harrenstein, Paul, Turrini, Paolo and Wooldridge, Michael (2015) *Electric Boolean games : redistribution schemes for resource-bounded agents*. In: AAMAS '15 Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, Istanbul, Turkey, 4-8 May 2015. Published in: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems pp. 655-663. ISBN 9781450334136.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/97675>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

"© ACM, 2015. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems pp. 655-663. ISBN 9781450334136 2015.

<https://dl.acm.org/citation.cfm?id=2772879.2773238>

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP url' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

Electric Boolean Games

Redistribution Schemes for Resource-Bounded Agents

Paul Harrenstein
Dept. of Computer Science
University of Oxford
paulhar@cs.ox.ac.uk

Paolo Turrini
Dept. of Computing
Imperial College London
p.turrini@imperial.ac.uk

Michael Wooldridge
Dept. of Computer Science
University of Oxford
mjw@cs.ox.ac.uk

ABSTRACT

In Boolean games, agents uniquely control a set of propositional variables, and aim at achieving a goal formula whose realisation might depend on the choices the other agents make with respect to the variables they control. We consider the case in which assigning a value to propositional variables incurs a cost, and moreover, we assume agents to be restricted in their choice of assignments by an initial endowment: they can only make choices with a lower cost than this endowment. We then consider the possibility that endowments can be redistributed among agents. Different redistributions may lead to Nash equilibrium outcomes with very different properties, and so certain redistributions may be considered more attractive than others. In this context we study *centralised* redistribution schemes, where a system designer is allowed to redistribute the initial energy endowment among the agents in order to achieve desirable systemic properties. We also show how to extend this basic model to a dynamic variant in which an electric Boolean game takes place over a series of rounds.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems;
J.4 [Computer Applications]: Social and Behavioral Sciences - Economics

General Terms

Economics; Theory

1. INTRODUCTION

Resource bounded agency is a core concern in artificial intelligence [7]. Every automated system directed to the realisation of some task must ultimately reckon with the usage of resources needed to accomplish it. In the multi-agent systems community special attention has been given to agents that are bounded-reasoners, in the sense that they are limited in processing higher-order beliefs [3], memory [2] or temporal horizon [11]. Somewhat less attention has been devoted to energy consumption affecting strategic ability, notable exceptions being [4, 23], which model how groups of agents can perform joint strategies depending on initial available resources. Also, in the computer-aided verification community, techniques have been developed that incorporate the idea of energy compliance, i.e., what a system can achieve under some limited resources [8]. These latter approaches focus however on general strategic ability, i.e., what a group of agents can achieve together, but do not address the issue of strategic *self-interested* behaviour. In brief, our aim in the present paper is to address this issue, studying systemic properties that can be achieved or avoided in equilibrium. Thus, our work is in a similar spirit to [24] and [16].

Our work is based on the framework of Boolean games [15, 5], a rich and natural model of interactive decision-making among goal-directed agents. In conventional Boolean games, each agent exercises unique control over a set of atomic propositions, in the sense that they can choose the values for these variables, and agents each seek the satisfaction of a goal, specified as formula of propositional logic. The model that we use in the present paper extends this basic setup by assuming that each choice has an associated cost (cf., [9, 16]). We intuitively understand the cost of a choice as being the energy requirement of the choice (hence *electric* Boolean game). Such an energy requirement affects the initial energy allocation each agent is endowed with, and actions can only be taken if they are consistent with the endowment.

The class of equilibrium outcomes that can be obtained in electric Boolean games is largely restricted by energy compliance issues. Therefore, we consider *centralised* redistribution schemes, in which an external authority or a system designer is allowed, before the game starts, to redistribute the initial energy endowment among the agents. The questions we ask concern the outcomes that can be rationally achieved or eliminated by making use of such schemes. It has to be noted that, due to the potentially infinite number of possible redistributions, the external authority is confronted with a computationally challenging task, and we therefore study useful algorithmic procedures to solve several related decision problems. We also show how to extend this basic model to *iterated electric Boolean games*, a dynamic variant of electric Boolean games based on *iterated Boolean games* as proposed by [13]. Iterated electric Boolean games take place over a series of rounds and agents, while bound by dynamically changing endowment constraints, iteratively choose an action and strive to achieve long term objectives expressed by LTL formulas.

The remainder of the paper is structured as follows. Section 2 introduces Boolean games, which we extend with a cost function and an energy endowment for each agent, giving what we call an *electric Boolean game*. In Section 3 we introduce the notion of an *endowment redistribution* as a centralised scheme to manipulate energy endowments. We show various algorithmic properties of decision problems that an external authority faces in one-shot games. In Section 4 we show how to extend the basic model to *iterated electric Boolean games*, the dynamic variant of electric Boolean games. In Section 5 we hint at possible directions for future work.

2. ELECTRIC BOOLEAN GAMES

A Boolean game is populated by a finite set of agents, each of which is given control over a finite set of propositional variables and assigned a formula of propositional logic: his *goal*. The game is played by agents independently and concurrently assigning a truth value to each propositional variable they control. The resulting set

of choices for each agent will define a valuation for the overall set of propositional variables, which will either satisfy or fail to satisfy each agent's goal formula. Clearly, agents will want to make choices that result in their goal being satisfied, but whether or not an agent's goal is satisfied will depend on the choices made by other agents.

Formally, a **Boolean game** [15, 14, 6] is defined on a finite set Φ of propositional variables as a tuple

$$B = (N, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n).$$

Here, $N = \{1, \dots, n\}$ is a finite set of agents, with typical element i and Φ_1, \dots, Φ_n is a partition of the propositional variables Φ over N , that is $\Phi = \bigcup_{i \in N} \Phi_i$, and $\Phi_i \cap \Phi_j = \emptyset$ for distinct agents i and j . For each agent i the set Φ_i collects the propositions under her unique control. Moreover, each γ_i is a propositional formula over the set Φ representing the **goal** of agent i .

A Boolean game is played by all agents choosing a truth-value assignment to the propositional variables they control. A **choice** for an agent i is a function $v_i: \Phi_i \rightarrow \{\perp, \top\}$. By V_i we denote the **set of choices** available to agent i . A **choice profile** or **outcome** is a tuple $\vec{v} = (v_1, \dots, v_n)$ of choices, one for each agent. By $\vec{V} = V_1 \times \dots \times V_n$ we refer to the set of all choice profiles or outcomes. Each outcome straightforwardly determines a valuation over all propositional variables. For $\vec{v} = (v_1, \dots, v_n)$ being a choice profile and $p \in \Phi_i$ for some (unique) $i \in N$, $\vec{v}(p)$ is the value assigned by v_i to atom p . We let \vec{v}_{-i} abbreviate $(v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$, a collection of choices of all agents but i . We will use $\vec{v}_{p\bar{q}\bar{r}}$ to denote the outcome in which variable p is set to true and variables q and r to false, and similarly for other outcomes.

Each choice function and, in particular, each choice profile, satisfies a set of formulas of propositional logic. We use \models to denote the satisfaction relation. For all valuations \vec{v} and \vec{v}' and each agent i with goal γ_i , we say that i **weakly prefers** choice profile v to choice profile v' , whenever

$$\vec{v}' \models \gamma_i \text{ implies } \vec{v} \models \gamma_i.$$

Thus, the agents' preferences are dichotomous, in the sense that they prefer outcomes satisfying their goals to outcomes that do not, otherwise they are indifferent.

Electric Boolean Games.

In a Boolean game agents can set variables to true or false as they see fit in pursuit of their goal. In many realistic settings, though, the execution of programs operating on value assignments comes with a cost, in terms of time or energy. Moreover, some actions may have such a high cost that they cannot be executed given the resources at an agent's disposal. Therefore, we consider Boolean games in which setting a propositional variable p to true or to false has a cost and in which each agent has only limited resources. Formally, a **cost function** c has the signature $c: \Phi \times \{\top, \perp\} \rightarrow \mathbb{Z}$. Intuitively, $c(p, b)$ is the cost that comes with setting propositional variable p to the Boolean value b . By c^0 we denote the **zero-cost function**, which assigns zero cost to each action. In Section 3, we will make the technical assumption that costs are non-negative, to avoid notational complications in the proofs. In a Boolean game, propositional variables are controlled by agents, and the cost of setting variable p to the value b falls to the agent i controlling p . Accordingly, with each choice v_i of agent i in a Boolean game, we associate the **aggregate cost** $c_i(v_i)$ resulting from this choice, i.e.,

$$c_i(v_i) = \sum_{p \in \Phi_i} c_i(p, v_i(p)).$$

The **total cost** of a choice profile $\vec{v} = (v_1, \dots, v_n)$ we define as,

$$c(\vec{v}) = \sum_{i \in N} c_i(v_i).$$

For an agent i to be able to make a choice v_i , she has to have sufficient resources at her disposal. Informally, we think of each agent having a battery and that she can only perform a particular action if that action does not consume more battery power than is available in the agent's battery.

The resources or battery power available to each agent are given by an **endowment function** e which associates with each agent i a value e_i in \mathbb{R}_0^+ . We denote the **total (aggregate) endowment** of e by $\mathbf{e} = \sum_{i \in N} e_i$. Further, let e^0 denote the **zero-endowment function**, which allots an endowment of 0 to each agent. We will assume that $\mathbf{e} \in \mathbb{N}_0$. The integer assumptions with respect to the cost functions and the total endowment are to avoid unnecessary representational issues and are thus purely for technical convenience. Recapping, while costs are modelled as integers, the endowments to agents can take real values. This, notice, means that there are infinitely many ways of allocating an initial amount of battery power among the agents.

The endowment that each agent is assigned indicates the level of energy they can consume. The idea is that agents can only take an action if the amount of energy consumed by that action does not exceed the level of battery power they have. Once an action is taken, the level of battery power of the agent taking it is updated accordingly. So, if Ann has an initial endowment of 10 units of energy then she can perform all actions that cost her no more than 10 units, but she will not be able to perform any action with a cost greater than 10, unless some additional resource is redistributed to her. If actions have negative cost, they should be understood as *recharging actions*, i.e., actions that have the effect of increasing the overall battery power by the agent taking them.

Formally, we define an **electric Boolean game** as a tuple (B, c, e) , consisting of a Boolean game $B = (N, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma)$, a cost function c , and an endowment function e . We usually denote the Boolean game (B, c, e) by $B^{c,e}$. The electric Boolean game consisting of Boolean game B , the zero-cost function c^0 and the zero-endowment function e^0 , we also denote by B^0 . As we will see later, in B^0 no restrictions are imposed on which choices the agents can make. It is important to note that the costs and endowments in an electric Boolean game *do not* affect the agents' preferences over the outcomes—these are the same as in the underlying Boolean game—but only which choices the agents can make.

Because of insufficient battery power, some choices might turn out to be prohibitively expensive for an agent. Thus, given an electric Boolean game $B^{c,e}$, we say that a choice v_i of agent i is **feasible for i** if $c_i(v_i) \leq e_i$. The set of feasible choices available to agent i in electric Boolean game $B^{c,e}$ we denote by $feasible_i(B^{c,e})$, so $feasible_i(B^{c,e}) \subseteq V_i$. Observe that in a Boolean game B^0 with the zero-cost function c^0 and zero-endowment function e^0 every choice for every agent is feasible. We say a choice profile $\vec{v} = (v_1, \dots, v_n)$ is feasible if for each agent i the choice v_i is feasible for i . We have $feasible(B^{c,e})$ denote the set of feasible choice profiles in $B^{c,e}$, which we also refer to as the **feasible region** of $B^{c,e}$. Thus,

$$feasible(B^{c,e}) = feasible_1(B^{c,e}) \times \dots \times feasible_n(B^{c,e}).$$

Feasible Nash equilibria.

In electric Boolean games, we assume that agents can only make feasible choices. Electric Boolean games restricted to their feasible regions can be analysed as strategic games, to which standard game-theoretic solution concepts may be applied [20]. When con-

		r			
		q	\bar{q}		
p	0, 0, 0	0, 1, 1			
\bar{p}	1, 1, 0	1, 0, 1			
	0	2			

		\bar{r}			
		q	\bar{q}		
	1, 0, 1	1, 1, 0	0		
	0, 1, 1	1, 1, 1	2		
	0	2			

$e_1 = e_2 = e_3 = 1$

Figure 1: Feasibility, equilibria, and redistribution

Considering individual deviations from a feasible outcome, we only consider deviations to *feasible choices*. So we define a **(pure) Nash equilibrium** of an electric Boolean game $B^{c,e}$ as a *feasible* outcome $\vec{v} = (v_1, \dots, v_n)$ such that for all agents i and all *feasible* choices $v'_i \in \text{feasible}_i(B^{c,e})$,

$$(\vec{v}_{-i}, v'_i) \models \gamma_i \text{ implies } \vec{v} \models \gamma_i.$$

The set of Nash equilibria of an electric Boolean game $B^{c,e}$ we denote by $NE(B^{c,e})$, and like in regular Boolean games, may be empty.

There are electric Boolean games with equilibria that might have undesirable social properties, e.g., inefficiency, and could lead to more desirable ones if endowments could be *redistributed* among agents. Consider, e.g., the three-agent electric Boolean game in Figure 1 with a total endowment $e = 3$ and three propositional variables p , q , and r . There, the row agent controls variable p , the column agent variable q , and the matrix agent variable r . The entries associated with agents' individual actions represent their cost, i.e., the row agent has cost 2 for setting p to false, while the entries in the cells encode whether agents satisfy their goal or not at that outcome. For instance, if p and q are set to true, but r to false then, the first (row) and the third (matrix) agents realise their goal, but the second (column) agent does not. This is represented by 1, 0, 1 in the respective cell. For endowment function e such that $e_1 = e_2 = e_3 = 1$, only outcome \vec{v}_{pqr} is feasible, whereas for redistribution e' with $e'_1 = 3$ and $e'_2 = e'_3 = 0$ both \vec{v}_{pqr} and $\vec{v}_{\bar{p}q\bar{r}}$ are. In fact, there are only four outcomes that are feasible under some redistribution, namely \vec{v}_{pqr} , $\vec{v}_{\bar{p}q\bar{r}}$, $\vec{v}_{p\bar{q}\bar{r}}$, and $\vec{v}_{\bar{p}q\bar{r}}$; all other outcomes will never be feasible. Under starting endowment e , outcome \vec{v}_{pqr} is the only Nash equilibrium, with no agent realising its goal. But now suppose an external authority redistributes e among the agents according to e' . This not only alters the set of feasible outcomes but also the set of equilibria. Under e' , outcome $\vec{v}_{\bar{p}q\bar{r}}$ would be the only Nash equilibrium, where two agents are realising their goal. Again, no possible redistribution can turn $\vec{v}_{\bar{p}q\bar{r}}$, the outcome maximising social welfare, into a Nash equilibrium.

3. REDISTRIBUTION SCHEMES

Different endowment distributions may significantly affect the reachability of potentially desirable properties in equilibrium. A key concept in this study of electric Boolean games is therefore that of *redistribution of endowments or battery power*. Intuitively, we give the external authority the power of an upfront reallocation of the initial battery power so as to steer the game towards more desirable outcomes. Formally, we say that an endowment function e' is a **redistribution** of an endowment function e , whenever e' assigns

		q	\bar{q}		
p	1, 1	1, 0	2		
\bar{p}	0, 1	0, 0	0		
	2	0			

$e_1 = e_2 = 1$

Figure 2: Feasibility and equilibria: no commutativity

the same total endowment to the agents as e does, i.e., if

$$\sum_{i \in N} e_i = \sum_{i \in N} e'_i.$$

When attempting to manipulate a game using redistributions it is important to notice that the set of Nash equilibria of an electric Boolean game is not simply given by the intersection of the feasible outcomes and the Nash equilibria of the underlying Boolean game. Although every Nash equilibrium of a Boolean game remains an equilibrium in the electric Boolean game no matter what costs on the actions are imposed or what endowment distribution is chosen (provided they are not rendered infeasible), the opposite inclusion does not hold.

PROPOSITION 1. *For every electric Boolean game $B^{c,e}$,*

$$NE(B^0) \cap \text{feasible}(B^{c,e}) \subseteq NE(B^{c,e}) \cap \text{feasible}(B^0).$$

The inclusion in the opposite direction, however, does not hold.

Proof: Consider an outcome $\vec{v} \in NE(B^0) \cap \text{feasible}(B^{c,e})$. As obviously $\vec{v} \in \text{feasible}(B^0)$, assume for contradiction that $\vec{v} \notin NE(B^{c,e})$. Then, there is an agent i and some feasible choice $v'_i \in \text{feasible}_i(B^{c,e})$ with $\vec{v} \not\models \gamma_i$ and $(\vec{v}_{-i}, v'_i) \models \gamma_i$. As trivially $v'_i \in \text{feasible}_i(B^0)$, it follows that $\vec{v} \notin NE(B^0)$, a contradiction. For the second part, consider the electric Boolean game depicted in Figure 2. There, $\vec{v}_{\bar{p}q} = NE(B^{c,e}) \cap \text{feasible}(B^0)$, whereas $NE(B^0) \cap \text{feasible}(B^{c,e}) = \emptyset$. \square

This observation could be of significance for a mediator who has the authority to (re-)distribute endowments over the agents. For instance, she may want to distribute resources over the agents in such a way so as to render particular outcomes feasible and other outcomes infeasible. In this way she can indirectly steer the agents' incentives and force some (desirable) outcomes to become Nash equilibria or other undesirable ones to cease to be in equilibrium. In this section, we therefore study the respective formal conditions under which, given an electric Boolean game $B^{c,e}$ and outcome \vec{v} ,

- (i) a redistribution e' of e exists such that $\vec{v} \in NE(B^{c,e'})$, and
- (ii) a redistribution e' of e exists such that $\vec{v} \notin NE(B^{c,e'})$.

The power of the central authority, however, may be bounded in the sense that she cannot render every outcome feasible or infeasible and *a fortiori* cannot manipulate the Nash equilibria entirely at will. In this section we restrict our attention to games with non-negative cost functions.

Taming the Number of Redistributions.

Despite their apparent simplicity, endowment redistributions in electric Boolean games introduce great complexity. Clearly, each redistribution yields its own feasible region, but not every feasible region can only be achieved by one redistribution. Still, the number of feasible regions that can be achieved by some redistribution in an electric Boolean game can be large. To give an indication, the number of *integer* redistributions of an *integer* endowment \mathbf{e} over n agents is given by

$$\binom{\mathbf{e} + n - 1}{n - 1}.$$

This is the number of ways one can partition a sum of integers into ordered non-negative integer summands. Many redistributions, however, yield the same set of feasible outcomes. Below, we show that there is a (representation) bound on the redistributions that we need to consider so as to be able to distinguish all feasible regions of an electric Boolean game.

Recall we are assuming that cost functions as well as the total endowment are both non-negative and integer-valued. As a consequence of this fact, we find that in electric Boolean games with n agents we can restrict attention to redistributions e for which the endowment e_i of each agent can be written as a fraction with n as denominator. The crucial point used in the proof is that, even if there is an infinite number of redistributions, we can, *without loss of generality*, restrict attention to a special subset of redistributions which can be represented polynomially but still can distinguish all feasible regions.

We say that choice v_i is **adjacent** to v'_i for an agent i if

- (i) $c_i(v_i) < c_i(v'_i)$, and
- (ii) there is no choice v''_i such that $c_i(v_i) < c_i(v''_i) < c_i(v'_i)$.

(One can think of the different choices of an agent being ordered by ascending cost.) We then obtain the following fact.

LEMMA 2. *Let $B^{c,e}$ be an electric Boolean game with n agents and a non-negative and integer-valued cost function c and endowment function e such that $\mathbf{e} \in \mathbb{N}_0$. Then, there is a redistribution e' of e such that,*

- (i) $e'_i \cdot n \in \mathbb{N}_0$, for each agent i , and
- (ii) $\text{feasible}(B^{c,e}) = \text{feasible}(B^{c,e'})$.

Sketch of proof: Consider a most expensive feasible outcome \bar{v}^* in $B^{c,e}$, i.e., $\bar{v}^* \in \arg \max_{\bar{v} \in \text{feasible}(B^{c,e})} \mathbf{c}(\bar{v})$. Then, $\mathbf{c}(\bar{v}^*) \leq \mathbf{e}$. If for some agent i there is no choice adjacent to v_i^* , then v_i^* is a most expensive choice for i among all his choices in V_i . Then, let, for each agent j ,

$$e'_j = \begin{cases} \mathbf{e} - \sum_{k \neq j} c_k(v_k^*) & \text{if } j = i, \\ c_j(v_j^*) & \text{otherwise.} \end{cases}$$

It is then easy to see that e' satisfies both (i) and (ii).

So for the remainder of the proof, we may assume that for every agent i there is a choice v'_i adjacent to v_i^* and let $\bar{v}' = (v'_1, \dots, v'_n)$. As $\mathbf{c}(\bar{v}') \in \mathbb{N}_0$, we know that for some integer $k \in \mathbb{N}_0$,

$$\mathbf{e} + k = \mathbf{c}(\bar{v}').$$

If $k = 1$, define e' such that $e'_i = c_i(v'_i) - \frac{1}{n}$ for each agent i . Thus, e' is a redistribution of e . Moreover, some reflection reveals that e' satisfies both (i) and (ii). The case for $k > 1$ is slightly more complicated than space allows, but runs along similar lines. \square

As an immediate consequence of Lemma 2, we find that the total number of feasible regions in an electric Boolean game with n agents and total endowment \mathbf{e} is upper bounded by

$$\binom{n\mathbf{e} + n - 1}{n - 1}.$$

Elimination.

Some equilibria might have properties we do not find desirable and a central authority may want to deploy redistributions to eliminate them. Thus, we say that an outcome \bar{v} is **eliminable** in electric Boolean game $B^{c,e}$ if \bar{v} is feasible and there exists a redistribution e' such that $\bar{v} \notin NE(B^{c,e'})$. We also say that a feasible outcome \bar{v} is **feasibly eliminable** in $B^{c,e}$ if \bar{v} is feasible and there exists a redistribution e' such that $\bar{v} \in \text{feasible}(B^{c,e'}) \setminus NE(B^{c,e'})$. Finally, we say that an outcome \bar{v} is **rationally eliminable** if \bar{v} is an equilibrium in $B^{c,e}$ and there exists a redistribution e' such that $\bar{v} \in \text{feasible}(B^{c,e'}) \setminus NE(B^{c,e'})$.

Despite the large number of possible redistributions, the external authority may find herself confronted with, we find that for the elimination of outcomes or equilibria, she may restrict her attention to significantly smaller set of redistributions.

Define for each electric Boolean game $B^{c,e}$ and each agent i a redistribution e^i that assigns the total endowment \mathbf{e} to i and 0 to all other agents, i.e., all agents j ,

$$e^i_j = \begin{cases} \mathbf{e} & j = i, \\ 0 & \text{otherwise.} \end{cases}$$

Now we have the following proposition and corollary, which say that to establish whether an outcome is a Nash equilibrium under every redistribution one only needs to check $|N|$ redistributions.

PROPOSITION 3. *Let \bar{v} be an outcome of an electric Boolean game $B^{c,e}$. Then, $\bar{v} \in NE(B^{c,e'})$ for all redistributions e' of e if and only if $\bar{v} \in NE(B^{c,e^i})$ for all agents i .*

Proof: The “if”-direction is immediate. Observe that, for every agent i , e^i is a redistribution of e . For the opposite direction, assume $\bar{v} \notin NE(B^{c,e'})$ for some redistribution e' . Then, either

- (i) $\bar{v} \notin \text{feasible}(B^{c,e'})$, or
- (ii) $\bar{v} \not\models \gamma_i$ and $(\bar{v}_{-i}, v'_i) \models \gamma_i$ for some $v'_i \in \text{feasible}_i(B^{c,e'})$.

If (i) there is some agent i with $c_i(v_i) > e'_i$. If i is the only agent, immediately $\bar{v} \notin \text{feasible}(B^{c,e^i})$. Otherwise, consider the redistribution e^j for some agent j distinct from i . As $c_i(v_i) > e'_i \geq 0$, also $c_i(v_i) > e^j_i = 0$. Hence, $\bar{v} \notin \text{feasible}(B^{c,e^j})$ and $\bar{v} \notin NE(B^{c,e^j})$. If (ii), we also have that $v'_i \in \text{feasible}_i(B^{c,e^i})$ and it follows that $\bar{v} \notin NE(B^{c,e^i})$. \square

COROLLARY 4. *Let \bar{v} be an outcome of an electric Boolean game $B^{c,e}$. Then, \bar{v} is eliminable if and only if $\bar{v} \notin NE(B^{c,e^i})$ for some agent i .*

To see if, given an electric Boolean game, an outcome \bar{v} is a Nash equilibrium for every redistribution *under which \bar{v} is feasible*, a similar argument applies. In this case, however, define for each agent i the redistribution $e^{\bar{v},i}$ such that, for all agents j ,

$$e^{\bar{v},i}_j = \begin{cases} \max(\mathbf{e} - \sum_{k \neq j} c_k(v_k), 0) & \text{if } j = i, \\ c_j(v_j) & \text{otherwise.} \end{cases}$$

Roughly speaking, $e^{\vec{v},i}$ gives agent i the maximal endowment without rendering \vec{v} infeasible. A proof along analogous lines as the one for Proposition 3 then yields the following result. Recall that we have assumed cost functions to be non-negative. Under this assumption, for each agent i , an outcome \vec{v} is infeasible in $B^{c,e^{\vec{v},i}}$ if and only if \vec{v} is infeasible under every redistribution of e . Again we find that we only need to consider $|N|$ redistributions to establish whether an outcome is feasibly eliminable.

PROPOSITION 5. *Let \vec{v} be an outcome of an electric Boolean game $B^{c,e}$ with non-negative cost function c . Then, $\vec{v} \in NE(B^{c,e'})$ for all redistributions e' of e such that $\vec{v} \in \text{feasible}(B^{c,e'})$ if and only if $\vec{v} \in NE(B^{c,e^{\vec{v},i}})$ for all agents i .*

COROLLARY 6. *Let \vec{v} be an outcome of an electric Boolean game $B^{c,e}$. Then, \vec{v} is feasibly eliminable if and only if $\vec{v} \notin NE(B^{c,e^{\vec{v},i}})$ for some agent i .*

We define the following three decision problems.

ELIMINATION

Given: Electric Boolean game $B^{c,e}$ and feasible choice profile $\vec{v} \in \text{feasible}(B^{c,e})$

Problem: Is v eliminable in $B^{c,e}$?

FEASIBLE ELIMINATION

Given: Electric Boolean game $B^{c,e}$ and feasible choice profile $\vec{v} \in \text{feasible}(B^{c,e})$

Problem: Is v feasibly eliminable in $B^{c,e}$?

RATIONAL ELIMINATION

Given: Electric Boolean game $B^{c,e}$ and feasible choice profile $\vec{v} \in NE(B^{c,e})$

Problem: Is v rationally eliminable in $B^{c,e}$?

Surprisingly, checking whether an equilibrium is (feasibly) eliminable is no harder than checking that an outcome is not a Nash equilibrium:

PROPOSITION 7. *ELIMINATION and FEASIBLE ELIMINATION are both NP-complete.*

Proof: For both ELIMINATION and FEASIBLE ELIMINATION NP-hardness follows by reducing the problem of deciding whether a given profile \vec{v} is a Nash equilibrium in a regular Boolean game—which is known to be coNP-complete¹—to the complementary problem. Given a Boolean game $B = (N, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n)$ construct electric Boolean game B^0 with the zero-cost function c^0 and the zero-endowment function e^0 . Observe that in B^0 all outcomes are feasible and e^0 is the only redistribution of e^0 itself. It is then easy to see that \vec{v} is a pure Nash equilibrium in B if and only if \vec{v} is not (feasibly) eliminable in B^0 .

For membership in NP of either ELIMINATION or FEASIBLE ELIMINATION, let $B^{c,e}$ be a given electric Boolean game and \vec{v} a given outcome. We can guess a redistribution e' of e , an agent i , and a choice v'_i . Observe that, by virtue of Lemma 2, this can be achieved in polynomial time. For ELIMINATION, we can then verify whether either $c_j(v_j) > e'_j$ for some $j \in N$ or $c_i(v'_i) \leq e'_i$, $\vec{v} \not\models \gamma_i$ and $(\vec{v}_i, v'_i) \models \gamma_i$. For FEASIBLE ELIMINATION verify whether $c_i(v'_i) > e^{\vec{v},i}$ and whether both $\vec{v} \not\models \gamma_i$ and $(\vec{v}_i, v'_i) \models \gamma_i$. As all this can be performed in polynomial time, we are done. \square

¹See [6] for a complexity analysis of Boolean games.

For RATIONAL ELIMINATION we have a similar result.

PROPOSITION 8. *RATIONAL ELIMINATION is NP-complete.*

Proof: As in the proof of Proposition 7 hardness is proved by reducing problem of deciding whether a given profile \vec{v} is a Nash equilibrium in a regular Boolean game. Given a Boolean game $B = (N, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n)$ and an outcome $\vec{v} = (v_1, \dots, v_n)$, construct the electric Boolean game

$$D^{c,e} = (N \cup \{0\}, \Phi_0, \Phi_1, \dots, \Phi_n, \gamma_0, \dots, \gamma_n, c, e),$$

where 0 is a dummy agent with $\Phi_0 = \emptyset$ (thus, 0 has one choice v_0 , namely, $v_0 = \emptyset$) and $\gamma_0 = \top$. Set cost function c such that for every agents $i \in N$ and every choice v'_i ,

$$c_i(v'_i) = \begin{cases} 0 & \text{if } v'_i = v_i, \\ 1 & \text{otherwise.} \end{cases}$$

for every agent $i \in N$ and observe that $c_0(\emptyset) = \sum \emptyset = 0$ by common convention. Let the endowment function e be such that, for all agents $i \in N \cup \{0\}$ and all choices v'_i ,

$$e_i = \begin{cases} |N| & \text{if } i = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Then, \vec{v} is the only feasible outcome in $D^{c,e}$ and $\vec{v} \in NE(D^{c,e})$. Let, furthermore, e' be the endowment function such that

$$e'_i = \begin{cases} 0 & \text{if } i = 0, \\ 1 & \text{otherwise.} \end{cases}$$

Obviously, e^* is a redistribution of e and every outcome is feasible in D^{c,e^*} . This holds, in particular, for \vec{v} . Moreover, as can easily be appreciated, $\vec{v} \notin NE(D^{c,e'})$ for some redistribution e' of e if and only if $\vec{v} \notin NE(D^{c,e^*})$. Hence,

\vec{v} is a Nash equilibrium in B

iff $\vec{v} \in NE(D^{c,e^*})$

iff $\vec{v} \in NE(D^{c,e'})$ for all redistributions e' with $\vec{v} \in \text{feasible}(D^{c,e'})$
iff \vec{v} is not rationally eliminable.

Membership in NP can be established as in the proof of NP-membership of FEASIBLE ELIMINATION (see Proposition 7). \square

Construction.

Rather than eliminating outcomes and equilibria with presumably undesirable properties, a central authority might want to see if she can introduce some desirable equilibria by choosing an appropriate redistribution. Thus, we come to consider the **construction problem**, i.e., the problem of deciding whether an endowment redistribution exists under which a given outcome is a Nash equilibrium in a given electric Boolean game. To facilitate the discussion, we say that choice v'_i is a **beneficial deviation** from an outcome \vec{v} for agent i if both $\vec{v} \not\models \gamma_i$ and $(\vec{v}_i, v'_i) \models \gamma_i$.

For a given electric Boolean game $B^{c,e}$ and a given outcome \vec{v} , we distinguish four cases, in the first three of which the construction problem has an immediate solution.

Case 1 $e < c(\vec{v})$,

Case 2 $e \geq c(\vec{v})$ and some agent i has a beneficial deviation v'_i from \vec{v} such that $c_i(v'_i) \leq c_i(v_i)$.

Case 3 Neither Case 1 nor Case 2 applies, and some agent i has no beneficial deviation v'_i from \vec{v} .

Case 4 Neither Case 1, Case 2, nor Case 3 applies, i.e., $\mathbf{e} \geq \mathbf{c}(\vec{v})$ and all agents i have a beneficial deviation v'_i from \vec{v} . Moreover, for all agents i and for all beneficial deviations v'_i , we have $c_i(v'_i) > c_i(v_i)$.

If Case 1 obtains, there is no redistribution e' of e under which \vec{v} is feasible and, hence, there is no redistribution under which \vec{v} is a Nash equilibrium either. Neither in Case 2 can \vec{v} be made into a Nash equilibrium through redistributing the total endowment. Observe that in this case, if agent i has sufficient endowment to choose v_i , she has also sufficient endowment to choose v'_i . Hence, no matter which endowment assigned to her, agent i can always profitably deviate to v'_i , whenever v_i is feasible.

By contrast, if Case 3 applies, for every agent i without a beneficial deviation from \vec{v} , outcome \vec{v} will be a Nash equilibrium under redistribution $e^{\vec{v},i}$. Recall that $e^{\vec{v},i}$ is the redistribution that assigns all agents j other than i their cost at \vec{v} , with agent i getting whatever remains. As in this case $\mathbf{c}(\vec{v}) \leq \mathbf{e}$, outcome \vec{v} is feasible in $B^{c,e^{\vec{v},i}}$. Moreover, as Case 2 does not apply, in $B^{c,e^{\vec{v},i}}$ no agent distinct from i can profitably deviate to a feasible choice and i cannot profitably deviate at all. Observe that Case 3 already applies if there is one agent who has her goal satisfied at \vec{v} .

This leaves us with Case 4. Now we may assume that all agents i have a beneficial deviation from \vec{v} to some choice v_i . We find that in this case the constructibility of a Nash equilibrium \vec{v} in an electric Boolean game $B^{c,e}$ can be reduced to a making a comparison of the total endowment $\mathbf{e}^{\vec{v}}$ of one particular endowment function $e^{\vec{v}}$ with the total endowment \mathbf{e} of e . Before we give this characterisation, we first introduce some more terminology and notation.

Define, for each agent i , the set $V_i^{\vec{v}}$ of beneficial deviations v'_i from \vec{v} that agent i has. Under the assumptions of Case 4, clearly, $V_i^{\vec{v}} \neq \emptyset$ for all agents i . Moreover, we may also assume that $c_i(v_i) < c_i(v'_i)$ for all $v'_i \in V_i^{\vec{v}}$. Now, define the endowment function $e^{\vec{v}}$, such that, for every agent i ,

$$e_i^{\vec{v}} = \min\{c_i(v'_i) : v'_i \in V_i^{\vec{v}}\}.$$

Thus, $e^{\vec{v}}$ assigns to every agent the minimum cost for a beneficial deviation from \vec{v} . Observe that $e^{\vec{v}}$ need not be a redistribution of e . We now obtain the following characterisation.

PROPOSITION 9. *Let \vec{v} be an outcome in an electric Boolean game $B^{c,e}$ with non-negative costs and for which Case 4 holds. Then, there is some redistribution e' of e such that $\vec{v} \in NE(B^{c,e'})$ if and only if*

$$\mathbf{c}(\vec{v}) \leq \mathbf{e} < \mathbf{e}^{\vec{v}}.$$

Proof: For the “only if”-direction, first assume for contraposition that $\mathbf{c}(\vec{v}) > \mathbf{e}$ and consider an arbitrary redistribution e' of e . Then, $\mathbf{e}' = \mathbf{e}$ and therefore $\vec{v} \notin \text{feasible}(B^{c,e'})$. It follows that there is no redistribution e' of e with $\vec{v} \in NE(B^{c,e'})$. Now assume for contraposition that $\mathbf{e} \geq \mathbf{e}^{\vec{v}}$ and again consider an arbitrary redistribution e' of e . Observe that there is an agent i with $e'_i \geq e_i^{\vec{v}}$. If $\vec{v} \notin \text{feasible}(B^{c,e'})$, then also $\vec{v} \notin NE(B^{c,e'})$. If, on the other hand, $\vec{v} \in \text{feasible}(B^{c,e'})$, then agent i can profitably deviate from \vec{v} to some $v'_i \in V_i^{\vec{v}}$. Thus, also in this case, $\vec{v} \notin NE(B^{c,e'})$.

For the “if”-direction, assume that $\mathbf{c}(\vec{v}) \leq \mathbf{e} < \mathbf{e}^{\vec{v}}$. Then, there is some redistribution e' of e with $c_i(v_i) \leq e'_i < e_i^{\vec{v}}$ for all agents i . Accordingly, $\vec{v} \in \text{feasible}(B^{c,e'})$. Moreover, for every agent i , every beneficial deviation v'_i from \vec{v} by i is infeasible in $B^{c,e'}$, that is, $e'_i < c_i(v'_i)$. It follows that $\vec{v} \in NE(B^{c,e'})$, as desired. \square

Our analysis of the construction problem can be leveraged to obtain a complexity result for the corresponding decision problem.

RATIONAL CONSTRUCTION

Given: Electric Boolean game $B^{c,e}$ and outcome \vec{v}

Problem: Is there a redistribution e' of e with $\vec{v} \in NE(B^{c,e'})$?

We find that **RATIONAL CONSTRUCTION** is coNP-hard, but cannot be computationally more complex than P^{NP} .²

PROPOSITION 10. **RATIONAL CONSTRUCTION** coNP-hard and included in P^{NP} .

Proof: We prove coNP-hardness by reducing the problem of deciding whether a given profile \vec{v} is a Nash equilibrium in a regular Boolean game. Given a Boolean game B , construct electric Boolean game B^0 with the zero-cost function c^0 and zero-endowment function e^0 . Then, all outcomes are feasible and there is only one redistribution of e^0 , namely, e^0 itself. Hence, trivially, \vec{v} is a pure Nash equilibrium in B if and only if \vec{v} is a Nash equilibrium in B^{c^0,e^0} for some redistribution e' of e^0 .

To see that **RATIONAL CONSTRUCTION** is in P^{NP} , we first show that, for a given n -agent electric Boolean game $B^{c,e}$ with endowment function e and outcome \vec{v} , each of the following can be achieved in coNP.

- D1 Check whether $\mathbf{e} \geq \mathbf{c}(\vec{v})$, i.e., check whether Case 1 applies.
- D2 Decide whether, for some agent i no deviation v'_i from \vec{v} is beneficial, i.e., check whether Case 3 applies.
- D3 Decide whether for all agents i and all beneficial deviations v'_i from \vec{v} of i it is the case that $c_i(v_i) < c_i(v'_i)$, i.e., check whether Case 4 applies.
- D4 Decide, for a given agent i and a given integer $x \in \mathbb{N}_0$, whether $e_i^{\vec{v}} = x$.
- D5 Given an endowment e' , check whether $\mathbf{e} < \mathbf{e}'$.

Clearly, both D1 and D5 can be achieved in polynomial time. For D2, we can guess for each agent i a choice v'_i , and check whether for all of these n guesses of v'_i both $\vec{v} \not\models \gamma_i$ and $(\vec{v}_{-i}, v'_i) \models \gamma_i$. Also this can be performed in polynomial time. For D3, we can guess an agent i along with a choice $v'_i \in V_i$ and check whether $\vec{v} \not\models \gamma_i$, $(\vec{v}_{-i}, v'_i) \models \gamma_i$, and $c_i(v'_i) \leq c_i(v_i)$. Again, this can be done in polynomial time. Finally, for D4, being given agent i and integer x , guess a choice $v'_i \in V_i$ and check whether both $\vec{v} \not\models \gamma_i$ and $(\vec{v}_{-i}, v'_i) \models \gamma_i$ as well as whether $c_i(v'_i) < x$.

Given answers for D1 through D5, we can proceed as follows. If D1 yields “no”, return “no”, i.e., Case 1 applies and there is no redistribution making \vec{v} into a Nash equilibrium. If D1 yields “no” and D2 “yes”, Case 3 applies and we can return “yes”. Finally, if D1, D2, and D3 all yield “no”, Case 2 applies and return “no”. Otherwise, Case 4 applies. Then, for each agent i , define e'_i such that $e'_i = x$ if and only if the oracle for D4 yields “yes” on i and x , by going through all integers $x \in \{0, \dots, c_i(v_i^*)\}$, where v_i^* is agent i 's costliest choice. Thus, $e^{\vec{v}} = (e'_1, \dots, e'_n)$. Then, we query the oracle for D5 and output “yes” if and only if $\mathbf{e} < \mathbf{e}'$. The soundness of this last step follows from Proposition 9. We conclude that **CONSTRUCTION** is in P^{NP} . \square

²Recall that P^{NP} is the class of languages recognised by a deterministic Turing machine in polynomial time with a polynomial number of queries of an NP-oracle.




Figure 3: An iterated electric Boolean game.

What we have shown is that, despite the prohibitive number of redistributions, an external authority does not have to check all of them to establish whether she can find one that eliminates an undesirable or creates a desirable equilibrium. By doing so, however, also other equilibria may be affected. An interesting issue for future research would therefore be to investigate the conditions and complexity of the problem to find a redistribution so as to eliminate all equilibria with a particular property φ without introducing any that do not satisfy φ .

4. THE DYNAMIC SETTING

In the previous section, agents were involved one-shot interactions, trying to achieve propositional goals by performing feasible choice profiles. In this section, we explore a dynamic setting, where agents repeatedly choose an assignment to the propositional variables they control. However, the cost an agent incurs with the choices she makes at each point of time has direct effect on her level of battery power, and thus on the choices she can make, in the next.

Consider the game in Figure 3, in which, over a prolonged period of time, two agents, *Row* and *Col*, have to make choices for p and q , respectively. As *Row* has an initial endowment of only 1, he can set p to true only once before his endowment sinks to 0. Then, he has to set p to false, an action with negative costs that ‘recharges’ his endowment by 3. With an initial endowment of 2, *Col* has to start setting q to false at least twice. Thus, under the initial endowment function e , it is inevitable that p and q will both be simultaneously set to false before the second round. If, however, an external authority wants to make sure that $p \vee q$ always holds, she could redistribute e and assign an initial endowment of 2 to *Row* and one of 0 to *Col*. Then, the sequence $(\vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}})^\omega = \vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}}, \dots$ would become feasible under the resource constraints imposed by this redistribution. Clearly, under any redistribution of the total endowment e , at any given time in any run *Col* will have set q to false at least twice as often as he has set q to true. Observe, however, that this would no longer be the case if, rather than merely redistributing the *initial* endowment, the external authority could redistribute the available total endowment available at any point of time. By then invariably assigning the total available endowment at each time to *Col* also the sequence $(\vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}})^\omega = \vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}}, \vec{v}_{p\bar{q}}, \dots$ would be feasible.

To formally reason about such settings, we extend the model of *iterated Boolean games* as proposed by Gutierrez et al. [13] by adding cost functions and endowment functions.

Linear-Time Temporal Logic (LTL).

When agents make choices iteratively, this leads to (infinite) sequences of valuations, which properties can be described using **linear-time temporal logic** (LTL) [10, 17, 18]. Besides the usual Boolean connectives, LTL also has temporal operators **X** (“next”),

F (“eventually”), **G** (“always”), and the binary modal operator **U** (“until”). Formally, the syntax of LTL is defined with respect to a finite set Φ of propositional variables as follows, where $p \in \Phi$.

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi$$

The remaining classical and LTL connectives are then defined in the standard way. In particular, $\mathbf{F}\varphi = \top \mathbf{U}\varphi$, and $\mathbf{G}\varphi = \neg \mathbf{F}\neg\varphi$.

Formulas of LTL are interpreted on runs and times. Formally, a **run** ρ is an infinite sequence of valuations in \vec{V}^ω . We use $t \in \mathbb{N}_0$ as a temporal index into ρ and write $\vec{v}[t] = (v_1[t], \dots, v_n[t])$ for the valuation at time t in ρ . Then, for $p \in \Phi$, run $\rho = \vec{v}[0], \vec{v}[1], \dots$, and $t \geq 0$, the semantics for LTL is as follows.

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

We write $\rho \models \varphi$ for $\rho, 0 \models \varphi$ and say that φ **satisfied** by ρ .

Iterated Electric Boolean Games.

An **iterated Boolean Game** $\mathbf{B} = (N, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n)$ is exactly like a Boolean game, except that each agent i (longterm) goal is given by an LTL formula γ_i [13]. An **iterated electric Boolean game** we define a tuple (\mathbf{B}, c, e) , where \mathbf{B} an iterated Boolean game and c and e a **cost function** and an **endowment function** as before. Again, we write $\mathbf{B}^{c,e}$ for (\mathbf{B}, c, e) . We assume that for each agent i there is a choice c_i with cost at most zero.

An iterated electric Boolean game is played in an infinite series of rounds, in each of which each agent i assigns values to the variables in Φ_i she controls giving rise to an infinite sequence of valuations. Each agent i tries to set her variables in each round in such a way so as to eventually get her goal γ_i achieved. Here, an agent can condition her choice on the choices of the other agents in previous rounds. Thus, a **strategy** for agent i is a function $\sigma_i : \vec{V}^* \rightarrow V_i$, which associates with each **history**—i.e., a finite and possibly empty sequence $\varrho = \vec{v}[0], \dots, \vec{v}[t]$ of valuations—a choice $\sigma_i(\varrho)$ in V_i . For $t < 0$, we stipulate that $\vec{v}[0], \dots, \vec{v}[t]$ is the empty sequence ϵ . A profile $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$ of strategies then induces a run $\rho(\vec{\sigma}) = \vec{v}[0], \vec{v}[1], \vec{v}[2], \dots$ such that, for $t \geq 0$,

$$\begin{aligned} \vec{v}[0] &= (\sigma_1(\epsilon), \dots, \sigma_n(\epsilon)), \\ \vec{v}[t+1] &= (\sigma_1(\vec{v}[0], \dots, \vec{v}[t]), \dots, \sigma_n(\vec{v}[0], \dots, \vec{v}[t])). \end{aligned}$$

The choices an agent makes at time t may affect her endowment at $t+1$: it may go down or up depending on whether her choice at t brings with it a positive or a negative cost. Formally, we define an **endowment scheme** as a function ε that associates with each agent i and each history ϱ in \vec{V}^* a value $\varepsilon_i(\varrho)$ in \mathbb{R} . Thus, the initial endowment function e of an iterated electric Boolean game $\mathbf{B}^{c,e}$ can naturally be extended to an endowment scheme, which we will here also denote by e , such that for each agent i and time $t \geq 0$,

$$\begin{aligned} e_i(\epsilon) &= e_i, \\ e_i(\vec{v}[0], \dots, \vec{v}[t]) &= e_i(\vec{v}[0], \dots, \vec{v}[t-1]) - c_i(v_i[t]). \end{aligned}$$

The **total endowment** may also vary over time and depend on a history; for $t \geq 0$, its formal definition is given by,

$$\begin{aligned} \mathbf{e}(\epsilon) &= \mathbf{e}, \\ \mathbf{e}(\vec{v}[0], \dots, \vec{v}[t]) &= \mathbf{e} - \sum_{0 \leq t' < t} \mathbf{c}(\vec{v}[t']). \end{aligned}$$




Figure 4: An iterated electric Boolean game whose feasible runs under dynamic redistribution are precisely those in which, at every point of time where p is set to true, q was set to true at least once more often than p in the past.

At any point of time, agents can only choose actions that their current endowment allows them to play. Formally, $\vec{v}[0], \vec{v}[1], \dots$ is a **feasible run** of an iterated electric Boolean game $\mathbf{B}^{c,e}$, if for every agent i and every $t \geq 0$,

$$\begin{aligned} c_i(v_i[0]) &\leq e_i(\epsilon), \\ c_i(v_i[t+1]) &\leq e_i(\vec{v}[0], \dots, \vec{v}[t]). \end{aligned}$$

Notice that at a feasible run no agent’s endowment will be negative at any point of time. The set of feasible runs of an iterated electric Boolean game $\mathbf{B}^{c,e}$ we denote by $feasible(\mathbf{B}^{c,e})$. Having assumed that all agents have an action with cost zero, the set of feasible runs in an iterated electric Boolean game is always non-empty.

Dynamic Redistribution Schemes.

When playing an iterated electric Boolean game, the individual endowments to the agents may fluctuate. Thus, it may occur that at some point one agent’s endowment is low and cannot play a particular action, whereas another agent has surplus of endowment. If this happens, some runs will be prevented from being feasible. If such a run is desirable, however, an external authority might want to allocate some of the surplus endowment available at the time to the one agent to the other agent. Thus, the setting of iterated electric Boolean games we come to consider redistribution schemes in which a central planner can redistribute over time. Formally, we define a **dynamic redistribution scheme** for $\mathbf{B}^{c,e}$ as an endowment scheme ϵ such that, for every history $\vec{v}[0], \dots, \vec{v}[t]$, both

- (i) $e(\vec{v}[0], \dots, \vec{v}[t]) = \sum_{i \in N} \epsilon_i(v[0], \dots, \vec{v}[t])$, and
- (ii) $e(\vec{v}[0], \dots, \vec{v}[t]) \geq 0$ implies $\epsilon_i(v[0], \dots, \vec{v}[t]) \geq 0$ for all $i \in N$.

Condition (i) requires that the total endowment available at a certain point in time is fully redistributed over all agents. Condition (ii) ensures that every agent gets a non-negative endowment if the total available endowment is non-negative.

Feasibility of a run in an iterated electric Boolean game extends naturally to feasibility in an iterated electric Boolean game under an redistribution scheme ϵ . Run $\rho = \vec{v}[0], \vec{v}[1], \dots$ is **feasible under ϵ** if, for every agent i and $t \geq 0$,

$$\begin{aligned} c_i(v_i[0]) &\leq \epsilon_i(\epsilon), \\ c_i(v_i[t+1]) &\leq \epsilon_i(\vec{v}[0], \dots, \vec{v}[t]). \end{aligned}$$

The runs of an electric Boolean game $\mathbf{B}^{c,e}$ under dynamic redistribution scheme ϵ we denote by $feasible(\mathbf{B}^{c,\epsilon})$. The game in Figure 3 shows that the sets $feasible(\mathbf{B}^{c,e})$ and $feasible(\mathbf{B}^{c,\epsilon})$ may be quite different properties. It might also be worth observing that, in contrast to the sets of runs of an iterated Boolean game, the set feasible

runs of an iterated electric Boolean game, under a dynamic redistribution scheme, need not be an ω -regular language over \vec{V} , i.e., a language recognisable by a Büchi automaton. For an example, see Figure 4. The key insight behind this phenomenon is that the level of endowment behaves as the stack of a pushdown automaton.

An external authority may be interested whether a given iterated electric Boolean game allows for feasible runs that satisfy a certain LTL-expressible property, and if there are none, whether she can design a dynamic redistribution scheme so that runs satisfying the property become feasible. Thus, we come to consider the following decision problem.

RATIONAL CONSTRUCTION

Given: Iterated electric Boolean game $\mathbf{B}^{c,e}$ and LTL formula φ

Problem: Does a dynamic redistribution scheme ϵ for $\mathbf{B}^{c,e}$ exist such that $\rho \models \varphi$ for some $\rho \in feasible(\mathbf{B}^{c,\epsilon})$?

This problem is closely related to the satisfiability problem for LTL, which is known to be PSPACE-complete (see, Sistla and Clarke [22]). We obtain the following results by a non-trivial adaptation of Sistla and Clarke’s original proof. The proof itself is here omitted for reasons of space.

PROPOSITION 11. FEASIBLE DYNAMIC REDISTRIBUTION is PSPACE-complete.

5. CONCLUSION AND FUTURE WORK

We have looked at Boolean games where agents can only perform actions provided they meet given resource requirements. In this context, we considered the possibilities a central planner authorised to redistribute the agents’ endowments in order to steer the game to more desirable outcomes. In spite of the prohibitive number of possible redistributions, we found that well-behaved procedures exist to verify whether a given outcome is eliminable, feasibly eliminable or rationally eliminable. We have also been able to obtain similar results for procedures introducing new equilibria, i.e., for the construction problem. We have studied what can be achieved in the static, one-shot setting and given details of how things work out in the iterated setting.

The most natural direction of future research is to investigate the formal properties of *decentralised* redistribution schemes, in which groups of agents can get together and decide to redistribute resources among themselves. To handle this case, solution concepts of a more cooperative nature—e.g., stability notions based on feasible group deviations, or, even more generally, on *effectivity functions* (see, e.g., [19, 1])—are called for. As feasibility introduces some extra structure in Boolean games, an interesting issue would be to characterise the conditions under which an effectivity function *corresponds* to some electric Boolean game (cf., e.g., [21, 12]).

Finally, appropriate stability concepts could also be defined for and applied to the dynamic setting. Thus, one could investigate how the behaviour of decentralised redistribution over time and how it relates to centralised redistribution.

6. ACKNOWLEDGMENTS

Michael Wooldridge and Paul Harrenstein are supported by the ERC under Advanced Grant 291528 (“RACE”). Paolo Turrini acknowledges the support of the IEF Marie Curie fellowship “Norms in Action: Designing and Comparing Regulatory Mechanisms for Multi-Agent Systems” (FP7- PEOPLE-2012-IEF, 327424 “NINA”).

REFERENCES

- [1] J. Abdou and H. Keiding. *Effectivity Functions in Social Choice Theory*. Kluwer Academic Publishers: Dordrecht, The Netherlands, 1991.
- [2] A. Albore, N. Alechina, P. Bertoli, C. Ghidini, B. Logan, and L. Serafini. Model-checking memory requirements of resource bounded reasoners. In *Proceedings of the Twenty-first National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, 2006*, pages 213–218. ACM Press, 2006.
- [3] N. Alechina and B. Logan. Ascribing beliefs to resource bounded agents. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS 2002*, pages 881–888. ACM Press, 2002.
- [4] N. Alechina, B. Logan, N. H. Nga, and A. Rakib. Logic for coalitions with bounded resources. *Journal of Logic and Computation*, 21(6):907–937, 2011.
- [5] E. Bonzon, M.-C. Lagasquie-Schiex, and J. Lang. Dependencies between players in Boolean games. *International Journal of Approximate Reasoning*, 50(6):899–914, 2009.
- [6] E. Bonzon, M.-C. Lagasquie-Schiex, J. Lang, and B. Zanuttini. Compact preference representation and Boolean games. *Autonomous Agents and Multi-Agent Systems*, 18(1):1–35, 2009.
- [7] M. E. Bratman, D. J. Israel, and M. E. Pollack. Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4:349–355, 1988.
- [8] A. Chakrabarti, L. de Alfaro, T. A. Henzinger, and M. Stoelinga. Resource interfaces. In *Proceedings of the Third International Conference on Embedded Software (EMSOFT) (LNCS Volume 2855)*, pages 117–133. Springer-Verlag: Berlin, Germany, 2003.
- [9] P. E. Dunne, S. Kraus, W. van der Hoek, and M. Wooldridge. Cooperative Boolean games. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2008)*, Estoril, Portugal, 2008.
- [10] E. A. Emerson. Temporal and modal logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science Volume B: Formal Models and Semantics*, pages 996–1072. Elsevier Science Publishers B.V.: Amsterdam, The Netherlands, 1990.
- [11] M. Fisher and C. Ghidini. Executable specifications of resource-bounded agents. *Autonomous Agents and Multi-Agent Systems*, 21(3):368–396, 2010.
- [12] V. Goranko, W. Jamroga, and P. Turrini. Strategic games and truly playable effectivity functions. *Autonomous Agents and Multi-Agent Systems*, 26(2):288–314, 2013.
- [13] J. Gutierrez, P. Harrenstein, and M. Wooldridge. Iterated Boolean games. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI-13)*, Beijing, China, 2013.
- [14] P. Harrenstein. *Logic in Conflict*. PhD thesis, Utrecht University, 2004.
- [15] P. Harrenstein, W. van der Hoek, J.-J. Meyer, and C. Witteveen. Boolean games. In J. van Benthem, editor, *Proceeding of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge (TARK VIII)*, pages 287–298, Siena, Italy, 2001.
- [16] P. Harrenstein, P. Turrini, and M. Wooldridge. Hard and soft equilibria in Boolean games. In A. Lomuscio, P. Scerri, A. Bazzan, and M. Huhns, editors, *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, pages 845–852, 2014.
- [17] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*. Springer-Verlag: Berlin, Germany, 1992.
- [18] Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems—Safety*. Springer-Verlag: Berlin, Germany, 1995.
- [19] H. Moulin and B. Peleg. Cores of effectivity functions and implementation theory. *Journal of Mathematical Economics*, 10:115–145, 1982.
- [20] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press: Cambridge, MA, 1994.
- [21] M. Pauly. *Logic for Social Software*. PhD thesis, University of Amsterdam, 2001. ILLC Dissertation Series 2001-10.
- [22] A. P. Sistla and E. M. Clarke. The complexity of propositional linear temporal logics. *Journal of the ACM*, 32(3):733–749, 1985.
- [23] M. Wooldridge and P. E. Dunne. On the computational complexity of coalitional resource games. *Artificial Intelligence*, 170(10):853–871, 2006.
- [24] M. Wooldridge, U. Endriss, S. Kraus, and J. Lang. Incentive engineering for boolean games. *Artificial Intelligence*, 195: 418–439, 2013.