

Tradeoffs between Incentive Mechanisms in Boolean Games

Vadim Levit, Zohar Komarovsky
 Ben-Gurion University of the Negev
 Be'er-Sheva, Israel
 {levitv,zoharkom}@cs.bgu.ac.il

Tal Grinshpoun
 Ariel University
 Ariel, Israel
 talgr@ariel.ac.il

Amnon Meisels
 Ben-Gurion University of the Negev
 Be'er-Sheva, Israel
 am@cs.bgu.ac.il

Abstract

Two incentive mechanisms for Boolean games were proposed recently – taxation schemes and side payments. Both mechanisms have been shown to be able to secure a pure Nash equilibrium (PNE) for Boolean games. A complete characterization of outcomes that can be transformed to PNEs is given for each of the two incentive mechanisms. Side payments are proved to be a weaker mechanism in the sense that the outcomes that they can transform to PNEs are a subset of those transformable by taxation. A family of social-network-based Boolean games, which demonstrates the differences between the two mechanisms for securing a PNE, is presented. A distributed search algorithm for finding the side payments needed for securing a PNE is proposed. An empirical evaluation demonstrates the properties of the two mechanisms on the family of social-network-based Boolean games.

1 Introduction

Boolean games are a family of games based on propositional logic [Harrenstein *et al.*, 2001; Bonzon *et al.*, 2006; Dunne *et al.*, 2008], where each participant (agent) holds a distinct set of Boolean variables and has some personal goal it attempts to satisfy. An agent's personal goal is represented as a propositional logic formula over some set of Boolean variables, where some of these variables are not necessarily held by the agent. The actions that an agent can take consist of assigning values to the variables it holds. Therefore, in order to decide on which action to perform, a rational agent must take into consideration two factors – its personal goal and the costs of its actions.

Over the last few years several *incentive mechanisms* were proposed for Boolean games [Wooldridge *et al.*, 2013; Turrini, 2013; Grant *et al.*, 2014]. These mechanisms attempt to influence the agents' preferences so that a certain desirable outcome will result. To be more specific, the goal of the proposed mechanisms was stabilization, i.e., securing the existence of a *pure Nash equilibrium (PNE)* [Osborne and Rubinstein, 1994]. There are two conceptually different ways of affecting the preferences of agents – changing the structure of their personal goals, or changing the costs of actions.

The mechanisms that are at the focus of the present paper influence agents by changing the costs of their actions. The taxation mechanism assumes the existence of an external agent – the *principal*. This principal has the ability to impose *tax* on agents' actions [Wooldridge *et al.*, 2013]. When an agent reasons on what action to perform, it must take into consideration both the marginal costs and the tax levied on its actions. Another mechanism for manipulating the costs of agents' actions in Boolean games is that of *side payments* [Turrini, 2013]. The side payments mechanism enables agents in Boolean games to sacrifice part of their payoff for some given outcome in order to convince other players to play a certain strategy.

The two mechanisms differ in their applicability to various scenarios. While side payments are more appropriate in situations where no central authority exists, such as social-network games [Clercq *et al.*, 2014b], there also exist real-world Boolean games applications that do rely on the presence of such authority. Examples include a robot performing tasks in an automated warehouse [Harrenstein *et al.*, 2014] and scheduling charging times of electric vehicles [Levit *et al.*, 2013a]. Nevertheless, there are applications where both incentive mechanisms are naturally applicable. For example, in a scenario where employees are being allocated to projects [Clercq *et al.*, 2014a], taxation may be naturally imposed by the management, while employees can "pay" their colleagues by assisting them with their tasks.

In case a Boolean game allows to apply both mechanisms, the deep knowledge regarding the abilities/restrictions of the above two mechanisms is needed in order to make a decision regarding the mechanism to choose. Addressing the question, Harrenstein *et al.* [2014] showed that an *equilibrium can be eliminated* (with additional restrictions) using side payments or taxation. On the other hand, Turrini [2013] theoretically showed that for every Boolean game and every transfer function (of side payments) one can find a taxation scheme that secures all equilibria that the transfer function does. Turrini's work does not provide any results in the other direction. In other words, what can one say about equilibria that can be secured by a taxation scheme – can they also be secured by side payments?

The present study focuses on this remaining question about the differences between the two mechanisms. First, a theoretical characterization of Boolean games for which a pure Nash

equilibrium can be secured using side payments is presented. Next, the analogous characterization of Levit *et al.* [2013b] for a taxation scheme is considered. The theoretical comparison of the two mechanisms shows that side payments are a weaker mechanism for securing a PNE in Boolean games.

Two separate important questions arise in view of the theoretical differences between taxation schemes and side payments. First, what are the differences between the PNEs that are secured by each of these two mechanisms? Second, what are the differences in the overall change of cost needed to secure a stable state? An extensive empirical evaluation addresses these two questions and uses *social-network-based Boolean games* which initially do not have a PNE, but for which there is at least one taxation scheme that can secure its existence. An effective distributed search algorithm for Asymmetric Distributed Constraint Optimization Problems (ADCOP) [Grinshpoun *et al.*, 2013] is used for finding the appropriate side payments.

The plan of the paper is as follows. First, Boolean games are presented in Section 2. Taxation and side payments, which are the incentive mechanisms at focus, are described in detail in Section 3. Section 4 studies the theoretical properties of the outcomes that can be transformed to a PNE by each of the mechanisms. Boolean games in social networks are introduced in Section 5, followed by a description of the search procedure for side payments. An extensive empirical evaluation of social-network-based Boolean games that compares between the two incentive mechanisms is given in Section 6. Section 7 outlines our conclusions.

2 Boolean games

A Boolean game consists of a set of agents $A = \{1, \dots, n\}$, the players of the game. Each agent $i \in A$ controls a set of Boolean variables (φ_i is the set of variables controlled by agent $i \in A$) [Harrenstein *et al.*, 2001; Bonzon *et al.*, 2006; Dunne *et al.*, 2008]. This means that only agent i can set the values for each variable $p \in \varphi_i$. $\varphi_1, \dots, \varphi_n$ form a partition of the game variables Φ .

Each agent has a *personal goal*, represented by a Boolean formula γ_i . Every goal γ_i may contain the variables of agent i and possibly variables controlled by other agents. It is assumed that actions of agents have *costs* defined by a cost function $c : \Phi \times \mathbb{B} \rightarrow \mathbb{R}_{\geq}$, where \mathbb{R}_{\geq} denotes the set of non-negative real numbers [Dunne *et al.*, 2008].

A *choice* of agent $i \in A$, defined by a function $v_i : \varphi_i \rightarrow \mathbb{B}$, is an allocation of truth (\top) or falsity (\perp) to all of the agent's variables, φ_i . Let V_i denote the set of all available choices for agent $i \in A$. The intuitive interpretation of V_i is that it defines all actions or strategies available to agent i .

An *outcome* $v = (v_1, \dots, v_n) \in V_1 \times \dots \times V_n$ is a collection of choices, one for each agent. Every outcome uniquely defines a valuation for all variables in the game.

Similarly to [Dunne *et al.*, 2008; Wooldridge *et al.*, 2013], a Boolean Game G is a $2n + 3$ tuple:

$$G := \langle A, \Phi, c, \gamma_1, \dots, \gamma_n, \varphi_1, \dots, \varphi_n \rangle$$

where $A = \{1, \dots, n\}$ is the set of agents, $\Phi = \{p, q, r, \dots\}$ is a finite set of Boolean variables, $c : \Phi \times \mathbb{B} \rightarrow \mathbb{R}_{\geq}$ is a cost

function of assignments, $\gamma_1, \dots, \gamma_n$ are the goals of agents A , and $\varphi_1, \dots, \varphi_n$ is a partition of variables Φ over agents A .

The primary aim of each agent $i \in A$ is to choose an assignment to the variables φ_i under its control, so as to satisfy its personal goal γ_i . If an agent can achieve its personal goal in more than one way, then it will prefer to minimize costs. If the agent cannot get its goal achieved it will prefer to choose an assignment that minimizes its costs.

We will adopt the agents' utility proposed by Wooldridge *et al.* [2013], that models the preference dichotomy using a *boost factor* μ_i for states satisfying agent i 's goal.

Let $c_i(v_i)$ denote the cost to agent i of its choice v_i :

$$c_i(v_i) := \sum_{p \in \varphi_i} c(p, v_i(p))$$

Let μ_i denote the cost of agent i 's worst outcome:

$$\mu_i := \max_{v_i \in V_i} (c_i(v_i))$$

The utility of i from outcome $v = (v_1, \dots, v_i, \dots, v_n)$ is:

$$u_i(v) := \begin{cases} 1 + \mu_i - c_i(v_i), & \text{if } v \models \gamma_i \\ -c_i(v_i), & \text{otherwise} \end{cases}$$

where $v \models \gamma$ means that the valuation defined by the outcome v satisfies formula γ .

The utility for agent i ranges from $1 + \mu_i$ (the best outcome for i , where it gets its goal achieved by performing actions that have zero cost) down to $-\mu_i$ (where i does not get its goal achieved but makes its most expensive choice). Since costs are never negative, agent's utility is positive at goal-achieving states.

Given the above formal definition of utility, one can define solution concepts in the standard game-theoretic way [Osborne and Rubinstein, 1994]. The present paper focuses on PNEs. An outcome $v = (v_1, \dots, v_i, \dots, v_n)$ is a PNE, if for every agent $i \in A$, there is no $v'_i \in V_i$ such that $u_i(v_1, \dots, v_i, \dots, v_n) < u_i(v_1, \dots, v'_i, \dots, v_n)$.

3 Incentive mechanisms

Several incentive mechanisms for Boolean games, which attempt to influence the preferences of rational agents, were proposed recently. One such mechanism that uses *taxation schemes* was proposed by Wooldridge *et al.* [2013]. This method introduces an external principal that imposes additional costs (taxes) on the actions of agents. Another mechanism, proposed by Turrini [2013], extends Boolean games with the machinery of *side payments*. The basic idea of side payments is that agents may sacrifice part of their payoff (for some given outcome) in order to convince other players to play a certain strategy. These potential sacrifices, termed *transfer functions*, are simultaneously decided upon by the different agents in a pre-play phase [Jackson and Wilkie, 2005]. Posterior to this pre-play phase follows the actual game, in which the utilities are updated according to the transfer functions. After the game is played, only the side payments relevant to its final outcome are actually being transferred between the players.

These two mechanisms are similar in nature, as they both influence the decision of rational agents by changing the costs of their actions. However, as shown next, there exist Boolean games for which one mechanism can secure a PNE while the other can not.

3.1 Taxation schemes

A taxation scheme [Wooldridge *et al.*, 2013; Levit *et al.*, 2013b] defines additional costs on actions, over those given by the cost function c . One can model taxes by a function $\tau : \Phi \times \mathbb{B} \rightarrow \mathbb{R}_{\geq}$, so that $\tau(p, b)$ is the tax that should be levied on the agent controlling variable $p \in \Phi$ in case the value $b \in \mathbb{B}$ is assigned. An external *principal*, which prefers certain outcomes over others, is at liberty to impose a taxation scheme to fit its requirements. Agents always seek to minimize their costs, so by assigning different taxation schemes the principal can incentivize agents to perform some actions over others.

We adopt the assumption of Wooldridge *et al.* [2013], under which taxation schemes cannot change the personal goal of an agent. Therefore, if an agent has a chance to achieve its goal, it will take it, no matter what the taxation incentives are.

Since the agents must be aware of the taxation while making a decision, an agent's utility must be extended so as to take the taxes into consideration [Wooldridge *et al.*, 2013]. The taxation functions are:

$$\tau_i(v_i) := \sum_{p \in \varphi_i} \tau(p, v_i(p))$$

where $\tau_i(v_i)$ represents the total tax that would be levied on player i if it chooses v_i . Let μ_i denote the cost to i of its most expensive course of action:

$$\mu_i := \max_{v_i \in V_i} (c_i(v_i) + \tau_i(v_i))$$

The utility to agent i of an outcome $v = (v_1, \dots, v_i, \dots, v_n)$, is:

$$u_i(v) := \begin{cases} 1 + \mu_i - (c_i(v_i) + \tau_i(v_i)), & \text{if } v \models \gamma_i \\ -(c_i(v_i) + \tau_i(v_i)), & \text{otherwise} \end{cases}$$

The utility of agent i ranges from $1 + \mu_i$ (the best outcome for i , where it achieves its personal goal by performing actions that have no tax or any other cost) down to $-\mu_i$ (where i does not get its goal achieved but makes its most expensive choice).

3.2 Side payments

Contrary to a taxation scheme, side payments do not assume the existence of a principal. They enable Boolean games to be transformed from the inside, by endowing agents with the possibility of sacrificing part of their payoff in order to convince other agents to play a certain strategy. We adopt the Boolean transfer functions $\beta_i : V \times A \rightarrow \mathbb{R}_{\geq}$ of Turriani [2013], where $V = \prod_{i \in A} V_i$. Each such function specifies the payoff that agent i secures to other agents in case some outcome occurs. We assume that a Boolean transfer function can be defined only for a PNE outcome and that the transfer will be performed only when a PNE is played. We term this family of functions *active transfer functions*.

Similarly to taxation schemes, an agent's utility is extended to take side payments into consideration while deciding on the action to take:

$$\beta_i(v) := \sum_{j \in A} \beta_i(v, j) - \sum_{j \in A} \beta_j(v, i)$$

where $\beta_i(v)$ represents the *net loss* that agent i incurs at outcome v when using the Boolean transfer function β . We restrict our attention to transfer functions that assure that for all agents $i \in A$ and all outcomes v ,

$$c_i(v_i) \geq -\beta_i(v)$$

put in words, the cost is always higher than the net gain. Another restriction relates to the maximal amount of payoff that an agent can sacrifice. Such a restriction mimics the traditional view of games with side payments [Jackson and Wilkie, 2005], where the actions have utilities rather than costs. In such situations it is natural to assume that an agent cannot offer a side payment that is larger than its utilitarian benefit. One can define the following boost factor:

$$\mu_i^\beta := \max_{v \in V} (c_i(v_i) + \beta_i(v)) \quad (1)$$

The utility that i obtains from outcome $v = (v_1, \dots, v_i, \dots, v_n)$ is:

$$u_i^\beta(v) := \begin{cases} 1 + \mu_i - (c_i(v_i) + \beta_i(v)), & \text{if } v \models \gamma_i \\ -(c_i(v_i) + \beta_i(v)), & \text{otherwise} \end{cases}$$

This means that the maximal amount of payoff that the agent can sacrifice is $2 \cdot \mu_i + 2$, since the maximal utility the agent can receive is $1 + \mu_i$ (it achieves its personal goal by performing actions that have no cost) while the worst utility is $-\mu_i$ (in case the agent takes its most expensive action and does not achieve its personal goal). Hereinafter, we use μ as defined in Equation 1. Nonetheless, other values of μ can be used to represent different limitations of the maximal sacrificable payoff restriction. In the extreme, $\mu = \infty$ reflects the dichotomous nature of Boolean games, while assuming the availability of infinite funds for each agent.

4 Existence of a PNE in Boolean games

Not every Boolean game has a PNE, but one can try to secure the existence of a PNE by using an incentive mechanism. The set of Boolean games (and outcomes in these games) for which a stable state can be secured using taxation schemes was characterized by Levit *et al.* [2013b].

Definition 1. A *Special Outcome (SO)* is an outcome $v = (v_1, \dots, v_i, \dots, v_n)$ in which there exists an agent ($i \in A$) for whom the following conditions hold:

- $(v_i, v_{-i}) \not\models \gamma_i$ – the agent does not achieve its personal goal
- $\exists v'_i \in V_i, (v'_i, v_{-i}) \models \gamma_i$ – it can achieve it by a unilateral deviation

Using the above definition, Levit *et al.* [2013b] proved that:

Proposition 1. An outcome can be transformed to a PNE (with the appropriate taxation scheme) if and only if this outcome is not a special outcome.

This result is needed for the comparison between taxation schemes and side payments.

4.1 Securing stable states with side payments

Let G be a Boolean game and $v = (v_1, \dots, v_i, \dots, v_n)$ an outcome. The agents in the game forms two distinct groups:

- $S = \{i \in A \mid \forall v'_i \in V_i, u_i(v_i, v_{-i}) \geq u_i(v'_i, v_{-i})\}$
agents that rationally chooses to stay in outcome v
- $D = \{i \in A \mid \exists v'_i \in V_i, u_i(v_i, v_{-i}) < u_i(v'_i, v_{-i})\}$
agents that rationally choose to deviate from outcome v

The *affecting cost change* for agent i in outcome v :

$$g_i(v) := \begin{cases} u_i(v) - \max_{v'_i \in V_i \setminus \{v_i\}} (u_i(v'_i, v_{-i})), & \text{if } i \in S \\ \max_{v'_i \in V_i \setminus \{v_i\}} (u_i(v'_i, v_{-i})) - u_i(v), & \text{if } i \in D \end{cases}$$

Intuitively, $g_i(v)$ defines the maximal net loss (from a transfer function) that an agent i can suffer but still rationally choose to stay at outcome v (in case $i \in S$), or the minimal net gain that agent i must obtain in order to rationally prefer an outcome v over other outcomes (in case $i \in D$). The total gain/loss of an outcome v are:

$$Gain(v) := \sum_{i \in S} g_i(v), \quad Loss(v) := \sum_{i \in D} g_i(v)$$

Proposition 2. *Given a Boolean game G and an outcome $v = (v_1, \dots, v_i, \dots, v_n)$, if*

- v is not a special outcome
- $Gain(v) \geq Loss(v)$

then there exists an active Boolean transfer function that transforms v to a PNE.

Proof. To show the correctness of Proposition 2, we construct an active Boolean transfer function. For each agent $i \in S$ we define the *rate of participation*:

$$p_i(v) = \frac{g_i(v)}{Gain(v)}$$

Note that $Gain(v) = 0$ implies either $Loss(v) = 0$ (i.e., no side payments needed) or there is no active Boolean transfer function that transforms v to a PNE. Hence, zero division will not occur when $p_i(v)$ is computed. Now, we define the active Boolean transfer function as follows:

$$\beta_i(v, j) = \begin{cases} p_i(v) \cdot g_j(v), & \text{if } i \in S, j \in D \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

Note that the definition of active transfer functions implies that $\beta_i(v', j) = 0$ for each outcome $v' \neq v$.

By Equation 2, every agent $j \in D$ has a net gain of: $-\beta_j(v) = \sum_{i \in S} p_i(v) \cdot g_j(v) = \sum_{i \in S} \frac{g_i(v)}{Gain(v)} \cdot g_j(v) = g_j(v)$. In other words, agent j obtains enough net gain to prevent it from deviating from outcome v .

In addition, every agent $i \in S$ suffers a maximal net loss of: $\beta_i(v) = \sum_{j \in D} p_i(v) \cdot g_j(v) = \sum_{j \in D} \frac{g_i(v)}{Gain(v)} \cdot g_j(v) = \frac{Loss(v)}{Gain(v)} \cdot g_i(v) \leq g_i(v)$ (the last inequality follows from $Gain(v) \geq Loss(v)$, hence $\frac{Loss(v)}{Gain(v)} \leq 1$). Every agent $i \in S$ suffers at most $g_i(v)$ net loss, therefore i still does not prefer to deviate from outcome v . Consequently, v is a PNE, since there is no agent (from both groups S and D) who wants to unilaterally deviate from it. \square

Proposition 3. *Given a Boolean game G and an outcome $v = (v_1, \dots, v_i, \dots, v_n)$, if there exists an active Boolean transfer function, that transforms v to a PNE then the following conditions must hold:*

- v is not a special outcome
- $Gain(v) \geq Loss(v)$

Proof. First, v is not a special outcome, otherwise some agent (i) would prefer to change its choice to satisfy its personal goal. Next, one needs to show that $Gain(v) \geq Loss(v)$. Denote by $T(v)$ the net sum of transfers between agents in S and agents in D :

$$T(v) = \sum_{i \in S, j \in D} \beta_i(v, j) - \beta_j(v, i)$$

Assume by contradiction that $T(v) > Gain(v)$, then the sum of updated utilities (according to the given transfer functions) of agents in S is decreased by more than $Gain(v)$. Hence, at least one agent will rationally choose to deviate from v in contradiction of v being a PNE.

Assume by contradiction that $Loss(v) > T(v)$, then the sum of the updated utilities (according to the given transfer functions) of agents in D is increased by less than $Loss(v)$. This means that at least one agent in D still rationally chooses to deviate from v , in contradiction of v being a PNE. Consequently, $Gain(v) \geq T(v)$ and $T(v) \geq Loss(v)$, which means $Gain(v) \geq Loss(v)$. \square

Corollary 4. *An outcome v can be transformed to a PNE by the use of active side payments if and only if the two conditions defined in Propositions 2 and 3 hold.*

4.2 Taxation vs. side payments

The difference between taxation schemes and side payments follows immediately from Proposition 1 and Corollary 4:

Corollary 5. *Every PNE that can be secured with active side payments can be also secured with a taxation scheme, but the opposite is not true.*

The differences in the functionality of the two incentive mechanisms are exemplified next.

Example 1. *Consider two agents $A = \{1, 2\}$. Each agent owns a single variable (i.e., $\varphi_1 = \{a\}$ and $\varphi_2 = \{b\}$). The personal goal of agent 1 is $\gamma_1 = \neg a \wedge \neg b$ and the personal goal of agent 2 is $\gamma_2 = \neg a \wedge b$. The costs associated with the actions are: $c(a, \perp) = 10, c(a, \top) = 0, c(b, \perp) = 0$ and $c(b, \top) = 4$.*

One can easily verify that there is no PNE in this Boolean game. Suppose the outcome $v = \{a = \top, b = \top\}$ is a desired one (to become a PNE). Let us see how this can be secured by using taxation and active side payments:

- Taxation scheme – simply apply a tax of at least 4 for the assignment $b = \perp$ (i.e., $\tau(b, \perp) \geq 4$)
- Active side payments – agent 1 must transfer to agent 2 a payoff of 4 (i.e., $\beta_1(\{a = \top, b = \top\}, 2) = 4$)

For this specific outcome (v), we were able to find both a taxation scheme and active side payments that secure the existence of a PNE. However, this is not true in general. Reconsider Example 1 and the outcome $v' = \{a = \perp, b = \top\}$. It is possible to find a taxation scheme that secures a PNE at outcome v' (e.g., $\tau(a, \top) = 10$). However, with the μ of Equation 1, there is no active side payment that can transform outcome v' to a PNE.

Note that taxation is an *indirect* operation, in the sense that when the principal chooses to levy a tax in order to perturb the agents' preferences, the tax is chosen in a manner that no agent directly pays it. Contrary to that, in active side payments the operation is *direct*, in the sense that the side payments are directly transferred among the agents.

5 Boolean games in social networks

Having investigated the theoretical differences between taxation schemes and side payments, it is natural to compare them empirically. In order to be able to perform such a comparison, we will first define a family of Boolean games that initially do not have a PNE. By applying one of the two incentive mechanisms, a stable state can be secured for this family of Boolean games.

5.1 Coordination in Boolean games

To make things simple, let us start from Boolean games of two agents $A = \{1, 2\}$, each owning a single variable ($\varphi_1 = \{a\}$ and $\varphi_2 = \{b\}$). If the personal goals of agents require pure *coordination*¹ (or anti-coordination), then every such Boolean game will initially have a PNE.

Definition 2. A 2-players Boolean game has a conflicting interaction if the personal goals take the form of $\gamma_1 = (a \wedge b) \vee (\neg a \wedge \neg b)$ and $\gamma_2 = (\neg a \wedge b) \vee (a \wedge \neg b)$. The agents have a disagreement among their personal goals (one agent wants pure coordination and the other anti-coordination).

The Boolean game of Definition 2 does not have a PNE and there is no taxation scheme that can secure its existence. To create a Boolean game that can serve for comparing the PNE securing schemes we will relax one of the agents' personal goals as follows:

Definition 3. A 2-players Boolean game has a manipulable conflicting interaction if the personal goals are $\gamma_1 = a \vee b$ and $\gamma_2 = (a \wedge b) \vee (\neg a \wedge \neg b)$.

There is only one outcome that is not a special outcome in this Boolean game. In addition, if $c(a, \perp) < c(a, \top)$ this game will have no stable state. This means, that the Boolean game at hand meets our requirements.

An interesting way to generate more complex inter-agent interactions is by introducing an underlying *social network*, where the nodes represent agents and the links represent the interactions between them. This is a natural representation for a Boolean game, where graphical connections stand for variables of agents that appear in other agents' personal goals. In the family of social-network-based Boolean games that will be used in the empirical evaluation, each dependency among

neighboring agents (i.e., connected by a link in the network) takes the form of the simple game described in Definition 3. Consequently, the personal goals of agents are constructed by a disjunction of all inter-neighbor interactions in the underlying social network.

Proposition 6. One can secure the existence of a stable state in a social-network-based Boolean game, generated according to the rules above by levying a taxation scheme.

Proof. Suppose every agent chooses \top as the assignment for its variable. This results in an outcome where every agent achieves its personal goal. Therefore, this outcome is not a special outcome and consequently there is a taxation scheme that can secure a PNE in this Boolean game. \square

5.2 ADCOP

An asymmetric distributed constraints optimization problem (ADCOP) [Grinshpoun *et al.*, 2013] is a tuple

$$\langle A, X, D, R \rangle$$

where $A = \{A_1, A_2, \dots, A_n\}$ is a finite set of agents. $X = \{X_1, X_2, \dots, X_m\}$ is a finite set of variables. Each variable is held by a single agent (an agent may hold more than one variable). $D = \{D_1, D_2, \dots, D_m\}$ is a set of domains. Each domain D_i consists of a finite set of values that can be assigned to variable X_i . R is the set of relations (constraints). Each constraint $C \in R$ is a function $C : D_{i_1} \times D_{i_2} \times \dots \times D_{i_k} \rightarrow \prod_{j=1}^k \mathbb{R}_{\geq}$ that defines a non-negative cost for every participant in every value combination of a set of variables. The *asymmetry* of constraints in the ADCOP model stems from the potentially different costs of constraints for every participant.

A *complete assignment* consists of assignments to all variables in X . A *solution* to an ADCOP is a complete assignment of minimal cost.

5.3 Searching for side payments with ADCOPs

The main idea is to construct an ADCOP and apply a search on it. The objective of the search is to find the active transfer function that secures the existence of a PNE and imposes the minimal payoff transfer. Given a Boolean game G we define the ADCOP as follows:

- The set of agents in the ADCOP is the same as in G .
- The variables of the ADCOP are those of G with the same variable allocation.
- Every domain D_i consists of two values (0 for \perp and 1 for \top).
- For every agent A_i construct a constraint. This constraint includes valuations of variables that appear in γ_i (in what follows the variables that appear in γ_i will be denoted by $(v_{i_1}, \dots, v_{i_k})$). The constraints of the ADCOP are not in the form of a table, but are rather on-the-fly computed during search from a formula that takes constant computation time. The details of the constraints are described next.

¹Coordination games [Cooper, 1999].

One wants to select only those outcomes v that are not a SO and for which $Gain(v) \geq Loss(v)$. For every special outcome or outcome with $Gain(v) < Loss(v)$ the relation cost should be larger than the maximal possible net sum of transfers. Active transfer function values may vary, but since the agents will rationally choose them, the values of these functions should be minimal. Thus, the maximal possible net sum of transfers M can be computed using the following equation:

$$M := \sum_{p \in \Phi} |c(p, \top) - c(p, \perp)|$$

This leads to the following constraint in the ADCOP:

$$C_i(v_{i_1}, \dots, v_{i_k}) := M + \epsilon \quad (3)$$

The constraint defined by Equation 3 is applied to every outcome v that is a SO or for which $Gain(v) < Loss(v)$. Out of all remaining outcomes we want to find the one that minimizes the net sum of transfers. Therefore, the cost of a constraint should be the needed (minimal) payoff transfer:

$$C_i(v_{i_1}, \dots, v_{i_k}) := \begin{cases} g_i(v) & \text{if } i \in D \\ 0 & \text{if } i \in S \end{cases} \quad (4)$$

Equation 4 describes the minimal payoff transfer to a single agent i . This transfer ensures that agent i will not have any incentive to change its choice.

The ADCOP's solution is a full assignment (v_1, v_2, \dots, v_n) that represents a PNE state when the appropriate active transfer function is used. The active side payment is calculated during the search process and can be stored along with its matching assignment.

6 Experimental evaluation

The differences between side payments and taxation schemes are evaluated on social-network-based Boolean games. Problems were randomly generated and the reported results are averages over 100 instances for each setting.

6.1 Problem generation

For each experiment a random problem was generated. First, an Erdős-Rényi [1959] random network was generated. Next, the Boolean game was constructed according to the rules described in Section 5.1, where the cost of assigning \top was chosen from the range $[100, 200)$ and the cost of assigning \perp from the range $[0, 100)$. Then, an ADCOP was generated from the Boolean game using the procedure described in Section 5.3. Finally, the problem was solved using the k -ary *SyncABB-1ph* algorithm [Levit *et al.*, 2013b].

6.2 Experimental results

The first part of the experimental evaluation compares the solution quality of the two incentive mechanisms. For this purpose two measures are considered.

Figure 1 presents the percentage of games that have a PNE state after applying each of the incentive mechanisms. Every game in this experimental set did not have a PNE initially. The games were designed so that it will be possible to ensure the existence of a PNE using taxation (see Proposition 6). We




Figure 1: Percentage of games that have a PNE after applying the incentive mechanisms

can see that side payments are able to secure a PNE in about 80% of the instances, on average. In addition, it is evident that the probability for ensuring a PNE with active side payments drops when the problems become larger.

The second measure is the size of the *overall cost change* with respect to the original costs of the game. The respective percentages of the cost change for taxation schemes and side payments are calculated as follows:

$$\frac{\sum_{p \in \Phi, b \in \mathbb{B}} \tau(p, b)}{\sum_{p \in \Phi, b \in \mathbb{B}} c(p, b)} \cdot 100\%, \quad \frac{\sum_{i \in S, j \in D} \beta_i(v, j)}{\sum_{p \in \Phi, b \in \mathbb{B}} c(p, b)} \cdot 100\%$$




Figure 2: Percentage of cost change

For this results we selected only those problems for which a PNE can be ensured with active side payments. Figure 2 shows that the overall cost change is very small (less than 6% in all of the problem settings in the evaluation). This means, that for the set of games which were used in our experiments, there are only a few agents who need to be manipulated in order to ensure existence of a stable state.

In order to evaluate the algorithms, we consider the mean number of *Non-Concurrent Constraint Checks (NCCCs)*, which is a commonly used measure for the runtime performance of distributed constraints search algorithms [Zivan and Meisels, 2006]. The exponential growth of the computational load with respect to the problem size is clearly seen in Figure 3 and is of no surprise as ADCOPs are NP-Hard problems. However, the type of incentive mechanism (i.e., taxation or side payments) does not significantly affect the performance.




Figure 3: Mean number of NCCCs

7 Conclusion

Two incentivising mechanisms for Boolean games where studied in the last three years – taxation schemes [Wooldridge *et al.*, 2013; Levit *et al.*, 2013b] and side payments [Turrini, 2013]. Both mechanisms were shown to be able to eliminate (with additional restrictions) pure Nash equilibria (PNEs) [Harrenstein *et al.*, 2014]. The present study focuses on the differences between the two mechanisms and shows that side payments are a weaker mechanism for securing a PNE. The set of outcomes of Boolean games that can be transformed to a PNEs by the use of side payments is fully characterized and is found to be a subset of outcomes that can become PNEs by the use of a taxation scheme [Levit *et al.*, 2013b].

An extensive empirical evaluation was performed on social-network-based Boolean games, which initially do not have a PNE. A distributed search algorithm looked for solutions of minimal cost change for securing a PNE by each of the two mechanisms. The evaluation provides two interesting insights about the two mechanisms. First, side payments secure a PNE in only $\sim 80\%$ of the problems. Second, the overall change of cost of the games is very small (less than 6%). This points to the fact that the incentive mechanisms need to only influence a small number of agents, in order to secure a PNE.

Acknowledgments

The research was partially supported by the Lynn and William Frankel Center for Computer Sciences at Ben-Gurion University and by the Paul Ivanier Center for Robotics Research and Production Management.

References

[Bonzon *et al.*, 2006] Elise Bonzon, Marie-Christine Lagasquie-Schiex, Jérôme Lang, and Bruno Zanuttini. Boolean games revisited. In *ECAI*, pages 265–269, 2006.

[Clercq *et al.*, 2014a] Sofie De Clercq, Kim Bauters, Steven Schockaert, Martine De Cock, and Ann Nowé. Using answer set programming for solving boolean games. In *KR*, page 4, 2014.

[Clercq *et al.*, 2014b] Sofie De Clercq, Kim Bauters, Steven Schockaert, Mihail Mihaylov, Martine De Cock, and Ann Nowé. Decentralized computation of pareto optimal pure nash equilibria of boolean games with privacy concerns. In *ICAART*, pages 50–59, 2014.

[Cooper, 1999] Russell Cooper. *Coordination games*. Cambridge University Press, 1999.

[Dunne *et al.*, 2008] Paul E. Dunne, Wiebe van der Hoek, Sarit Kraus, and Michael Wooldridge. Cooperative Boolean games. In *AAMAS*, pages 1015–1022, 2008.

[Erdős and Rényi, 1959] Paul Erdős and Alfréd Rényi. On random graphs. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.

[Grant *et al.*, 2014] John Grant, Sarit Kraus, Michael Wooldridge, and Inon Zuckerman. Manipulating games by sharing information. *Studia Logica*, 102(2):267–295, 2014.

[Grinshpoun *et al.*, 2013] Tal Grinshpoun, Alon Grubshtein, Roie Zivan, Arnon Netzer, and Amnon Meisels. Asymmetric distributed constraint optimization problems. *Journal of Artificial Intelligence Research (JAIR)*, 47:613–647, 2013.

[Harrenstein *et al.*, 2001] Paul Harrenstein, Wiebe van der Hoek, John-Jules Meyer, and Cees Witteveen. Boolean games. In *TARK*, pages 287–298, 2001.

[Harrenstein *et al.*, 2014] Paul Harrenstein, Paolo Turrini, and Michael Wooldridge. Hard and soft equilibria in boolean games. In *AAMAS*, pages 845–852, 2014.

[Jackson and Wilkie, 2005] Matthew O Jackson and Simon Wilkie. Endogenous games and mechanisms: Side payments among players. *The Review of Economic Studies*, 72(2):543–566, 2005.

[Levit *et al.*, 2013a] Vadim Levit, Tal Grinshpoun, and Amnon Meisels. Boolean games for charging electric vehicles. In *IAT*, pages 86–93, 2013.

[Levit *et al.*, 2013b] Vadim Levit, Tal Grinshpoun, Amnon Meisels, and Ana L. C. Bazzan. Taxation search in boolean games. In *AAMAS*, pages 183–190, 2013.

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

[Turrini, 2013] Paolo Turrini. Endogenous boolean games. In *IJCAI*, pages 390–396, 2013.

[Wooldridge *et al.*, 2013] Michael Wooldridge, Ulle Endriss, Sarit Kraus, and Jérôme Lang. Incentive engineering for boolean games. *Artificial Intelligence*, 195:418–439, 2013.

[Zivan and Meisels, 2006] R. Zivan and A. Meisels. Message delay and DisCSP search algorithms. *Annals of Mathematics and Artificial Intelligence (AMAI)*, 46:415–439, 2006.