

Delegating Decisions in Strategic Settings

Paul E. Dunne, Paul Harrenstein, Sarit Kraus, and Michael Wooldridge

Abstract—We formalise and investigate the following problem. A number of decisions must be delegated to a collection of agents; once the decisions are delegated, the agents to whom the decisions are delegated will then make these decisions rationally and independently, in pursuit of their own preferences. A principal is able to determine how decisions will be delegated, and seeks to do so in such a way that, when the decisions are ultimately made, some overall goal is satisfied. The *principal delegation problem*, is then, given such a setting, whether it is possible for the principal to delegate decisions in such a way that, if all the agents to whom decisions have been delegated then make their respective decisions rationally, the principal’s goal will be achieved in equilibrium. We also distinguish the *distributed allocation problem* where the agents can delegate decisions among themselves. Here we not only require that the principal’s goal will be achieved in equilibrium, but moreover that the allocation to the agents is *stable*, in the sense that no coalition of agents can redistribute the decisions delegated to them among themselves so as to be better positioned to satisfy their individual goals in equilibrium. We formalise these problems using *Boolean games*, which provides a very natural framework within which to capture the delegation problem: decisions are directly represented as Boolean variables, which the principal assigns to agents. After motivating and formally defining the principal and distributed delegation problems, we investigate these computational complexity of several varieties of these problems, along with some issues surrounding it.

Impact Statement—Our research concerns the ubiquitous problem of how to incentivise agents to decide on courses of action that are simultaneously desirable from the global perspective and rational from a local game-theoretic perspective. The design of mechanisms so as to achieve this is through steering the strategic capabilities of interested agents by delegating decisions to them in specific ways. We distinguish the cases in which the delegation is effectuated by a principal and where the agents are delegating decisions among one another. Our reliance on the mathematical framework of Boolean games to formalise the principal and distributed delegation problems offers valuable insights into the computational complexity of these problems. The delegation problem touches on many important applications, not only in everyday life, but also in Artificial Intelligence, in particular Multi-Agent Systems.

Index Terms—Game theory; mechanism design; logic; knowledge representation; computational complexity.

I. INTRODUCTION

Throughout our lives, we must inevitably delegate to other agents decisions whose outcome will affect us, even though we know full well that the agents we delegate the decisions to are self-interested, and will make these decisions in their own interest. Examples of this problem are legion. Thus, one can think of the chair of a university department, who must allocate teaching and admin responsibilities to faculty members, who may have the interests of their research area

Paul E. Dunne is with the Department of Computer Science, University of Liverpool, UK, Paul Harrenstein and Michael Wooldridge are with the Department of Computer Science, University of Oxford, UK, and Sarit Kraus is with Department of Computer Science, Bar Ilan University, Israel.

	coal phase-out	no coal phase-out	coal phase-out	no coal phase-out
high taxes	sophisticated green economy in A ● ○ ●	low-emission plants ● ○ ○	sophisticated green economy in B ● ○ ○	subsidised emission market ○ ● ●
low taxes	wind/solar energy ● ● ○	status-quo ○ ● ○	nuclear energy ○ ● ●	regulated energy market ● ○ ●
	free-trade block A		free-trade block B	

Fig. 1. The PM’s predicament. Different policy decisions in the departments of Finance, Energy, and Trade affect the future course the country is to take. The three circles in the bottom right of each cell indicate the preferences of the three ministerial candidates, the environmentalist, the free-trade libertarian, and populist, respectively. A filled circle indicates approval, an empty circle disapproval. The PM strives for a joint equilibrium decision that at a majority of the ministers approve.

prevail over a well-balanced curriculum, diversity of research, or a well-balanced budget. Or of a film studio producing a film and has to appoint a director, a screen writer, as well as a casting manager. In computer science, the designer of a multi-agent system can be seen to delegate tasks to multiple largely autonomous agents. In each case some leeway in the execution of the delegated task has to be granted, and we face the problem how to best delegate which responsibilities to which decision makers. In this paper, we formally analyse this issue, which we will refer to as the *principal delegation problem*, so as to investigate its computational properties. We also study a more cooperative setting, where the agents have to assign responsibilities among one another in the absence of a principal. This setting would be more appropriate, for instance, in the case where a couple of friends who are organising a party and who have to divide the tasks of finding a venue, hiring a DJ, and engaging a catering service. Or in the case where four allied powers occupying an enemy city have to decide who is to be responsible for the city’s food supply, its infrastructure, and its public security. This problem we refer to as the *distributed delegation problem*.

To illustrate the principal delegation problem we use the following fictional political scenario as a running example.

Example 1: Consider the case of a Prime Minister of a not further specified Western European country, who has to reshuffle their cabinet. There are vacancies in the departments of Finance, of Trade, and of Energy. In the upcoming months, in each of these departments there is an important decision to be made. The minister of Finance will have to decide

whether or not to raise taxes, the minister of Trade will have to decide whether to renegotiate a free-trade agreement with trade-block A or enter a new one with trade-block B, and the Energy minister will have to commit to a coal-phase out or pursue other measures. How these decisions will affect the (environmental) future of the country is tabulated in Fig. 1.

There are three candidates, an environmentalist, a free-trade libertarian, and a populist, each of whom has different (but dichotomous) preferences over what the country's future should look like. Thus the environmentalist prefers green (italicised) outcomes, and the free-trade libertarian prefers not to raise taxes, unless the country ends up in trade-block B without a coal phase-out. The populist wants to join trade-block B, unless there is going to be a green economy, in which case only A is acceptable.

The prime minister is confronted with the task of appointing ministers for each of the three ministries from among the three candidates. Different ways of appointing the three ministers lead to different outcomes, as the candidates will have to pursue their goals with different strategic capabilities. Moreover, the prime minister themselves may have their own objectives, in our case, a stable cabinet in which a majority of the three ministers supports the political course that results as a consequence of the ministers' decisions. How can the prime minister achieve this?

By making the libertarian Finance minister, assigning the populist to Energy, and putting the environmentalist in charge of the Trade department, it can easily be checked that there is a unique Nash equilibrium, in the sense that it is the only outcome where none of the ministers would like to revise their decision unilaterally. This equilibrium would result in a future based on solar and wind energy. Alternatively, if the environmentalist, the libertarian, and the populist are responsible for Energy, Trade, and Finance, respectively, another but also unique equilibrium ensues, resulting in a sophisticated green economy in block A. In either case, the prime minister's objective of majority support in the cabinet is fulfilled as well. This will not always be the case though. For instance, if the populist were to decide which trade-block to join, the libertarian whether to raise taxes, and the environmentalist on the coal phase-out, then there is no equilibrium, let alone one that is acceptable to the Prime Minister. \triangleleft

The formal model we use to frame the delegation problem is based on *Boolean games*, which were originally introduced by [1, 2, 3, 4, 5]. In a Boolean game, the agents are players, and each player i is assumed to have a goal, represented as a propositional formula γ_i over some set Φ of Boolean variables. Each player i has some subset Φ_i of the variables Φ , with the idea being that the variables Φ_i are under the unique control of player i . The choices/strategies available to i correspond to all the possible allocations of truth (\top) or falsity (\perp) to the variables in Φ_i . A player would ideally like to choose an allocation for its variables so as to satisfy its goal γ_i , but whether i 's goal is in fact satisfied will depend on the choices made by others; and whether their goal is satisfied will in turn

	t	\bar{t}
a	1, 0, 1	<u>1, 1, 0</u>
\bar{a}	1, 0, 0	0, 1, 1
	c	\bar{c}

(a) Allocation $\alpha^{a|t|c}$

	a	\bar{a}
c	<u>1, 0, 1</u>	1, 0, 0
\bar{c}	1, 0, 0	0, 1, 1
	t	\bar{t}

(b) Allocation $\alpha^{c|a|t}$

	t	\bar{t}
c	1, 0, 1	1, 1, 0
\bar{c}	1, 0, 0	0, 1, 0
	a	\bar{a}

(c) Allocation $\alpha^{c|t|a}$

	c	\bar{c}
a	<u>1, 0, 1</u>	<u>1, 0, 0</u>
\bar{a}	1, 0, 0	0, 1, 1
	t	\bar{t}

(d) Allocation $\alpha^{a|c|t}$

Fig. 2. Partial game of Example 1 under four different allocations, $\alpha^{a|t|c}$, $\alpha^{c|a|t}$, $\alpha^{t|c|a}$, and $\alpha^{a|c|t}$, of departments to ministers. In each game, the environmentalist chooses rows, the free-trade libertarian columns, and the populist matrices. Nash equilibria are underlined.

depend on the choice made by i , and others.¹

Example 2: Thus, in Example 1, the binary decisions relating to tax levels, the coal phase-out, and the choice of free-trade agreement can be formalised, respectively, as three propositional variables t , c , and a , which each can take the value \top (tax increase, coal phase out, free-trade with A) or \perp (no tax increase, no coal phase-out, free-trade with B). The players are the candidates and each can be uniquely appointed minister to one or more of the three departments, or to none. \triangleleft

Boolean games have two main features that make them well-suited to studying delegation problems. First, the strategies of the players are given by subsets of variables over which they have unique control. This affords a natural set-theoretic structure. Setting the truth-value of a given propositional variable could be regarded as a subtask that is delegated to the player controlling that variable. The subtask of setting the truth-value of a propositional variable can in principle be delegated to any of the players and independently of the other variables assigned to their control. Second, in a Boolean game, player satisfaction depends in a systematic way on the truth-values assigned to the propositional variables, but not on the identity of the players who assigned the truth-value to the variables. Loosely speaking, different ways in which control over the variables is distributed over the players may affect the *strategic structure* of the game, but not its *preference structure*. Also see Fig. 2 for an illustration of this point.

In the variant of Boolean games that we use to model the delegation problem, we therefore assume that some of the variables in Φ may be initially unallocated, i.e., not assigned to any player's variable set Φ_i . We have seen how the unallocated variables are assigned to the players may essentially affect the Nash equilibria of the resulting game where all variables are allocated. An external *principal* (corresponding to the prime

¹As one reviewer aptly observed, the principal delegation problem could also be modelled as a *Stackelberg game* [6, 7, pp.97–98], where the principal is the leader and the players the followers. Boolean games, however, allow us to focus more on the internal structure of the allocations rather than on the leader-follower hierarchy.

	q	\bar{q}		q	\bar{q}
p	1, 0, 1	0, 1, 0		0, 1, 0	1, 0, 1
\bar{p}	0, 0, 0	<u>0, 0, 0</u>		1, 0, 1	0, 1, 0
	r	\bar{r}		r	\bar{r}

(a) Partial game P under α

	p	\bar{p}		p	\bar{p}
q	<u>1, 0, 1</u>	0, 0, 0		0, 1, 0	1, 0, 1
\bar{q}	0, 1, 0	0, 0, 0		1, 0, 1	0, 1, 0
	r	\bar{r}		r	\bar{r}

(b) Partial game P under α'

Fig. 3. Partial game of Example 3 under allocations α and α' . Player 1 chooses rows, player 2 chooses columns, and player 3 chooses matrices.

minister in the example above) must allocate these variables to players within the game, i.e., the decision about which unallocated variable is assigned to which player is determined by the principal. Once the principal has made an allocation, then the resulting Boolean game is played in the normal way. Thus, a player i is able to determine values for all the variables Φ_i that it is initially allocated, as well as values for the variables that were allocated to them by the principal. Note that the principal is *not* part of the resulting Boolean game: the values chosen for variables Φ are made by the players in the game. Thus the *only* way the principal can influence a game is in choosing the allocation of originally unallocated variables to players. The principal will make an allocation with some overall objective in mind. We represent the objective by a Boolean formula Υ over the variables Φ . If the principal is successful in allocating variables to players, then the result is that players will rationally choose values for variables so that the objective Υ is satisfied in equilibrium. Thus the overall problem faced by the principal is as follows: *Can I assign the unallocated variables to players in such a way that if the players then play the resulting game rationally, my objective Υ will be satisfied in equilibrium?* We refer to this as the *principal delegation problem*. This problem was investigated before in [8] under the name of the *delegation problem*.

The observation that the way variables are allocated to the players essentially affects the Nash equilibria of the games after allocation lies at the heart of the principal delegation problem. It need not only be the principal who may have an interest in how variables are allocated and which Nash equilibria result as a consequence. This could also hold for the players to whom tasks are being delegated. To illustrate this point, consider the following example.

Example 3: Consider the situation wherein three students are to prepare a common meal and have to decide who is to get tomatoes (p) or beans (\bar{p}) from the greengrocer's, cheese (q) or cream (\bar{q}) from the creamery, and spaghetti (r) or rice (\bar{r}) from the grocer's shop. They have varying preferences over the meals they can make with the ingredients bought.

Fig. 3 depicts the two Boolean games that result in two different ways, α and α' , in which these tasks are allocated. In the former, the variables p , q , and r are controlled by player 1, player 2, and player 3, respectively. In the latter, players 1 and 2 have exchanged control over p and q . Suppose both player 1's and player 3's goal is given by $(r \rightarrow (p \wedge \neg q)) \wedge (\neg r \rightarrow (p \leftrightarrow \neg q))$, and player 2's goal by $(r \rightarrow (p \wedge \neg q)) \wedge (\neg r \rightarrow (p \leftrightarrow q))$.

In the game under allocation α , there is one Nash equilibrium, which renders true $\neg p \wedge \neg q \wedge r$ and satisfies no player's goal. By contrast, the unique equilibrium in the game under α' satisfies $p \wedge q \wedge r$ and as such is strictly better for both players 1 and 3 with respect to the equilibrium under α , whereas player 2 is indifferent between the equilibria that result under α and α' . Therefore, the three players could have an incentive to form a coalition and reallocate the variables they control among one another. We could therefore say that allocation α is *blocked* by the players, and therefore fails to be *stable* in the game-theoretic sense of the word. Observe that, by the same reasoning, the coalition consisting of players 1 and 2 could be said to be similarly incentivised to collude and deviate. Interestingly, and as the reader may verify, this is not the case for the coalition with players 1 and 3, the two beneficiaries of the first reallocation! Finally, given allocation α' , no coalition can be found that would like to reallocate their variables so as to guarantee a better Nash equilibrium: any such coalition would have to involve player 2 and at least one other player, but the latter would be worse off in any outcome where player 2 is better off. \triangleleft

In Section V, we formally address this possibility of coalitions of players *blocking* allocations, i.e., the possibility that some coalitions cooperate by reallocating allocated variables among its members, and to do so to their advantage. More precisely, we consider the following *distributed delegation problem*: *Can we assign the unallocated variables to players in such a way that is both stable, i.e., no coalition is incentivised to reallocate the variables among themselves, and will satisfy our objective Υ in equilibrium?*

In the remainder of this paper, we formalise and study the principal and distributed delegation problems, focussing particularly on computational issues. After this introduction, we present the formal preliminaries in Section II. Section III is devoted to the fundamental principal delegation problem, both in its weak and strong forms. In Section IV, we investigate a variant of the principal delegation problem in which the principal seeks an allocation that will result in an equilibrium that maximises some objective function. Section V deals with the distributed delegation problem. In Section VI, we conclude by evaluating our results and discussing related work, and pointing out directions for future research.

II. PRELIMINARY DEFINITIONS

We now introduce the variation of Boolean games that we use in the present paper, which is directly descended from previous Boolean game models (compare, e.g., [1, 3, 4, 5]).

a) Propositional Logic: Let $\{\top, \perp\}$ be the set of Boolean truth-values, with “ \top ” being truth and “ \perp ” being falsity. Let $\Phi = \{p, q, \dots\}$ be a finite, fixed, and non-empty vocabulary of Boolean variables. The set of well-formed formulae of propositional logic over Φ is then constructed using the conventional Boolean operators “ \wedge ”, “ \vee ”, “ \rightarrow ”, “ \leftrightarrow ”, and “ \neg ”, as well as the constants “ \top ” and “ \perp ”. We abuse notation by using \top and \perp to denote both the syntactic constants for truth and falsity respectively, as well as their semantic counterparts. We will often find it useful to abbreviate clauses such as $p \wedge \neg q \wedge r$

to pqr . A *valuation* is a total function $v : \Phi \rightarrow \{\top, \perp\}$, assigning truth or falsity to every Boolean variable. We write $v \models \varphi$ to mean that the propositional formula φ is true under, or satisfied by, valuation v , where the satisfaction relation \models is defined in the standard way. Let $\mathcal{V}(\Phi)$ denote the set of all valuations over Φ , omitting the parameter Φ when clear from the context. We write $\models \varphi$ to mean that φ is valid, i.e., that $v \models \varphi$ holds for all valuations v . We denote the fact that $\models \varphi \leftrightarrow \psi$ by $\varphi \equiv \psi$.

b) *Quantified Boolean Formulas*: As well as propositional logic, we make use of *Quantified Boolean Formulas* (QBFs). QBFs extend propositional logic with quantifiers $\exists X$ and $\forall X$, where $X \subseteq \Phi$. A formula $\exists X \varphi$ asserts that there is some assignment of truth-values to the variables X such that φ is true under this assignment, while a formula $\forall X \varphi$ asserts that φ is true under all assignments of truth-values to the variables X . QBFs are very powerful: for example the satisfiability of a propositional formula φ over variables Φ can be expressed by the QBF $\exists \Phi \varphi$. Quantifiers in QBFs can be nested: the formula $\exists X \forall Y \varphi$ asserts that there is some assignment of values to X such that no matter how values for variables Y are assigned, the formula φ will be satisfied.

c) *Boolean Games and Partial Boolean Games*: A *partial (Boolean) game*, which is formally given by a tuple

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

is populated by a set N of n *agents*, the players of the game. Each player i is assumed to have a *goal*, characterised by a propositional formula γ_i . Each player i *controls* a possibly empty subset Φ_i of the overall finite set of Boolean variables $\Phi = \{p, q, \dots\}$, i.e., i has the unique ability to set the value of each variable p in Φ_i to either \top or \perp . We will require that no variable is controlled by more than one player, i.e., for distinct players i and j , we assume $\Phi_i \cap \Phi_j = \emptyset$. Readers who are familiar with Boolean games might also be expecting to see the requirement that every variable is controlled by an player, but for the moment, we do *not* make this assumption. Thus, in a partial game, it is possible that some variables in Φ are not allocated to players, i.e., that $\Phi_1 \cup \dots \cup \Phi_n \neq \Phi$. We let $\Phi_U = \Phi \setminus (\Phi_1 \cup \dots \cup \Phi_n)$ denote the set of *unallocated variables*. Two extremal points are worth identifying: if $\Phi_U = \Phi$ and if $\Phi_U = \emptyset$. If the former, $\Phi_1 = \dots = \Phi_n = \emptyset$, and so *all* variables are unallocated. In the latter case, there are *no* unallocated variables and so every variable is assigned to some player. Then, $\Phi_U = \emptyset$, and we also refer to the partial game as a *Boolean game*.

d) *Outcomes, Preferences, and Choices*: The *outcomes* of a (partial) Boolean game over Φ are given by the set of valuations $\mathcal{V}(\Phi)$. The preferences of each player i over the outcomes (valuations) are defined by the formula γ_i in a very straightforward way: player i strictly prefers all those outcomes that satisfy its goal γ_i over all those that do not, but is indifferent between outcomes that both satisfy their goal, and between outcomes that both do not satisfy their

goal.² It is convenient to define for each player i a *utility function* $u_i : \mathcal{V}(\Phi) \rightarrow \{0, 1\}$ over outcomes which captures these preferences such that, for all valuations v in $\mathcal{V}(\Phi)$,

$$u_i(v) = \begin{cases} 1 & \text{if } v \models \gamma_i, \\ 0 & \text{otherwise.} \end{cases}$$

Given outcomes v_1 and v_2 , we write $v_1 \succeq_i v_2$ to mean that $u_i(v_1) \geq u_i(v_2)$, with the corresponding strict \succ_i and indifferent \sim_i subrelations defined in the usual way.

When playing a (partial) Boolean game, a player i will aim to choose an assignment of values for the variables Φ_i under their control so as to satisfy their goal γ_i . However, γ_i may contain variables controlled by player distinct from i , who will also be trying to get their goals satisfied; and their goals in turn may be dependent on the variables Φ_i .

Formally, a *choice* for an player i is a function $v_i : \Phi_i \rightarrow \{\top, \perp\}$, i.e., an allocation of truth or falsity to all the variables under i 's control. Let \mathcal{V}_i denote the set of choices for player i . If an player i controls no Boolean variables at all, i.e., if $\Phi_i = \emptyset$, we assume that i has only one action, which we denote by \emptyset . A *strategy profile* is a tuple $\vec{v} = (v_1, \dots, v_n)$ in $\mathcal{V}_1 \times \dots \times \mathcal{V}_n$ consisting of exactly one choice for one player. The strategy profile $(v_1, \dots, v_{i-1}, v'_i, v_{i+1}, \dots, v_n)$ we also denote by (\vec{v}_{-i}, v'_i) .

There is a natural correspondence between outcomes and strategy profiles, and we often treat outcomes for Boolean games as valuations, for example writing $(v_1, \dots, v_n) \models \varphi$ to mean that the valuation defined by the outcome (v_1, \dots, v_n) satisfies formula φ . Of course, for partial Boolean games, where $\Phi_U \neq \emptyset$, this correspondence between outcomes and strategy profiles does not hold in general.

e) *Allocations*: Partial games can be made into fully fledged Boolean games by allocating any unallocated Boolean variables to the players. Formally, an *allocation* is a total function $\alpha : \Phi_U \rightarrow N$, with the intended interpretation that, for $p \in \Phi_U$, under allocation α , variable p is allocated to player $\alpha(p)$. With a small abuse of notation, we let α_i denote the set of variables allocated to player i under allocation α , i.e., $\alpha_i = \{p \in \Phi_U : \alpha(p) = i\}$. Let $\mathcal{A}(P)$ denote the set of all allocations over partial game P . Note that $\mathcal{A}(P)$ will have cardinality $|N|^{|\Phi_U|}$, which is exponential in $|\Phi_U|$.

If $\Phi_U = \emptyset$, we will assume that there is a single “empty” allocation possible. Consequently, the set of allocations is non-empty even for Boolean games. A partial game $P = (N, \Phi, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n)$ together with an allocation α defines a Boolean game $G(P, \alpha) = (N, \Phi, \Phi'_1, \dots, \Phi'_n, \gamma_1, \dots, \gamma_n)$ where $\Phi'_i = \Phi_i \cup \alpha_i$ for all players i . We also say that Boolean game $G(P, \alpha)$ *extends* partial game P via α . We let $\mathcal{G}(P)$ denote the set of Boolean games that may be obtained from partial game P through some allocation, i.e., $\mathcal{G}(P) = \{G(P, \alpha) : \alpha \in \mathcal{A}(P)\}$.

²We thus assume that players' preferences are dichotomous, which is traditional in the literature on Boolean games. Several richer models of player preferences have been proposed as well (e.g., [9, 10, 11]). We expect that adopting these models would only slightly affect the computational results presented in this paper.

f) Nash Equilibrium: The well-known notion of (pure) Nash equilibrium (see, e.g., [12]) is readily defined for Boolean games. For Boolean games, we say a strategy profile $\vec{v} = (v_1, \dots, v_n)$ is a Nash equilibrium if there is no player i and no strategy $v'_i \in \mathcal{V}_i$ for i such that $(\vec{v}_{-i}, v'_i) \succ_i (\vec{v}_{-i}, v_i)$. We denote the Nash equilibrium outcomes of a Boolean game G by $\mathcal{N}(G)$. As we are dealing with pure, i.e., non-randomised, strategies, it can very well be that a (partial) Boolean game has no equilibria, that is, it could be that $\mathcal{N}(G) = \emptyset$ for a given game G . If $\vec{v} = (v_1, \dots, v_n)$ is a Nash equilibrium and v the corresponding valuation, then we also say that v is *sustained by a Nash equilibrium*, or, with some abuse of terminology, also that the valuation v is a Nash equilibrium. Moreover, let $\mathcal{S}(P)$ denote the set of allocations over partial game P such that under these allocations, the resulting game has a Nash equilibrium, i.e.,

$$\mathcal{S}(P) = \{\alpha \in \mathcal{A}(P) : \mathcal{N}(G(P, \alpha)) \neq \emptyset\}.$$

As we argued above, the natural correspondence between strategy profiles and valuations does not hold for partial Boolean games. With preferences being defined over valuations, this renders the definition of a Nash equilibria in partial games as a property of strategy profiles problematic. Nevertheless, we can still define an outcome or valuation $v : \Phi \rightarrow \{\top, \perp\}$ to be *sustained by a Nash equilibrium* in a partial game P if there is no player i and no choice $v'_i \in \mathcal{V}_i$ with $(v_{-i}, v'_i) \succ_i v$, where (v_{-i}, v'_i) is the valuation v'' such that, for all $p \in \Phi$,

$$v''(p) = \begin{cases} v'_i(p) & \text{if } p \in \Phi_i, \\ v(p) & \text{otherwise.} \end{cases}$$

In an effort to avoid convoluted formulations, we also say, with a further abuse of terminology, that in such a case the valuation v itself is a Nash equilibrium of the partial game P . We now have the following simple but useful proposition, which intuitively says that allocating variables to players never results in new equilibria emerging. Rather, the outcomes sustained by an equilibrium in a Boolean game are a subset of the outcomes sustained by a Nash equilibrium of any underlying partial game.

Proposition 1: Let $P = (N, \Phi, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n)$ be a partial Boolean game, and $\alpha : \Phi_U \rightarrow N$ an allocation. Then, every outcome that is sustained by a Nash equilibrium of $G(P, \alpha)$ is sustained by a Nash equilibrium of P .

Proof: Let v be an outcome and assume for contraposition that v is not sustained by a Nash equilibrium. Then, there is some player i and strategy v'_i such that $(v_{-i}, v'_i) \succ_i v$. In that case, $v'_i \subseteq \Phi_i$ and hence also $v'_i \subseteq \Phi_i \cup \alpha_i$. But then we may also conclude that v is not sustained by a Nash equilibrium in $G(P, \alpha)$ either. \square

g) Computational Complexity: Although largely self-contained, our technical presentation is necessarily terse, and readers may find it useful to have some acquaintance with the theory of computational complexity (see, e.g., [13, 14]). Throughout the paper, we assume familiarity with the classes P, NP, coNP, and, more generally, the polynomial hierarchy,

i.e., the classes $\Delta_0^P = \Sigma_0^P = \Pi_0^P = P$ and $\Delta_{k+1}^P = P^{\Sigma_k^P}$, $\Sigma_{k+1}^P = NP^{\Sigma_k^P}$, and $\Pi_{k+1}^P = coNP^{\Pi_k^P}$, for $k \geq 0$.

III. THE PRINCIPAL DELEGATION PROBLEM

We now come to the first main problem we consider in this paper. We start with a partial game $P = (N, \Phi, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n)$, with associated unallocated variable set Φ_U . The set Φ_U will represent the decisions that are to be delegated to players in the game. The allocation of Φ_U to players is done by a *principal*, who has complete freedom to allocate the variables Φ_U to the players in N .³

Once an allocation is made, the partial game becomes a Boolean game, and the players will then make rational choices, resulting in some outcome. Recall that the way in which the variables can be allocated to the players may affect the *strategic* structure of the resulting Boolean game, in particular which outcomes are sustained by Nash equilibria.

Now, we assume that the principal will in fact make the allocation with a particular *objective* in mind, which we represent by a Boolean formula Υ . The idea is that the principal will try to choose an allocation so that, if the players then play the resulting Boolean game rationally, they will choose an outcome satisfying Υ .

Following [5]—who study the implementation of a principals' goal in Boolean games with taxes rather than through delegation—we will study two variations of the delegation problem, which we refer to as *weak* and *strong*. In the weak variation, the principal's objective Υ is required to be satisfied in *some* Nash equilibrium of the resulting Boolean game, while in the strong variation, Υ is required to be satisfied in *all* Nash equilibria.

A. Weak Principal Delegation

Formally, the WEAK PRINCIPAL DELEGATION problem is defined as follows:

WEAK PRINCIPAL DELEGATION

Given: A partial game P with Φ_U as the set of unallocated variables, and an objective Υ in $L(\Phi)$

Problem: Does there exist an allocation $\alpha : \Phi_U \rightarrow N$ such that Υ is satisfied in at least one Nash equilibrium of the Boolean game $G(P, \alpha)$?

We say this problem is “weak” because we only require that Υ is satisfied in *one* Nash equilibrium of $G(P, \alpha)$. We will consider stronger versions below. Notice that WEAK DELEGATION is equivalent to checking the following condition:

$$\exists G \in \mathcal{G}(P) : \exists \vec{v} \in \mathcal{N}(G) : \vec{v} \models \Upsilon.$$

The outermost existential quantifier emphasises that the task of the principal can be understood as *choosing a game* from the space of possible games $\mathcal{G}(P)$.

If (P, Υ) is a positive instance of the WEAK DELEGATION problem, then we say that Υ *can be weakly implemented* in P . Again following [5], we refer to weakly implementing

³ For both the players and the principal, we assume the same epistemic preconditions that hold for Nash equilibrium. In particular, we assume principal and the players have common knowledge over all players's goals.

alloc.	weak	strong	alloc.	weak	strong	alloc.	weak	strong
$\alpha^{tca }$	+	—	$\alpha^{a tc }$	+	+	$\alpha^{c ta}$	+	+
$\alpha^{tca }$	+	—	$\alpha^{t ca}$	+	+	$\alpha^{t ca}$	+	+
$\alpha^{tca }$	+	+	$\alpha^{c ta }$	+	+	$\alpha^{t ca}$	+	+
$\alpha^{tca }$	+	—	$\alpha^{t ca }$	+	—	$\alpha^{t ca }$	+	+
$\alpha^{tca }$	+	+	$\alpha^{c ta }$	+	+	$\alpha^{t ca }$	+	+
$\alpha^{tca }$	+	—	$\alpha^{c ta }$	+	—	$\alpha^{c ta }$	—	—
$\alpha^{tca }$	+	—	$\alpha^{a tc }$	+	+	$\alpha^{c ta }$	+	+
$\alpha^{tca }$	+	+	$\alpha^{a tc }$	+	+	$\alpha^{a tc }$	+	+
$\alpha^{tca }$	+	—	$\alpha^{c ta }$	+	+	$\alpha^{a tc }$	+	—

TABLE I
ALLOCATIONS IN EXAMPLE 1 TOGETHER WITH THEIR WEAK AND STRONG IMPLEMENTABILITY.

a tautology—i.e., implementing Υ where $\Upsilon \equiv \top$ —as *stabilisation*. The rationale for this terminology is that weakly implementing a tautology will result in a game that has at least one Nash equilibrium. It is easy to see that \top can be weakly implemented in P if and only if there is at least one allocation α such that $G(P, \alpha)$ allows for a Nash equilibrium. To illustrate weak delegation, we recall Example 1.

Example 4: The prime minister’s predicament in the Introduction can be modelled as a partial Boolean game $P = (\{1, 2, 3\}, \Phi_U, \Phi_1, \Phi_2, \Phi_3, \gamma_1, \gamma_2, \gamma_3)$ with the environmentalist, the libertarian, and the populist being players 1, 2, and 3, respectively, $\Phi_U = \{t, c, a\}$, $\Phi_1 = \Phi_2 = \Phi_3 = \emptyset$, and

$$\gamma_1 = tca \vee (a \leftrightarrow (t \leftrightarrow \bar{c})) \quad \gamma_2 = t \leftrightarrow (\bar{a} \wedge \bar{c}) \quad \gamma_3 = \bar{a} \leftrightarrow (t \wedge c).$$

There are $3^3 = 27$ allocations in total, each of which we denote by $\alpha^{X|Y|Z}$ where X , Y , and Z are understood as the variables assigned to 1, 2, and 3, respectively. Thus, e.g., for $\alpha^{a|tc}$ we have that $\alpha_1^{a|tc} = \{a\}$, $\alpha_2^{a|tc} = \emptyset$, and $\alpha_3^{a|tc} = \{t, c\}$. The four Boolean games depicted in Fig. 2 are obtained by combining P with the allocations $\alpha^{a|tc}$, $\alpha^{c|ta}$, $\alpha^{t|ca}$, and $\alpha^{a|tc|}$.

The prime minister’s (principal’s) objective—majority support for the energy policy to be followed—is given by $\Upsilon_{PM} = (c \leftrightarrow a) \vee \bar{t}c\bar{a}$. Inspecting Fig. 2, we thus find that assignments $\alpha^{a|tc}$, $\alpha^{c|ta}$, and $\alpha^{a|tc|}$ all weakly implement the prime minister’s objective, be it that the witnessing equilibrium for $\alpha^{a|tc}$ is a future with solar and wind energy, whereas for $\alpha^{c|ta}$ and $\alpha^{a|tc|}$ this is a sophisticated green economy in trade-block A. Also observe that $\alpha^{a|tc|}$ also allows for an equilibrium that does not meet the PM’s objective. By contrast, $\alpha^{c|ta|}$ does not weakly implement Υ_{PM} , as it does not allow for any equilibria at all. Table I tabulates all allocations in which Υ_{PM} can be weakly implemented. \triangleleft

Since we have a domain with Boolean formulae, and there are clearly exponentially many allocations of variables to players, it comes as no surprise that the WEAK DELEGATION problem is computationally hard. However, the good news is that it is no harder than the problem of determining the existence of pure strategy Nash equilibria in Boolean games, as we now show.

Theorem 2: WEAK PRINCIPAL DELEGATION is Σ_2^P -complete.

Proof: Allocations α are clearly small with respect to the

size of the partial game, and verifying that outcomes are Nash equilibria in Boolean games is in coNP (cf., [3]). Recall that $\Sigma_2^P = \text{NP}^{\text{coNP}}$. It then follows that the WEAK DELEGATION problem is in Σ_2^P : guess an allocation α and an outcome \vec{v} , and verify both that $\vec{v} \models \Upsilon$ and that \vec{v} is a Nash equilibrium in Boolean game $G(P, \alpha)$.

For hardness, we reduce the problem of checking whether a Boolean game has any pure strategy Nash equilibria, which is known to be Σ_2^P -complete (see, [3]). Given a Boolean game G that we wish to check for the existence pure strategy Nash equilibria simply define the corresponding partial game P to be game G —so that $\Phi_U = \emptyset$ —with objective $\Upsilon = \top$. Notice that the only possible allocation is the empty allocation, which defines the identity under the function $G(\dots)$. Now consider the WEAK DELEGATION problem:

$$\exists \alpha \in \mathcal{A}(P) : \exists \vec{v} \in \mathcal{N}(G(P, \alpha)) : \vec{v} \models \Upsilon.$$

Since $\Upsilon = \top$ and the only allocation possible is the empty allocation, this reduces to: $\exists \vec{v} \in \mathcal{N}(G(P, \alpha)) : \vec{v} \models \top$. Since the empty allocation is the identity under $G(\dots)$, and $\vec{v} \models \top$ for all \vec{v} , this further reduces to $\exists \vec{v} \in \mathcal{N}(G)$, which is exactly the problem of checking for the existence of pure strategy Nash equilibria in Boolean games. \square

B. Strong Principal Delegation

The strong principal delegation problem differs from the weak version in that it requires the objective Υ to be satisfied in *all* Nash equilibria of the Boolean game that results from an allocation. Formally, we define STRONG DELEGATION as follows:

STRONG PRINCIPAL DELEGATION

Given: A partial game P with Φ_U as the set of unallocated variables, and an objective Υ in $L(\Phi)$

Problem: Does there exist an allocation α such that both:

- (i) $G(P, \alpha)$ has at least one Nash equilibrium,
- (ii) all Nash equilibria of $G(P, \alpha)$ satisfy Υ ?

To see how the strong delegation problem differs from the weak delegation problem, we consider again the partial game of Example 1 from the Introduction.

Example 5: In Table I we summarise which allocations strongly implement the PM’s objective Υ_{PM} , and which ones do not. Strong implementability requires the existence of a Nash equilibrium after the principal’s allocating the variables (cf., condition (ii) in STRONG PRINCIPAL DELEGATION). This is why Υ_{PM} is not implemented by $\alpha^{c|ta|}$, the only allocation that does not allow any equilibria. By contrast, the reason why Υ_{PM} is not strongly implemented by $\alpha^{a|tc|}$ either is that the existence of the equilibrium $a\bar{c}t$ (low emission plants), which is supported by one minister only. Interestingly, there is no allocation for this partial game under which *all* equilibria only have minority support even though equilibria exist. \triangleleft

Before appraising the computational complexity of strong principal delegation, it is useful to introduce two related auxiliary decision problems, OBLIVIOUS PARTIAL GAME and STRONG PRINCIPAL DELEGATION- \emptyset . Before we do so, we first recall the following result from [3], where a strategy v_i for

player i is understood to be a *winning strategy*, if the (partial) valuation defined by v_i is completely under i 's control and, irrespective of the choices made by other player or players, v_i ensures that i 's goal is satisfied.

Proposition 3 (Bonzon *et al.*, [3]): Consider a two-player zero-sum⁴ Boolean game $G = (\{1, 2\}, \Phi, \Phi_1, \Phi_2, \gamma_1, \gamma_2)$. Then, $\vec{v} = (v_1, v_2)$ is a pure strategy Nash equilibrium for G if and only if v_1 is a winning strategy for player 1 or v_2 is a winning strategy for player 2.

Now we introduce *oblivious partial games* as an auxiliary concept as follows. Given a principal's objective Υ , we say that partial Boolean game P with unallocated variables Φ_U is *oblivious*, if for every allocation α , the resulting Boolean game $G(P, \alpha)$ has *at least one* Nash equilibrium satisfying Υ , i.e., if for every $\alpha \in \mathcal{A}(P)$, there is at least one $\vec{v} \in \mathcal{N}(G(P, \alpha))$ with $\vec{v} \models \Upsilon$. Informally, oblivious games have the property that, whichever delegation of decisions to players is made, there will be at least one acceptable valuation available, i.e., which is both an equilibrium and satisfies Υ .

We now consider the following decision problem OBLIVIOUS PARTIAL GAME.

OBLIVIOUS PARTIAL GAME

Given: A partial game P with Φ_U as the set of unallocated variables, and an objective Υ in $L(\Phi)$

Problem: Is P an oblivious partial game given Υ ?

We find that OBLIVIOUS PARTIAL GAME is Π_3^P -complete.

Proposition 4: OBLIVIOUS PARTIAL GAME is Π_3^P -complete. The problem remains Π_3^P -hard even when restricted to three-player games.

Proof: First observe that an instance (P, Υ) of OBLIVIOUS PARTIAL GAME is accepted if and only if

$$\forall \alpha \in \mathcal{A}(P) \exists \vec{v} \in \mathcal{V}(\Phi) \text{ such that } \vec{v} \in \mathcal{N}(G(P, \alpha)) \text{ and } \vec{v} \models \Upsilon$$

This computation required to check this condition can easily be carried out within Π_3^P .

For hardness we use a variant of the well-known quantified Boolean formula validity problem QBF_{3,∇}, where instances of the form $\psi(X, Y, Z)$ are accepted if for all assignments $x: X \rightarrow \{\top, \perp\}$ of values to X , there is an assignment $y: Y \rightarrow \{\top, \perp\}$ of values to Y such that for every assignment $z: Z \rightarrow \{\top, \perp\}$ of values to Z we have $(x, y, z) \models \psi$.

The variation we use, and which we call HALF-QBF_{3,∇}, restricts attention to valuations $x: X \rightarrow \{\top, \perp\}$ that set exactly half of the variables in X to \top , i.e., an instance $\psi(X, Y, Z)$ of HALF-QBF_{3,∇} is accepted if and only if for all $x: X \rightarrow \{\top, \perp\}$ with $|\{p \in X : x(p) = \top\}| = |X|/2$, there is a $y: Y \rightarrow \{\top, \perp\}$ such that for all $z: Z \rightarrow \{\top, \perp\}$:

$$(x, y, z) \models \psi(X, Y, Z).$$

We can straightforwardly show that, modified thus, the problem remains Π_3^P -hard, for instance, by using the ‘padding argument’ methods in [15] (in particular, Lemma 3.7).

Given an instance $\psi(X, Y, Z)$ of HALF-QBF_{3,∇}, we construct the three-player partial Boolean game P_ψ , with play-

ers 1, 2, and 3, and

$$\Phi_1 = Y \cup \{r\} \quad \Phi_2 = Z \cup \{s\} \quad \Phi_3 = \emptyset \quad \Phi_U = X.$$

Furthermore, define the players' preferences as follows:

$$\begin{aligned} \gamma_1 &= (EQ(X) \rightarrow \psi) \vee (r \leftrightarrow s) \\ \gamma_2 &= (EQ(X) \wedge \neg\psi \wedge (r \leftrightarrow \neg s)) \vee \neg MAJ(X) \\ \gamma_3 &= EQ(X) \end{aligned}$$

Here r and s are fresh variables. Moreover, $EQ(X)$ denotes the propositional function that evaluates to true if and only if the assignment sets *exactly* half of the variables in X to \top . Similarly, $MAJ(X)$ is the propositional function that evaluates to true if and only if the assignment sets *at least half* of X to \top . Note that the propositional functions $EQ(X)$ and $MAJ(X)$ can be encoded in polynomial size formulae (see, e.g., [16]). Accordingly, the construction of G_φ can be effected in polynomial time. Finally, set $\Upsilon = \top$.

As clearly every valuation thus satisfies Υ , it now suffices to show that $\psi(X, Y, Z)$ is an accepting instance of HALF-QBF_{3,∇} if and only if for every allocation $\alpha: X \rightarrow \{1, 2\}$, the game $G(P_\psi, \alpha)$ has at least one Nash equilibrium.

First assume that $\psi(X, Y, Z)$ is an accepting instance of HALF-QBF_{3,∇} and consider an arbitrary allocation of $\alpha: \Phi_U \rightarrow \{1, 2, 3\}$ of X . We show that there is a Nash equilibrium in $G(P_\psi, \alpha)$. We distinguish two cases: (i) $|\alpha_1 \cup \alpha_3| < |X|/2$, and (ii) $|\alpha_1 \cup \alpha_3| \geq |X|/2$.

If (i), it follows that $|\alpha_2| > |X|/2$. Let $\vec{v}^* = (v_1^*, v_2^*, v_3^*)$ be any strategy profile such that $v_2^*(p) = \perp$ for all $p \in \Phi_2 \cup \alpha_2$, and $v_1^*(r) = \perp$. Then, $(v_1^*, v_2^*, v_3^*) \models \neg MAJ$ and $(v_1^*, v_2^*) \models r \leftrightarrow s$. Hence, (v_1^*, v_2^*, v_3^*) satisfies both γ_1 and γ_2 , but not γ_3 . Observe, however, that player 3 cannot get their goal satisfied by playing any other strategy, and we may conclude that \vec{v}^* is a Nash equilibrium of $G(P_\psi, \alpha)$.

Now assume (ii). Then there is a $X' \subseteq \alpha_1 \cup \alpha_3$ such that $|X'| = |X|/2$. Let $x^*: X \rightarrow \{\top, \perp\}$ be the partial valuation such that $x^*(p) = \top$ if and only if $p \in X'$. Having assumed that $\psi(X, Y, Z)$ is accepted as an instance of HALF-QBF_{3,∇}, there is some $y^*: Y \rightarrow \{\top, \perp\}$ with $(x^*, y^*, z) \models \psi$ for all $z: Z \rightarrow \{\top, \perp\}$. Define the strategy profile $\vec{v}^* = (v_1^*, v_2^*, v_3^*)$ such that, for all $p \in \Phi_1 \cup \alpha_1$ and $q \in \alpha_3$,

$$\begin{aligned} v_1^*(p) &= \begin{cases} \top & \text{if } p \in X' \text{ and } x^*(p) = \top, \text{ or } p \in Y \text{ and } y^*(p) = \top, \\ \perp & \text{otherwise} \end{cases} \\ v_3^*(q) &= \begin{cases} \top & \text{if } q \in X' \text{ and } x^*(q) = \top, \\ \perp & \text{otherwise.} \end{cases} \end{aligned}$$

Let furthermore $v^*(p) = \perp$ for all $p \in \Phi_2 \cup \alpha_2$. Observe that thus both $\vec{v}^* \models EQ(X)$ and $\vec{v}^* \models \psi$. Accordingly players 1 and 3 have their respective goals γ_1 and γ_3 achieved at \vec{v}^* , and have no incentive to deviate. Finally, to see that player 2 does not have such an incentive either, and thus that \vec{v}^* is a Nash equilibrium, consider an arbitrary strategy v'_2 for player 2. If there is some $p \in \alpha_2 \cap X$ such that $v'_2(p) \neq v_2^*(p)$, then $v'(p) = \top$. Hence neither $(v_1^*, v'_2, v_3^*) \models EQ(X)$ nor $(v_1^*, v'_2, v_3^*) \models \neg MAJ(X)$. Now consider the case where $v'_2(p) = v_2^*(p)$ for all $p \in \alpha_2 \cap X$ and $v'_2(p) \neq v_2^*(p)$ for some $p \in Y \cup \{s\}$. Because $\psi(X, Y, Z)$ is accepted as an

⁴That is to say, for every valuation v , *exactly one* of γ_1 or γ_2 is satisfied.

instance of HALF-QBF_{3,∇}, it then holds that $(v_1^*, v_2', v_3^*) \models \psi$. In either case we may conclude that $(v_1^*, v_2', v_3^*) \not\models \gamma_2$ and that player 2 does not want to deviate.

For the opposite direction, assume for contraposition that $\psi(X, Y, Z)$ is *not* accepted as an instance of HALF-QBF_{3,∇}. Then, there is some $x: X \rightarrow \{\top, \perp\}$ with $x(p) = \top$ for exactly half of the variables $p \in X$ such that for all $y: Y \rightarrow \{\top, \perp\}$ we can find a $z: Z \rightarrow \{\top, \perp\}$ such that $(x, y, z) \models \psi$. Let $X_\top = \{p \in X : x(p) = \top\}$ and $X_\perp = \{p \in X : x(p) = \perp\}$. Now consider the allocation α that assigns control over all variables in X to player 3, i.e., $\alpha_1 = \alpha_2 = \emptyset$ and $\alpha_3 = X$. We show that the Boolean game $G(P_\psi, \alpha)$ does not admit a Nash equilibrium. To this end, consider an arbitrary strategy profile $\vec{v} = (v_1, v_2, v_3)$. First observe that, if $\vec{v}(p) \models EQ(X)$, it is not the case that $v_3(p) = \top$ for exactly half of the variables in X . Hence, player 3 would deviate to such a strategy to get their goal achieved, and \vec{v} is not a Nash equilibrium.

For the remainder of the proof we may therefore assume that $\vec{v} \models EQ(X)$. Now suppose furthermore that $\vec{v} \models \psi$ and, without loss of generality, that $\vec{v} \models r$. Having assumed that $\psi(X, Y, Z)$ is *not* accepted as an instance of HALF-QBF_{3,∇}, it follows that there is a strategy v_2' for player 2 such that $(v_1, v_2', v_3) \models \neg\psi$. We may also stipulate that $v_2'(s) = \perp$, and hence $(v_1, v_2', v_3) \models (r \leftrightarrow \neg s)$. Accordingly, $(v_1, v_2', v_3) \models \gamma_2$ and again it follows that \vec{v} is not a Nash equilibrium.

Finally assume that $\vec{v} \models EQ(X) \wedge \neg\psi$. Then, if $\vec{v} \models r \leftrightarrow s$, it can easily be seen that player 2 would like to deviate by setting s to the opposite value of r under \vec{v} . If, on the other hand, $\vec{v} \not\models r \leftrightarrow s$, player 1 would like to deviate by matching the value of r to that of s under \vec{v} . In either case, \vec{v} is not a Nash equilibrium, which concludes the proof. \square

One consequence of Proposition 4 is that we can now easily establish Σ_3^P -hardness for the variant of the strong delegation problem where we allow $\mathcal{N}(P, \alpha)$ to be empty. This problem we refer to as STRONG PRINCIPAL DELEGATION- \emptyset and is formally defined as follows.

STRONG PRINCIPAL DELEGATION- \emptyset

Given: A partial game P with Φ_U as the set of unlocated variables, and an objective Υ in $L(\Phi)$

Problem: Does there exist an allocation $\alpha: \Phi_U \rightarrow N$ such that Υ is satisfied in all Nash equilibria of the Boolean game $G(P, \alpha)$?

We now have the following intermediate result.

Proposition 5: STRONG PRINCIPAL DELEGATION- \emptyset is Σ_3^P -complete. The problem remains Σ_3^P -hard even when restricted to three-player games.

Proof: Let P be a partial game and Υ a principal's objective. Then, by the laws of first-order logic, we have that there is an allocation $\alpha \in \mathcal{A}(P)$ such that $\vec{v} \models \Upsilon$ for all Nash equilibria \vec{v} of $G(P, \alpha)$ if and only if it is not the case that for all allocations $\alpha \in \mathcal{A}(P)$ there is a Nash equilibrium \vec{v} of $G(P, \alpha)$ with $\vec{v} \models \neg\Upsilon$. Hence, the function f that maps every pair (P, Υ) to $(P, \neg\Upsilon)$ serves as a many-one reduction from STRONG PRINCIPAL DELEGATION- \emptyset to the complement of OBLIVIOUS PARTIAL GAME. Observing that for $\Upsilon' = \neg\Upsilon$ we have $\neg\Upsilon' = \neg\neg\Upsilon \equiv \Upsilon$, we find that it serves equally well

as a many-one reduction from the complement of OBLIVIOUS PARTIAL GAME to STRONG PRINCIPAL DELEGATION- \emptyset . As by Proposition 4 we know that OBLIVIOUS PARTIAL GAME is Π_3^P -complete, we may conclude that STRONG PRINCIPAL DELEGATION- \emptyset is Σ_3^P -complete as well. \square

Although STRONG PRINCIPAL DELEGATION- \emptyset admits problem instances which are accepted in the case that $\mathcal{N}(P, \alpha) = \emptyset$, whereas the form of STRONG PRINCIPAL IMPLEMENTATION does not allow this, it is not difficult to show that the additional constraint cannot make the strong delegation problem increase in complexity. Specifically, we are now in a position to prove the main result of this section.

Theorem 6: STRONG PRINCIPAL IMPLEMENTATION is Σ_3^P -complete. The problem remains Σ_3^P -hard even for three-player games.

Proof: We recall that instances of STRONG PRINCIPAL IMPLEMENTATION are given by pairs (P, Υ) with $P = (N, \Phi, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n)$ a partial Boolean game and Υ a propositional function. Let $\Phi_U \subseteq \Phi$ the set of unlocated decisions. Then, (P, Υ) is accepted as an instance of STRONG PRINCIPAL IMPLEMENTATION if and only if

$$\exists \alpha (\forall v (v \in \mathcal{N}(P, \alpha) \rightarrow v \models \Upsilon) \wedge \mathcal{N}(P, \alpha) \neq \emptyset),$$

which is equivalent to,

$$\exists \alpha (\forall v (v \notin \mathcal{N}(P, \alpha) \vee v \models \Upsilon) \wedge \mathcal{N}(P, \alpha) \neq \emptyset).$$

In order for $\mathcal{N}(P, \alpha)$ to be non-empty, it must contain at least one valuation w . Hence we can rewrite the expression as,

$$\exists (\alpha, w) (\forall v (v \notin \mathcal{N}(P, \alpha) \vee v \models \Upsilon) \wedge (w \in \mathcal{N}(P, \alpha)))$$

To keep notation brief, we express the condition above as $\exists (\alpha, w) \forall v \chi_1(\alpha, w, v)$. Similarly the condition $v \notin \mathcal{N}(P, \alpha)$ is witnessed by an player i together with a partial valuation u_i of decisions in their control, i.e., $v \notin \mathcal{N}(P, \alpha)$ is captured by,

$$\exists (u, i) (v \not\models \gamma_k \wedge (v_{-i}, u_i) \models \gamma_k \wedge u \in \{\top, \perp\}^{\Phi_i \cup \alpha_i}),$$

which we subsequently express as $\exists (u, k) \chi_2(\alpha, v, u, i)$.

Finally that $w \in \mathcal{N}(P, \alpha)$ requires that no j has a beneficial deviation in the event of $w \not\models \gamma_j$ by controlling the values of $\Phi_i \cup \alpha_i$. This leads to $w \in \mathcal{N}(P, \alpha)$ being of the form

$$\forall (s_1, s_2, \dots, s_n) \bigwedge_{j=1}^n (w \models \gamma_j \vee s_j \notin \{\top, \perp\}^{\Phi_i \cup \alpha_j} \vee (w_{-j}, s_j) \models \gamma_j)$$

which will be denoted as $\forall \vec{s} \chi_3(\vec{s}, w, \alpha)$. Combining these expressions, it is immediate that (P, Υ) will define a positive instance of PRINCIPAL STRONG DELEGATION if and only if

$$\exists (\alpha, w) \forall (v, \vec{s}) \exists (u, i) (\chi_1(\alpha, w, v) \wedge \chi_2(\alpha, v, u, i) \wedge \chi_3(\vec{s}, w, \alpha))$$

It remains to note that the individual tests χ_1 , χ_2 and χ_3 can be performed in time polynomial in the size of the representing formulae.

For Σ_3^P -hardness, we reduce HALF-QBF_{3,∇} to the complement of STRONG PRINCIPAL DELEGATION using a construction that varies on the one presented in the proof for hardness of OBLIVIOUS PARTIAL GAME in Proposition 4. Given an instance $\psi(X, Y, Z)$ of HALF-QBF_{3,∇}, we thus construct a partial game P'_ψ and define an objective Υ' such

that $\psi(X, Y, Z)$ is an accepting instance of $\text{HALF-QBF}_{3,\forall}$ if and only if for all allocations $\alpha \in \mathcal{A}(P'_\psi)$, either there is some $\vec{v} \in \mathcal{N}(G(P'_\psi, \alpha))$ with $\vec{v} \models \Upsilon'$ or $\mathcal{N}(G(P'_\psi, \alpha)) = \emptyset$.

Thus let $\psi(X, Y, Z)$ be an instance of $\text{HALF-QBF}_{3,\forall}$. Then, we construct the five-player partial Boolean game P'_ψ , with players 1, 2, and 3, such that

$$\Phi_1 = Y \cup \{r\} \quad \Phi_2 = Z \cup \{s\} \quad \Phi_3 = \{t\} \quad \Phi_U = X.$$

Furthermore, define the players' preferences as follows:

$$\begin{aligned} \gamma'_1 &= ((EQ(X) \rightarrow \psi) \vee (r \leftrightarrow s)) \wedge \neg t \\ \gamma'_2 &= ((EQ(X) \wedge \neg \psi \wedge (r \leftrightarrow \neg s)) \vee \neg MAJ(X)) \wedge \neg t \\ \gamma'_3 &= EQ(X) \end{aligned}$$

Here r , s , and t are fresh variables, and $MAJ(X)$ and $EQ(X)$ are as before. Let, moreover, $\Upsilon' = t$. Again, this construction can be effected in polynomial time.

Now assuming that $\psi(X, Y, Z)$ is an accepted instance of $\text{HALF-QBF}_{3,\forall}$, we can reason along analogous lines as in the proof of Proposition 4, that a Nash equilibrium \vec{v} of $G(P'_\psi)$ with $\vec{v} \models t$ is guaranteed to exist.

For the opposite direction, assume that $\psi(X, Y, Z)$ is not an accepted instance of $\text{HALF-QBF}_{3,\forall}$. By considering the allocation $\alpha \in \mathcal{A}(P'_\psi)$ with $\alpha_3 = X$, we find, in an analogous way as in the proof of Proposition 4, that $G(P'_\psi, \alpha)$ has no Nash equilibrium \vec{v} with $\vec{v} \models t$. Now, conclude the proof by noting that any strategy profile $\vec{w} = (w_1, w_2, w_3)$ such that $w_3(t) = \top$ and $w_3(p) = \top$ for exactly half of the propositional variables $p \in X$, will be a Nash equilibrium in $G(P'_\psi, \alpha)$, be it one with $\vec{w} \models t$. To see this, observe that in \vec{w} , player 3 has their goal fulfilled. This is not the case for players 1 and 2, but no matter how they deviate, t will remain true and falsify their respective goals γ'_1 and γ'_2 . Hence, $\vec{w} \in \mathcal{N}(G(P'_\psi, \alpha))$, and *a fortiori* $\mathcal{N}(G(P'_\psi, \alpha)) \neq \emptyset$, as desired. \square

IV. DELEGATION AS OPTIMISATION

So far, we have assumed that the principal is motivated to choose an allocation so that an objective Υ is satisfied in equilibrium; the idea being that the objective represents what the principal wants to achieve through delegation. We now generalise this approach, by assuming that in delegating decisions, the principal is attempting to maximise some objective function of the form:

$$f : \mathcal{V}(\Phi) \rightarrow \mathbb{R}_+,$$

where $\mathcal{V}(\Phi)$ is the set of valuations over Φ . Thus, an objective function assigns a positive real number $f(v)$ to every valuation v in $\mathcal{V}(\Phi)$. Recalling that a strategy profile \vec{v} for a Boolean game corresponds to a valuation in $\mathcal{V}(\Phi)$, we will also write $f(\vec{v})$ to mean the value through f of the valuation corresponding to \vec{v} . Notice that our original formulation of delegation with respect to objective formulae Υ is a special case of the setting we are now considering, where the function f_Υ is defined for an objective formula $\Upsilon \in L$ as follows:

$$f_\Upsilon(\vec{v}) = \begin{cases} 1 & \text{if } \vec{v} \models \Upsilon \\ 0 & \text{otherwise.} \end{cases}$$

The objective function f gives the value to the principal of every outcome \vec{v} . But how can we use f to obtain the value of an allocation? An allocation α for a partial game P will define a Boolean game $G(P, \alpha)$, and this game will in turn have an associated set of Nash equilibria $\mathcal{N}(G(P, \alpha))$. Our basic idea is to define the value of an allocation α for a partial game P through an objective function f to be the value of the *worst* Nash equilibrium in $\mathcal{N}(G(P, \alpha))$ (cf., [17]). However, there is a catch: what happens if $\mathcal{N}(G(P, \alpha)) = \emptyset$? In this case, we say the value of α through f is undefined. Formally, given a partial game P , allocation $\alpha \in \mathcal{A}(P)$, and objective function $f : \mathcal{V}(\Phi) \rightarrow \mathbb{R}_+$, we denote the value of α through f as $\hat{f}(P, \alpha)$:

$$\hat{f}(P, \alpha) = \begin{cases} \min\{f(\vec{v}) : \vec{v} \in \mathcal{N}(G(P, \alpha))\} & \text{if } \mathcal{N}(G(P, \alpha)) \neq \emptyset, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Now, given a partial game P and an objective function f as above, the *optimal allocation* will intuitively be the one that maximises the value of \hat{f} . Here, however, we must deal with the situation where there is no allocation that leads to a game with Nash equilibria. Recalling that $\mathcal{S}(P) = \{\alpha \in \mathcal{A}(P) : \mathcal{N}(G(P, \alpha)) \neq \emptyset\}$, we are thus interested in allocations that maximise the function $\hat{f}(\cdot, \cdot)$ on $\mathcal{S}(P)$, i.e., allocations α such that

$$\alpha \in \arg \max_{\alpha \in \mathcal{S}(P)} \hat{f}(P, \alpha).$$

To fix ideas and notations, we have the following example.

Example 6: We work with the partial game introduced in Example 1. Consider the objective function f defined such that $f(v) = |\{x \in \Phi : v \models x\}|$, i.e., f counts the number of variables assigned \top in a valuation. Intuitively, f could be seen as measuring the level of effort needed to get the government's policy through parliament. Hence,

$$\begin{aligned} f(\bar{t}\bar{c}\bar{a}) &= 0 & f(\bar{t}\bar{c}a) &= f(\bar{t}c\bar{a}) = 1 & f_1(\bar{t}ca) &= 2 \\ f(t\bar{c}\bar{a}) &= 1 & f(t\bar{c}a) &= f(tc\bar{a}) = 2 & f_1(tca) &= 3 \end{aligned}$$

Inspecting Fig. 2, we find that $\hat{f}(P, \alpha^{a|t|c}) = 2$ because $a\bar{t}\bar{c}$ is the unique Nash equilibrium in $G(P, \alpha^{a|t|c})$ and $f(a\bar{t}\bar{c}) = 2$. By similar reasoning, we find that $\hat{f}(P, \alpha^{c|a|t}) = 3$. As f assumes a maximal value of 3, it thus follows that $\alpha^{c|a|t}$ is optimal with respect to f . Observe that $\hat{f}(P, \alpha^{t|c|a})$ is undefined as there are no equilibria in $G(P, \alpha^{t|c|a})$. \triangleleft

Next, we want to consider the problem of computing $\arg \max_{\alpha \in \mathcal{S}(P)} \hat{f}(P, \alpha)$. A key difficulty here is with respect to the issue of representing the objective function f . Representing the function f in a problem instance by explicitly listing all input/output pairs $(\vec{v}, f(\vec{v}))$ will not be practicable, as there will be $2^{|\Phi|}$ such pairs in total. We need a *compact* representation for f , and for the purposes of this paper, we use a well-known scheme based on weighted Boolean formulae (see, e.g., [9]).

Formally, we will say a *feature* is a pair (φ, x) , where $\varphi \in L$ is a propositional formula, and $x \in \mathbb{R}_+$ is a positive real number. A *feature set*, \mathcal{F} , is simply a finite set of features, i.e., $\mathcal{F} = \{(\varphi_1, x_1), \dots, (\varphi_k, x_k)\}$. Every feature set \mathcal{F} induces an objective function $f_{\mathcal{F}}$, as follows:

$$f_{\mathcal{F}}(\vec{v}) = \sum \{x_i : (\varphi_i, x_i) \in \mathcal{F} \text{ and } \vec{v} \models \varphi_i\}$$

The feature set corresponding to our original objective formula Υ would be a singleton set $\{(\Upsilon, 1)\}$. Now, standard arguments from Boolean function theory tell us that (i) the feature set representation is *complete*, in the sense that for every objective function f there exists a feature set \mathcal{F} such that $f = f_{\mathcal{F}}$, and (ii) the feature set representation is *more compact* than the explicit representation for many objective functions f ; however (iii) there are objective functions for which the smallest equivalent feature set will be broadly of the same size, i.e., some objective functions will require exponentially many features.

We first consider the following decision variant of the optimal delegation problem.

OPTIMAL DELEGATION

Given: A partial game P , feature set \mathcal{F} , and $k \in \mathbb{R}_+$

Problem: If the term $\max_{\alpha \in \mathcal{S}(P)} \hat{f}_{\mathcal{F}}(P, \alpha)$ is defined, is it the case that $\max_{\alpha \in \mathcal{S}(P)} \hat{f}_{\mathcal{F}}(P, \alpha) \geq k$?

It is straightforward to establish the following:

Proposition 7: OPTIMAL DELEGATION is Σ_2^P -complete.

Proof sketch: Let partial game P , feature set \mathcal{F} , and allocation α be given. Then, the problem of determining whether $\hat{f}_{\mathcal{F}}(P, \alpha)$ is defined is Σ_2^P -complete. To see this, merely observe that this problem is directly equivalent to determining whether the game $G(p, \alpha)$ has a Nash equilibrium. At this point, notice that the problem of deciding whether $\max_{\alpha \in \mathcal{S}(P)} \hat{f}_{\mathcal{F}}(P, \alpha)$ is equivalent to checking whether $\mathcal{S}(P) \neq \emptyset$, which is Σ_2^P -complete by the proof of Theorem 2. The claim then follows immediately from these two observations. \square

With this result in place, we can also address the computational complexity of the following function problem.

OPTIMAL DELEGATION*

Given: A partial game P and feature set \mathcal{F}

Provide: An allocation α that maximises the function $\hat{f}_{\mathcal{F}}(P, \cdot)$, i.e., $\alpha \in \arg \max_{\alpha \in \mathcal{S}(P)} \hat{f}_{\mathcal{F}}(P, \alpha)$, if $\mathcal{S}(P) \neq \emptyset$, and \emptyset otherwise.

We can now state the main result of this section.

Theorem 8: OPTIMAL DELEGATION* is in $\text{FP}^{\Sigma_2^P}$.

Proof: Check whether $\alpha^*(P, f_{\mathcal{F}})$ is defined, and if not return 0. Otherwise, define a value μ as follows:

$$\mu = \sum_{(\varphi_i, x_i) \in \mathcal{F}} x_i$$

That is, μ is the largest value that an optimal allocation could possibly take. We thus have:

$$0 \leq \max_{\alpha \in \mathcal{S}(P)} \hat{f}_{\mathcal{F}}(P, \alpha) \leq \mu.$$

We can then use binary search to find the value $\hat{f}_{\mathcal{F}}(P, \alpha^*)$ of an optimal allocation α^* maximising $\hat{f}_{\mathcal{F}}$, by invoking a Σ_2^P -oracle for the decision variant OPTIMAL DELEGATION of the problem. We start by asking whether $\max_{\alpha \in \mathcal{S}(P)} \hat{f}_{\mathcal{F}}(P, \alpha) \geq \frac{1}{2}\mu$; if the answer is “no”, then we

ask whether $\max_{\alpha \in \mathcal{S}(P)} \hat{f}_{\mathcal{F}}(P, \alpha) \geq \frac{3}{4}\mu$, while if the answer is “yes”, we ask whether $\max_{\alpha \in \mathcal{S}(P)} \hat{f}_{\mathcal{F}}(P, \alpha) \geq \frac{3}{4}\mu$, and so on. We will converge to the value of the optimal allocation with at most polynomially many queries to a Σ_2^P -oracle for the decision variant OPTIMAL DELEGATION of the problem (cf., [13], p. 416). Given the value of the optimal allocation, we can then find an optimal allocation with at most a further $|N \times \Phi|$ queries to a Σ_2^P -oracle. \square

V. THE DISTRIBUTED DELEGATION PROBLEM

Given a partial Boolean game, it may very well happen that for some player the Nash equilibria under one allocation yields them a higher payoff than the Nash equilibrium that arises under another one. Thus, intuitively, the players’ preferences over outcomes induce preferences over allocations. As players may have a joint interest in some allocation being implemented rather than another, they may benefit from reallocating the propositional variables amongst themselves. Accordingly, the game-theoretic concept of *core stability*, in the sense of stability against coalition deviations, becomes relevant.

It is worth observing that in Example 3 under both α and α' , player 3 is assigned control over the same propositional variable r . Still, player 1’s goal is satisfied in the Nash equilibrium that emerges under α' , but not in the one under α . Thus, their preferences over allocations are not only dependent on the variables assigned to her control, but also on the way control over the remaining variables is distributed over the other players. The allocation of a particular set of propositional variables to one player can thus be said to impose *externalities* on the other players. In this respect, our setting differs from many other allocation settings studied in the literature.

In order to systematically extend the players’ preferences over outcomes to preferences over allocations, a couple of points deserve attention.

First, whether a group of players would prefer a reallocation α' of the variables they have under their control under an allocation α depends essentially on the outcomes that ensue under α and α' . In the previous parts we used the concept of Nash equilibrium to single these out. The allocation problem, however, is in principle independent of this choice for Nash equilibrium, in the sense that every other way to select outcomes under specific allocations defines its own set of delegation problems. Second, Nash equilibria in Boolean games, even when all variables are allocated, are guaranteed neither to exist nor to be unique. In view of these concerns, we augment the model with an *outcome function*, which associates with each allocation of a given partial game a *unique* outcome. This enables us to approach the (distributed) delegation problem from a generic point of view, while still allowing us concentrate on more restrictive outcome functions that select Nash equilibria whenever they exist. In the latter case, the outcome function basically plays the role of an oracle breaking ties among the equilibria. If under some allocation no Nash equilibria exist, we stipulate the outcome function to select an outcome that is least preferred by all players.

Formally, an *outcome function* for a partial Boolean game P is a function

$$h: \mathcal{A}(P) \rightarrow \mathcal{V}(\Phi) \cup \{\mathbf{0}\},$$

	q	\bar{q}	q	\bar{q}		r	\bar{r}	r	\bar{r}
p	<u>0, 0, 0</u>	<u>1, 0, 1</u>	<u>1, 1, 0</u>	<u>0, 1, 1</u>	p	<u>0, 0, 0</u>	<u>1, 1, 0</u>	<u>1, 0, 1</u>	<u>0, 1, 1</u>
\bar{p}	<u>0, 1, 1</u>	<u>1, 1, 0</u>	<u>1, 0, 1</u>	<u>0, 0, 0</u>	\bar{p}	<u>0, 1, 1</u>	<u>1, 0, 1</u>	<u>1, 1, 0</u>	<u>0, 0, 0</u>
	r	\bar{r}	q	\bar{q}		q	\bar{q}	pr	$\bar{p}\bar{r}$
(a) Allocation $\alpha^a = \alpha^{p q r}$					(b) Allocation $\alpha^b = \alpha^{p r q}$				
	r	\bar{r}	r	\bar{r}		r	\bar{r}	pr	$\bar{p}\bar{r}$
q	<u>0, 0, 0</u>	<u>1, 1, 0</u>	<u>0, 1, 1</u>	<u>1, 0, 1</u>	q	<u>0, 0, 0</u>	<u>1, 1, 0</u>	<u>0, 1, 1</u>	<u>1, 0, 1</u>
\bar{q}	<u>1, 0, 1</u>	<u>0, 1, 1</u>	<u>1, 1, 0</u>	<u>0, 0, 0</u>	\bar{q}	<u>1, 0, 1</u>	<u>0, 1, 1</u>	<u>1, 1, 0</u>	<u>0, 0, 0</u>
	p	\bar{p}	p	\bar{p}		p	\bar{p}	pr	$\bar{p}\bar{r}$
(c) Allocation $\alpha^c = \alpha^{q r p}$					(d) Allocation $\alpha^d = \alpha^{q pr }$				

Fig. 4. The partial Boolean game described in Example 7 under four different allocations. Robber 1 chooses rows, robber 2 chooses columns, and robber 3 chooses matrices. In the pure Nash equilibria are indicated by the payoffs to the players being underlined.

where $\mathbf{0}$ is an additional *null outcome* such that $u_i(\mathbf{0}) = 0$ for all players i . We say an outcome function h for a partial Boolean game is *Nash-consistent* if $h(\alpha)$ selects a Nash equilibrium of $G(P, \alpha)$, if there is one, and $\mathbf{0}$, otherwise.

We are now able to formally define the notions of stability that we have sketched out above. Thus, we say that, given outcome function h for a partial Boolean game P and allocation α , a coalition S of players (*strongly*) *blocks* α under h if there is some allocation α' with $\alpha_{-S} = \alpha'_{-S}$ such that $h(\alpha') \succ_i h(\alpha)$ for all players $i \in S$. We also say that S *weakly blocks* α under h if both $h(\alpha') \succeq_i h(\alpha)$ for all $i \in S$, and $h(\alpha') \succ_i h(\alpha)$ for some $i \in S$. Intuitively, S blocks an allocation α if the players in S can reallocate the propositional variables under their control among themselves such that h yields a outcome that better for them all. Coalition S weakly blocks α , if the players in S can reallocate the propositional variables under their control among themselves such that h yields a outcome that is not worse for any of the players in S and strictly better for some.

An allocation α is now said to be (*weak*) *core stable under outcome function* h whenever there is no coalition (*strongly*) blocking α under h . Allocation α is *strong core stable*, if no coalition S weakly blocks it. To illustrate these notions, let us consider an example.

Example 7: Three robbers have to cross a forest, through which lead two roads, the high road and the low road. The robbers, however, are quarrelsome and the forest is infested with wolves, who will eat any solitary robber. If all three robbers go one road together, they will quarrel, and subsequently be mangled by the wolves. Any group of two, however, will survive. They now have to decide whether to take the high road or the low road, or they can have one of the others make the decision for them. The question is now, who has to decide for whom which road to take? Obviously, at most two of the robbers can survive and we assume that surviving is the only thing they are interested in.

The situation can be modelled as a partial Boolean game, where p , q , and r are the decision variables for which road

robber 1, robber 2, and robber 3 takes: if a variable is set to true, then the high road is taken, else the low road. For allocations, we adopt the same conventions as in Example 4 and have $\alpha^{X|Y|Z}$ denote the allocation in which robbers 1, 2, and 3 are assigned the variables in X , Y , and Z , respectively. Fig. 4 depicts the Boolean games under allocations $\alpha^a = \alpha^{p|q|r}$, $\alpha^b = \alpha^{p|r|q}$, $\alpha^c = \alpha^{q|r|p}$, and $\alpha^d = \alpha^{q|pr|}$.

First consider the outcome function h such that,

$$h(\alpha) = \begin{cases} pq\bar{r} & \text{if } \alpha(p) = 2, \\ \bar{p}qr & \text{if } \alpha(p) = 3, \\ p\bar{q}r & \text{if } \alpha(p) = 1. \end{cases}$$

Then, trivially, *every* allocation is (weakly) core stable, as it can easily appreciated that under h there are always at least two robbers' goals are satisfied at $h(\alpha)$ for every allocation α and at least two dissatisfied robbers are required for an outcome to be strongly blocked. By contrast, the game does not allow for core stable allocations under h . To see this let α be arbitrary allocation. Then, α is blocked by coalition $\{2, 3\}$ if $\alpha(p) = 2$, by $\{1, 3\}$ if $\alpha(p) = 3$, and by $\{1, 2\}$, if $\alpha(p) = 1$. Now, observe, however, that h is *not* Nash consistent. For example, for allocation α^b we have $\alpha^b(p) = 1$. Hence $h(\alpha^b) = p\bar{q}r$, which however fails to be a Nash equilibrium of $G(P, \alpha^b)$ (see Fig. 4(b)).

Now let g be a Nash consistent outcome function with $g(\alpha) = pq\bar{r}$ and $g(\alpha') = p\bar{q}\bar{r}$. Then, coalition $\{2, 3\}$ (weakly) blocks α^a and allocation α^a is *not* strong Nash core stable under g . This is in contrast to allocation α^b , which is strong Nash core stable under g , as, with some effort, can be verified by the reader. \triangleleft

A first question that naturally arises is if, and under which conditions, stable allocations are guaranteed to exist. The following proposition settles this issue, and features an interesting contrast between weak and strong core stability.

Proposition 9: Strong core stable allocations are not guaranteed to exist for partial Boolean games, not even under Nash-consistent outcome functions. By contrast, core stable allocations are guaranteed to exist in partial Boolean games, even under general outcome functions.

Proof: For the first part, consider the partial game $P = (N, \Phi, \Phi_1, \Phi_2, \Phi_3, \gamma_1, \gamma_2, \gamma_3)$ depicted in Fig. 5, where $N = \{1, 2, 3\}$, $\Phi = \{p, q, r, s\}$, $\Phi_1 = \{p\}$, $\Phi_2 = \{q\}$, and $\Phi_3 = \{r\}$. Accordingly, $\Phi_U = \{s\}$, and there are three allocations, α^1 , α^2 , and α^3 , each of which assigns control over s to player 1, player 2, and player 3, respectively, i.e., $\alpha^i(s) = i$ for $1 \leq i \leq 3$. The three Nash equilibria of P are given by $pqr s$, $\bar{p}\bar{q}\bar{r} s$, and $\bar{p}\bar{q}\bar{r} \bar{s}$. Now consider outcome function h :

$$f(\alpha) = \begin{cases} pqr s & \text{if } \alpha = \alpha^1 \\ \bar{p}\bar{q}\bar{r} s & \text{if } \alpha = \alpha^3 \\ \bar{p}\bar{q}\bar{r} \bar{s} & \text{if } \alpha = \alpha^2 \end{cases}$$

Moreover, it is easy to check that coalition $\{1, 3\}$ weakly blocks α^1 , coalition $\{2, 3\}$ weakly blocks α^3 , and coalition $\{1, 2\}$ weakly blocks α^2 . We may conclude that strong core stable allocations are not guaranteed to exist for partial Boolean games. Noting that h defined to be Nash-consistent,

	q	\bar{q}	q	\bar{q}	
p	<u>1, 1, 0</u>	1, 0, 0	1, 0, 0	0, 1, 0	s
\bar{p}	0, 1, 0	1, 0, 0	0, 0, 0	<u>1, 0, 1</u>	
p	0, 0, 1	0, 1, 0	0, 1, 0	1, 0, 0	\bar{s}
\bar{p}	0, 0, 0	<u>0, 1, 1</u>	1, 0, 0	0, 1, 0	
	r		\bar{r}		

Fig. 5. Partial Boolean game showing that strong Nash stable allocations are not guaranteed to exist. Players 1, 2, and 3 control variables p (choosing rows), q (choosing columns), and r (choosing left or right matrix), respectively. Variable s is unassigned, i.e., $\Phi_U = \{s\}$, and the player who controls s decides whether the outcome falls in the upper two matrices (s) or in the lower two matrices \bar{s} . The three underlined outcomes, pqr s, $\bar{p}\bar{q}\bar{r}$ s, and $\bar{p}\bar{q}r\bar{s}$, are the Nash equilibria of the *partial* game, and they remain Nash equilibria of the Boolean games unless s is assigned to player 3, player 2, or player 1, respectively.

it moreover follows that strong Nash core stable allocations are not guaranteed to exist either.

For the second part, consider an arbitrary partial game P and outcome function h . Let α be an allocation that allocates all unassigned propositional variables in Φ_U to one player i . If α is core stable under h , we are done, so assume that this is not the case. Then, there is a coalition S and allocation α' with $\alpha'_{-S} = \alpha_{-S}$ such that $u_j(f(\alpha)) = 0$ and $u_j(f(\alpha')) = 1$ for all $j \in S$. Observe that $i \in S$ and that $\alpha'_k = \emptyset$ for all $k \notin S$. Now assume for contradiction that there is a coalition T blocking α' , i.e., there is an allocation α'' with $\alpha''_{-T} = \alpha'_{-T}$ such that $u_k(f(\alpha')) = 0$ and $u_k(f(\alpha'')) = 1$ for all $k \in T$. Without loss of generality, we may assume that T is not empty, and hence $\alpha'' \neq \alpha'$. Observe that $T \subseteq N \setminus S$. Accordingly, $\alpha'_k = \emptyset$ for all $k \in T$. It follows that $\alpha'' = \alpha'$, a contradiction. We may conclude the proof by observing that this argument still holds if h is assumed to be Nash consistent. \square

We now define the distributed delegation problem, which asks, for a given partial game, whether a *stable* allocation can be found that satisfies the principal's objective. We say that outcome function h is *polynomial* if it yields $h(\alpha)$ in polynomial time on input α . We state the problem generically, but one can vary as to whether weak core stability or strong core stability is to be considered, and whether the outcome functions are required to be Nash consistent or not.

DISTRIBUTED DELEGATION

Given: A partial Boolean game P and polynomial outcome function h , and a formula Υ in $L(\Phi)$

Problem: Does there exist an allocation α that is core stable under h such that Υ is satisfied at $h(\alpha)$?

Before we show that DISTRIBUTED DELEGATION is not harder than WEAK PRINCIPAL DELEGATION, we first establish the computational complexity of the related problem of deciding whether a given allocation is core stable in a given game under

a given outcome function. Formally, we define the problem ALLOCATION STABILITY as follows.

ALLOCATION STABILITY

Given: Partial game P , polynomial outcome function h , and allocation α

Problem: Is allocation α core stable under h in P ?

We now have the following auxiliary coNP-completeness result, of which the hardness part can be interpreted as saying that allocation is coNP-hard *even* when restricted to polynomial and Nash consistent outcome functions. For computationally more complex outcome functions the problem may very well become harder as well.

Proposition 10: ALLOCATION STABILITY is coNP-complete, even if restricted to Nash-consistent outcome functions and strong core stability.

Proof: To see that ALLOCATION STABILITY is in coNP, observe that for a given coalition S , and allocation α' , check whether $\alpha_{-S} = \alpha'_{-S}$ and whether $h(\alpha) \not\models \gamma_i$ and $h(\alpha') \models \gamma_i$ for all $i \in S$. Under the prevailing assumptions, these checks can all be performed in polynomial time. Then, α is *not* core stable if and only if all these checks yield a positive answer.

For coNP-hardness, we reduce from the complement of SATISFIABILITY. Given a propositional formula φ on the propositional variables $\Phi = \{p_1, \dots, p_k\}$, construct partial game P_φ with five players, 0, 1, 2, i_\perp , and i_\top and defined on $\{p, q, r\} \cup \Phi$, where p , q , and r are fresh variables not in Φ . Let $\Phi_0 = \{r\}$, and $\Phi_1 = \Phi_2 = \Phi_{i_\perp} = \Phi_{i_\top} = \emptyset$. Hence, $\Phi_U = \Phi \cup \{p, q\}$. Let the players' goals be given by:

$$\gamma_0 = \perp \quad \gamma_1 = \gamma_2 = \varphi \wedge pqr \quad \gamma_{i_\perp} = \gamma_{i_\top} = pqr$$

Define v_0 as the valuation with $v_0(x) = \perp$ for all $x \in \Phi \cup \{p, q, r\}$. Observe that v_0 is a Nash equilibrium of P_φ under α for every allocation α . We now associate a valuation v_α with every allocation α . If α is such that $\alpha_1 = \{p\}$ and $\alpha_2 = \{q\}$ or $\alpha_1 = \{q\}$ and $\alpha_2 = \{p\}$, then let v_α be such that $v_\alpha(r) = \top$ and for $x \in \Phi \cup \{p, q\}$:

$$v_\alpha(x) = \begin{cases} \top & \text{if } x \in \Phi \text{ and } \alpha(x) = i_\top, \text{ or} \\ & \text{if } x \in \{p, q\}, \alpha_1 = \{q\}, \text{ and } \alpha_2 = \{p\}, \\ \perp & \text{otherwise.} \end{cases}$$

For all other allocations α , set $v_\alpha = v_0$. Now, define the outcome function h^* such that $h^*(\alpha) = v_\alpha$ for all allocations α . It is worth observing that h^* is Nash-consistent. Moreover, on input α , the valuation $h^*(\alpha)$ can be computed in time polynomial in the size of φ .

Now, let α^* be any allocation such that $\alpha_0 = \emptyset$, $\alpha_1^* = \{p\}$, and $\alpha_2^* = \{q\}$. Accordingly, $\alpha_{i_\perp}^* \cup \alpha_{i_\top}^* = \Phi$. Note that $v_{\alpha^*} \models \bar{p}\bar{q}$ and that v_{α^*} is sustained by a Nash equilibrium in $G(P_\varphi, \alpha^*)$. We now prove that:

α^* is Nash core stable under h^* iff φ is not satisfiable.

First assume that φ is satisfiable and that $w: \Phi \rightarrow \{\perp, \top\}$ is a witnessing valuation. Let α' be the allocation such that $\alpha'_0 = \emptyset$, $\alpha'_1 = \{q\}$, $\alpha'_2 = \{p\}$, and

$$\alpha'_{i_\perp} = \{x \in \Phi : w(x) = \perp\} \quad \alpha'_{i_\top} = \{x \in \Phi : w(x) = \top\}.$$

Hence, $v_{\alpha'} \models pq$ and $v_{\alpha'} \models \varphi$. Let $S = \{1, 2, i_{\perp}, i_{\top}\}$. Then, $\alpha^*_{-S} = \alpha'_{-S}$. As $v_{\alpha^*} \not\models \gamma_i$ whereas $v_{\alpha'} \models \gamma_i$ for all $i \in S$, coalition S blocks α^* . Accordingly, α^* is not core stable. As h^* is Nash-consistent, α^* is not Nash-core stable either.

For the opposite direction, assume that φ is not satisfiable and for contradiction that there is a coalition S and an allocation α' such that $\alpha^*_{-S} = \alpha'_{-S}$, $h^*(\alpha^*) \not\models \gamma_i$ and $h^*(\alpha') \models \gamma_i$ for all $i \in S$. As $v \not\models \gamma_i$ for all valuations v and all $i \in \{0, 1, 2\}$, it follows that $S = \{i_{\top}, i_{\perp}\}$. However, $h^*(\alpha') \not\models pqr$ for all allocations α' with $\alpha^*_{\{0,1,2\}} = \alpha'_{\{0,1,2\}}$, and a contradiction ensues. Hence, α^* is Nash core stable under h^* .

For strong Nash stability, the arguments are analogous, modulo a couple of obvious technical details. \square

Using similar proof techniques we now prove the main result of this section. Here, as with ALLOCATION STABILITY, the hardness part provides a lower bound that holds *even* for polynomial and Nash-consistent outcome functions. For outcome functions that are harder to compute, the problem may become harder accordingly.

Theorem 11: DISTRIBUTED DELEGATION is Σ_2^P -complete, even for Nash-consistent outcome functions and strong Nash core stability.

Proof: As allocations are small relative to the input, we can check for a given allocation α whether $h(\alpha) \models \Upsilon$, and, by virtue of Proposition 10, we can consult a coNP-oracle to check whether α is (strong) (Nash) core stable. Then, a (strong) (Nash) core stable allocation α exists under h if and only if these two checks yield positive answers. Overall, we obtain membership of Σ_2^P .

To see that DISTRIBUTED DELEGATION is in Σ_2^P -hard, we reduce from QBF_{2,3}. Let $Q = \exists X \forall Y \varphi$ be a QBF_{2,3} instance. We first construct a partial Boolean game $P_Q = (N, \Phi, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n)$ along with a polynomial and Nash-consistent outcome function h^* and formula Υ such that Q is valid if and only if there is a Nash core stable allocation α^* under h^* in P_Q .

Let P_Q have nine players $0, i_1, i_2, i_{\top}, i_{\perp}, j_1, j_2, j_{\top}$, and j_{\perp} . Furthermore, let $\Phi = X \cup Y \cup \{p, q, r\}$, where p, q , and r are fresh variables not occurring in $X \cup Y$. Furthermore, set $\Phi_0 = \{r\}$ and $\Phi_i = \emptyset$ for any player i other than 0. Thus, $\Phi_U = X \cup Y \cup \{p, q\}$. Define the players' goal be given by:

$$\begin{aligned} \gamma_0 &= \perp & \gamma_{i_1} &= \gamma_{i_2} = \varphi \vee pqr & \gamma_{i_{\perp}} &= \gamma_{i_{\top}} = pqr \\ \gamma_{j_1} &= \gamma_{j_2} = \neg \varphi \wedge \bar{p}\bar{q}r & \gamma_{j_{\perp}} &= \gamma_{j_{\top}} = \bar{p}\bar{q}r \end{aligned}$$

With every allocation α we now associate a valuation v_{α} . Let us call an allocation α “live” if $\alpha_{i_1} \cup \alpha_{i_2} = \{p\}$, $\alpha_{j_1} \cup \alpha_{j_2} = \{q\}$, $X \subseteq \alpha_{i_{\perp}} \cup \alpha_{i_{\top}}$, and $Y \subseteq \alpha_{j_{\perp}} \cup \alpha_{j_{\top}}$. For live allocations α we set $v_{\alpha}(r) = \top$, and for all $z \in X \cup Y \cup \{p, q\}$:

$$v_{\alpha}(z) = \begin{cases} \top & \text{if } z \in X \cup Y \text{ and } z \in \alpha_{i_{\top}} \cup \alpha_{j_{\top}}, \text{ or} \\ & \text{if } z \in \{p, q\}, p \in \alpha_{i_1}, \text{ and } q \in \alpha_{j_1}, \text{ or} \\ & \text{if } z \in \{p, q\}, p \in \alpha_{i_2}, \text{ and } q \in \alpha_{j_2}, \\ \perp & \text{otherwise.} \end{cases}$$

For all other allocations α , set $v_{\alpha} = v_0$. Now, define the outcome function h^* such that $h^*(\alpha) = v_{\alpha}$ for all allocations α . Observe that h^* is Nash-consistent. Moreover,

on input α , the valuation $h^*(\alpha)$ can be computed in time polynomial in the size of α , which is again small relative to φ . Finally, let $\Upsilon = \varphi \wedge pqr$.

We prove that there is a core stable allocation α under h^* with $h^*(\alpha) \models \varphi \wedge pqr$ if and only if $Q = \exists X \forall Y \varphi$ is valid.

For the “if”-direction, assume that $Q = \exists X \forall Y \varphi$ is valid. Then, there is a $v: X \rightarrow \{\perp, \top\} \cup \{p, q, r\}$ such that for all $w: Y \rightarrow \{\perp, \top\}$, we have that $v \cup w \models \varphi$ and $v'(p) = v'(q) = v'(r) = \top$. Then, there is also a live allocation α such that $h^*(\alpha) = v$. Hence $h^*(\alpha)$ satisfies $\varphi \wedge pqr$. Now assume for contradiction there is some coalition S blocking α , i.e., there is some α' with $\alpha_{-S} = \alpha'_{-S}$ such that $u_i(h^*(\alpha')) \succ_i u_i(h^*(\alpha))$ for all $i \in S$. Because of the latter, we find that α' must be a live allocation. Some reflection reveals that $\{j_1, j_2\} \subseteq S \subseteq \{j_1, j_2, j_{\perp}, j_{\top}\}$ and that $h^*(\alpha') \not\models \varphi$. It follows that $h^*(\alpha)(x) = h^*(\alpha')(x)$ for all $x \in X$. But then it would follow that $h^*(\alpha') \models \varphi$, a contradiction.

For the “only if”-direction, assume that $Q = \exists X \forall Y \varphi$ is not valid. Then for every $v: X \rightarrow \{\perp, \top\}$, there is some $w: Y \rightarrow \{\perp, \top\}$ such that $v \cup w \not\models \varphi$. To show that there is no Nash stable allocation α with $h^*(\alpha)$ satisfying $\varphi \wedge pqr$, consider an arbitrary α . Without loss of generality we may assume that α is live, otherwise $h^*(\alpha) = v_0$ and $h^*(\alpha) \not\models \varphi \wedge pqr$. Let $S = \{j_1, j_2, j_{\perp}, j_{\top}\}$ and observe that $h^*(\alpha) \not\models \gamma_j$ for $j \in S$. Without loss of generality, we may assume that $\alpha(p) = i_1$ and $\alpha(q) = j_1$. By assumption, there is some valuation v' with $v' \models \neg \varphi$ and $v'(x) = h^*(\alpha)$ for all $x \in X$. Some reflection reveals that we can define an allocation α' such that $h^*(\alpha') = v'$ and $h^*(v) \models \neg \varphi \wedge \bar{p}\bar{q}r$. Hence, $h^*(\alpha') \models \gamma_j$ for all $j \in S$. As moreover, $\alpha_{-S} = \alpha'_{-S}$, it follows that S blocks α under h^* , and we may conclude that there is no Nash stable allocation α with $h^*(\alpha) \models \varphi \wedge pqr$.

A very similar construction shows that DISTRIBUTED ALLOCATION remains Σ_2^P -hard for *strong* Nash core stability. \square

It is worth observing that the outcome function is part of the input of DISTRIBUTED DELEGATION. For *specific* (Nash-consistent) outcome functions, the computational complexity of the distributed delegation may need to be assessed independently. For instance, one could be specifically be interested in Nash-consistent choice functions that for each allocation choose a Nash equilibrium that maximises social welfare.⁵

VI. DISCUSSION, RELATED WORK, AND FUTURE DIRECTIONS

We have introduced and investigated the problem of how to optimally delegate decisions to self-interested agents. We modelled this delegation problem using the setting of Boolean games. We argued that Boolean games provide a natural framework within which to model the delegation of Boolean decisions: individual decisions naturally map to Boolean variables, owned by individual agents. We distinguished between the

⁵ One could also define stability on *allocation-equilibrium pairs*: (α, \vec{v}) would then be stable if for all coalitions S , all allocations α' with $\alpha_{-S} = \alpha'_{-S}$, and all equilibria \vec{v}' under α , we have $\vec{v} \succeq_i \vec{v}'$ for all $i \in S$. By an argument analogous to the proof of Theorem 2, it can then be shown that the accompanying decision problem whether a stable pair (α, \vec{v}) with \vec{v} satisfying objective Υ exists in a partial game is Σ_2^P -complete, and reduces to deciding the existence of a Nash equilibrium in a Boolean game.

principal delegation problem, where one principal delegates decisions to subordinate agents, and the distributed delegation problem, where the subordinate agents can form coalitions and allocate unallocated variables as they see fit.

Our main results have been mainly negative, in the sense that, even in the elemental Boolean games framework, they all point to a high computational complexity of the various delegation problems we have investigated in this paper. WEAK PRINCIPAL DELEGATION, STRONG PRINCIPAL DELEGATION and DISTRIBUTED DELEGATION are all at least Σ_2^P -hard, while OPTIMAL DELEGATION is $\text{FP}^{\Sigma_2^P}$ -complete. By inspecting the proof of Theorem 2 it can be recognised that the hardness of WEAK PRINCIPAL DELEGATION reduces to the hardness of deciding whether a Nash equilibrium exists in a Boolean game. This raises the question whether the computational complexity of the (principal) delegations problems could be brought down by considering other solution concepts instead of Nash equilibrium to single out the relevant outcomes under a given allocation.

In this context it is worth emphasising that Theorem 11 establishes that the computational complexity of DISTRIBUTED DELEGATION is the same for general polynomial and (polynomial) Nash-consistent outcome functions. This indicates that the intractability of the distributed delegation problem cannot (solely) be attributed to the computational hardness of checking and finding Nash equilibria in Boolean games. The sheer number of allocations, which is exponential in the number of unallocated variables, seems to be a major culprit, and justifies the expectation that many efficient algorithms for the distributed delegation problem will be based on limiting the number of allocations that will have to be considered. In an similar vein, one could explore the different ways in which the problems presented in this paper can be restricted, e.g., by considering restrictions on the players' goals or the principal's objective function.

In our formal setting, we allow for allocations that concentrate control over variables in only a very few agents, which may be undesirable for many applications. Moreover, Section V suggests the existence of a trade-off between stability of an allocation and extreme concentration of power in a few agents. It would therefore be an interesting line of future research to investigate the delegation problems in connection with the various concepts of fairness, (social) welfare, and efficiency (e.g., Pareto optimality) that have been suggested in the literature (cf., e.g., [18], Part II).⁶

Our work can be seen to belong to a stream of work in *mechanism design* in Boolean games, which investigates mechanisms so as to incentivise players to choose strategies that are simultaneously *desirable*, commonly from a principal's point of view, and *rational* from a game theoretic perspective. The mechanisms in point usually aim to modify the Nash equilibria of the Boolean games by imposing taxes on the strategies that are available to the players. This line of research was initiated by [19] and was followed up by, e.g., [5, 20, 21, 22, 23, 24]. Our work on the principal and distributed delegation problems in this paper has a cognate motivation, but importantly differs

in the way the players' incentives are engineered, namely by manipulating the strategies they have at their disposal rather than by levying taxing on playing them.

The delegation problem is also closely related to the *principal-agent problem* studied in economics (see, e.g., [25]). A typical setting for the principal agent problem is where a principal engages the services of an agent to work on behalf of the principal, typically for a fee. The basic issue studied in the principal-agent problem is that the agent will be self-interested, and it may not be feasible for the principal to observe the actions of the agent; in which case, how can the principal be certain that the agent is indeed acting in the principals interests? Typical solutions to the principal-agent involve designing incentive schemes that will help to align the preferences of the agent with those of the principal.

Our work is also relevant to *logics of propositional control* [26, 27, 28]. Originally developed in [26], these logics are specifically intended to reason about scenarios in which a collection of agents each have control over some set of Boolean variables. In [27], Gerbrandy extended this work by also allowing agents to share control over propositional variables. We should also point to other formalisms for reasoning about delegation [29, 30]. However, in these other works, the focus is rather different to our own. For example, in [29] the focus is on decentralised trust management, while [31] presents a logic of delegation based on the STIT ("see to it that") operator. The use of logical systems such as QBF and DCL-PC to analyse delegation would also be worth pursuing in more detail.

Other directions for future research include looking at the problem from the perspective of Stackelberg (leader-follower) games, introducing costs to delegation actions, so that we can consider secondary preferences over allocations, and of course investigating tractable instances of the problem, and developing efficient heuristics for the decision problems we present.

ACKNOWLEDGMENTS

Sarit Kraus was supported by the Google Inter-university center for Electronic Markets and Auctions. Michael Wooldridge and Paul Harrenstein were supported by ERC Advanced Grant 291528 (RACE) and by the Alan Turing Institute, London. Paul Harrenstein was furthermore supported by ERC Grant 639945 (ACCORD).

REFERENCES

- [1] P. Harrenstein, W. van der Hoek, J.-J. Meyer, and C. Witteveen, "Boolean games," in *TARK VIII*, 2001, pp. 287–298.
- [2] P. Harrenstein, "Logic in conflict," Ph.D. dissertation, Utrecht University, 2004.
- [3] E. Bonzon, M. Lagasquie, J. Lang, and B. Zanuttini, "Boolean games revisited," in *ECAI '06*, 2006, pp. 265–269.
- [4] P. E. Dunne, S. Kraus, W. van der Hoek, and M. Wooldridge, "Cooperative Boolean games," in *AA-MAS 2008*, vol. 2, 2008, pp. 1015–1022.

⁶We are grateful to an anonymous reviewers for this suggestion.

- [5] M. Wooldridge, U. Endriss, S. Kraus, and J. Lang, “Incentive engineering for Boolean games,” *Artificial Intelligence*, vol. 195, pp. 418–439, 2013.
- [6] H. von Stackelberg, *Marktform und Gleichgewicht*. Julius Springer Verlag, 1934.
- [7] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. The MIT Press: Cambridge, MA, 1994.
- [8] S. Kraus and M. Wooldridge, “Delegating decisions in strategic settings,” in *ECAI ’12*, 2012, pp. 468–473.
- [9] J. Uckelman, Y. Chevaleyre, U. Endriss, and J. Lang, “Representing utility functions via weighted goals,” *Mathematical Logic Quarterly*, vol. 55, no. 4, pp. 341–361, 2009.
- [10] M. Mavronicolas, B. Monien, and K. W. Wagner, “Weighted Boolean formula games,” in *Algorithms, Probability, Networks, and Games*, C. Zaroliagis, G. Pantziou, and S. Kontogiannis, Eds. Springer, 2015, pp. 49–86.
- [11] D. Han, P. Harrenstein, S. Nugent, J. Philpott, and M. Wooldridge, “Behavioural strategies in weighted Boolean games,” *Information and Computation*, 2020.
- [12] M. Maschler, E. Solan, and S. Zamir, *Game Theory*. Cambridge U.P., 2013.
- [13] C. H. Papadimitriou, *Computational Complexity*. Addison-Wesley, 1994.
- [14] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-completeness*. Freeman, 1979.
- [15] P. E. Dunne, “The complexity of central slice functions,” *Theoretical Computer Science*, vol. 44, pp. 247–257, 1986.
- [16] M. S. Paterson, N. J. Pippenger, and U. Zwick, “Optimal carry save networks,” in *Boolean Function Complexity*, ser. L.M.S. Lecture Note Series, M. S. Paterson, Ed. Cambridge U.P., 1992, vol. 169, pp. 174–201.
- [17] E. Koutsoupias and C. Papadimitriou, “Worst-case equilibria,” in *STACS’99*, 1999, pp. 404–413.
- [18] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, Eds., *Handbook of Computational Social Choice*. Cambridge U.P., 2015.
- [19] U. Endriss, S. Kraus, J. Lang, and M. Wooldridge, “Designing incentives for Boolean Games,” in *AAMAS’11*, 2011.
- [20] P. Turrini, “Endogenous Boolean games,” in *IJCAI’13*, 2013, pp. 390–396.
- [21] —, “Endogenous games with goals: side-payments among goal-directed agents,” *Autonomous Agents and Multi-Agent Systems*, vol. 30, no. 5, pp. 765–792, 2016.
- [22] V. Levit, Z. Komarovskiy, T. Grinshpoun, and A. Meisels, “Tradeoffs between incentive mechanisms in Boolean games,” in *IJCAI’15*, 2015, pp. 68–74.
- [23] P. Harrenstein, P. Turrini, and M. Wooldridge, “Electric Boolean games: redistribution schemes for resource-bounded agents,” in *AAMAS’15*, 2015, pp. 655–663.
- [24] —, “Characterising the manipulability of Boolean games,” in *IJCAI’17*, 2017, pp. 1081–1087.
- [25] J.-J. Laffont and D. Martimort, *The Theory of Incentives: The Principal-Agent Model*. Princeton U.P., 2001.
- [26] W. van der Hoek and M. Wooldridge, “On the logic of cooperation and propositional control,” *Artificial Intelligence*, vol. 64, pp. 81–119, 2005.
- [27] J. Gerbrandy, “Logics of propositional control,” in *AA-MAS’06*, 2006, pp. 193–200.
- [28] W. van der Hoek, D. Walther, and M. Wooldridge, “Reasoning about the transfer of control,” *JAIR*, vol. 37, pp. 437–477, 2010.
- [29] N. Li, B. N. Grosz, and J. Feigenbaum, “Delegation logic: A logic-based approach to distributed authorization,” *ACM Transactions on Information and System Security*, vol. 6, no. 1, pp. 128–171, 2003.
- [30] G. Boella and L. Torre, “Delegation of power in normative multiagent systems,” in *DEON’06*, 2006, pp. 36–52.
- [31] T. J. Norman and C. Reed, “Group delegation and responsibility,” in *AAMAS’02*, 2002, pp. 491–498.

Paul E. Dunne Paul Dunne studied at the Universities of Edinburgh and Warwick and is a Professor of Computer Science at the University of Liverpool where he has been based since 1985. He has published and researched in a wide range of fields including Boolean Function Complexity, Combinatorial analysis of phase-transition phenomena, algorithmic and complexity aspects of multiagent systems, and complexity of formal argumentation. He is the author of three books, in addition to over one hundred specialist journal articles.

Paul Harrenstein is a postdoctoral research associate in the Department of Computer Science at the University of Oxford, where he is also Stipendiary Lecturer at Hertford College. His work focusses on both the cooperative and non-cooperative aspects of interaction in multiagent systems, while applying methods from computational social choice and game theory. He has co-authored 28 journal papers in both economics and computer science, along with one book chapter.

Sarit Kraus Sarit Kraus is a Professor of Computer Science at Bar-Ilan University, Israel. Her research is focused on intelligent agents and multi-agent systems (including people and robots). Kraus was awarded the IJCAI *Computers and Thought Award*, the ACM SIGART Autonomous Agents Research award, ACM Athena Lecturer, the EMET prize and was twice the winner of the IFAAMAS influential paper award. She is a fellow of AAAI, ECCAI, and ACM, and has been a recipient of an ERC Advanced grant.

Michael Wooldridge is a Professor of Computer Science and Head of Department of Computer Science at the University of Oxford, where he is a Fellow of Hertford College. He has been an AI researcher since 1989, and has published more than 400 scientific articles on the subject. He is a Fellow of the Association for Computing Machinery (ACM), the Association for the Advancement of AI (AAAI), and the European Association for AI (EurAI). From 2014–16, he was President of the European Association for AI, and from 2015–17 he was President of the International Joint Conference on AI (IJCAI).