

Incentive Design for Rational Agents

David Hyland¹, Munyque Mittelmann², Aniello Murano², Giuseppe Perelli³, Michael Wooldridge¹

¹University of Oxford

²University of Naples Federico II

³Sapienza University of Rome

{david.hyland, mjw}@cs.ox.ac.uk, {munyque.mittelmann, aniello.murano}@unina.it,
perelli@di.uniroma1.it

Abstract

We introduce *Incentive Design*: a new class of problems for equilibrium verification in multi-agent systems. In our model, agents attempt to maximize their utility functions, which are expressed as formulae in $LTL[\mathcal{F}]$, a quantitative extension of Linear Temporal Logic with functions computable in polynomial time. We assume agents are rational, in the sense that they adopt strategies consistent with game theoretic solution concepts such as Nash equilibrium. For each solution concept we consider, we analyze the problems of verifying whether an incentive scheme achieves a societal objective and finding one that does so, whether it be social welfare or any other aggregate measure of collective well-being. We study both static and dynamic incentive schemes, showing that the latter are more powerful than the former. Finally, we solve the incentive verification and synthesis problems for all the solution concepts we consider, and analyze their complexity.

1 Introduction

The design of mechanisms for aggregating preferences in multi-agent systems is a key problem in algorithmic game theory (Nisan et al. 2007). One classic approach is to design a system such that it satisfies the designer’s objective (e.g., a property such as efficiency) *under the assumption that agents act rationally*. In traditional approaches to formal verification, one is interested in determining whether a given property, expressed as a temporal logic formula, holds on some or all of the possible executions of a system (Baier and Katoen 2008). However, when components of the system are *rational agents*, we can instead focus our attention on only those outcomes of the system that are rational from the perspective of the agents. This is the central motivation behind the fields of rational verification and rational synthesis (Condurache et al. 2016; Filiot, Gentilini, and Raskin 2018; Almagor, Kupferman, and Perelli 2018; Fisman, Kupferman, and Lustig 2010; Kupferman and Shenwald 2022; Brice, Raskin, and van den Bogaard 2023).

In many settings, the underlying structure that describes the agents’ interactions with the environment and other agents is already known. While those systems may not comply with the designer’s objective, a complete redesign is not always feasible. For instance, a country may have environmental legislation but fail to achieve its sustainability objec-

tives because the incentives of private individuals and companies are not aligned with these higher-level objectives. In such cases, policymakers may, for example, resort to establishing taxes based on a company’s pollution rate, or subsidizing public transportation fees to incentivize its use.

In this paper, we are interested in designing incentives in multi-agent systems to motivate rational agents to achieve the designer’s objective. To achieve this, we introduce a new class of problems for equilibrium verification that we call *Incentive Design*. As usual in game theory, we consider agents who seek to maximize their personal utility function. In our setting, such utilities are expressed as formulae of $LTL[\mathcal{F}]$ (Almagor, Boker, and Kupferman 2016), a quantitative extension of Linear Temporal Logic (Pnueli 1977) with functions that are computable in polynomial time. This enables the representation of agents’ utilities as *quantitative temporal specifications*. The rationality of an outcome is defined with respect to game theoretic solution concepts, including *dominant strategy equilibrium*, *Nash equilibrium*, *immune*, and *resilient equilibria*. For each of these solution concepts, we study the problem of finding and verifying an appropriate *incentive mechanism* to achieve a societal objective, be it the social welfare or any other aggregating measure of collective wellbeing. To address these challenges, we introduce two problems: incentive verification and incentive synthesis. In *incentive verification*, the aim is to verify whether the application of a given incentive scheme guarantees some level of satisfaction for a given objective. In *incentive synthesis*, we are interested in finding an incentive scheme, if it exists, which guarantees that the satisfaction value of some objective attains at least a certain level.

We focus on two major classes of incentive mechanisms: *static* and *dynamic*. The former assigns new satisfaction values to some of the variables in the system, regardless of the execution history. The latter dynamically reassigns satisfaction values based on the history of the game thus far. We demonstrate that dynamic incentives are more powerful than their static counterpart. For all solution concepts that we consider, we show that incentive verification is 2EXPTIME-complete for both static and dynamic incentives, and is hence no harder than the corresponding qualitative rational verification problems. Finally, we solve the incentive synthesis problem and prove that it is 2EXPTIME-complete in the static case and in 3EXPTIME in the dynamic case.

Related work. Our work is inspired by a rich line of research analyzing and designing multi-agent systems composed of rational agents, known as rational verification and synthesis.

Also related is the problem of repairing the goals of players (Almagor, Avni, and Kupferman 2015) and synthesis of dynamic norms (Alechina et al. 2022). This problem consists at modifying the objectives or strategies of the players in order to obtain stability or to improve social welfare. Recent work has also proposed the application of formal methods to the verification and synthesis of mechanisms for social choice (Maubert et al. 2021; Mittelmann et al. 2022; Belardinelli et al. 2022; Narayanasamy 2022).

The problem of designing incentives has been considered from the perspective of economics, control theory, and machine learning (Ratliff et al. 2019). The economics approach typically focuses on the treatment of information asymmetries and the design of incentive-compatible mechanisms (Parkes et al. 2010; Laffont and Martimort 2009). The control theory approach introduces variables to encode information about the evolution of the environment, as it depends on agents' choices (Olsder 2009; Weber 2011). Learning approaches to dynamic incentives in multi-agent settings focus solely on adaptively modifying the rewards for quantitative reward-maximising agents (Centeno and Billhardt 2011; Ratliff and Fiez 2020; Yang et al. 2022). Differently from this line of work, our approach is rooted in formal methods, which guarantees the correctness of incentive schemes with respect to the input specification. Instead of specifying a particular global objective (such as utilitarian social welfare maximisation) and then designing a solution to achieve it in a wide variety of environments, our approach is more general in that we allow arbitrary global objectives specified in $LTL[\mathcal{F}]$ to be provided as input to our procedures.

The closest work to ours is a line of research on *Incentive Engineering*. In this direction, (Wooldridge et al. 2013) first considered taxation schemes for one-shot Boolean games, an approach that was later extended to concurrent games where players have lexicographic LTL and mean-payoff objectives (Hyland, Gutierrez, and Wooldridge 2023b). A variant of the problem, where the incentive designer can only observe some subset of the system variables, was studied by (Hyland, Gutierrez, and Wooldridge 2023a). These works investigate how to incentivize agents by introducing costs to their actions. Our notion of incentives is more general, as we allow for changes to any of the atomic features of the game states. For the first time, we consider the setting in which both the agents and the incentive designer have goals expressed in a quantitative temporal logic.

2 Preliminaries

For the remainder of the paper, we fix a set of atomic propositions AP , a set of agents Ag , and a set of strategy variables Var . We also let $\mathcal{F} \subseteq \{f: [-1, 1]^m \rightarrow [-1, 1] \mid m \in \mathbb{N}\}$ be a set of functions computable in polynomial time over $[-1, 1]$ of possibly different arities, that will parameterize the logics we consider. With slight abuse of notation, we denote by $f \in \mathcal{F}$ both the function and the corresponding

function symbol. It will be clear from the context what the symbol corresponds to.

In this paper, we consider a variant of the classic model of concurrent game structures (CGS) (Alur, Henzinger, and Kupferman 2002), where atomic propositions are associated with a weight in different states:

Definition 1. A *weighted concurrent game structure* (wCGS) is a tuple $G = (AP, Ag, (Ac_a)_{a \in Ag}, V, v_i, \delta, \ell)$ where (i) AP is a finite set of *atomic propositions*; (ii) $Ag = \{1, \dots, n\}$ is a finite set of n *agents*; (iii) Ac_a is a finite set of *actions* for agent a and $Ac = \times_{a \in Ag} Ac_a$ is the set of *joint actions*; (iv) V is a finite set of *positions*; (v) $v_i \in V$ is an *initial position*; (vi) $\delta: V \times \times_{a \in Ag} Ac_a \rightarrow V$ is a *transition function*; (vii) $\ell: V \times AP \rightarrow [-1, 1]$ is a *weight function*.

In a position $v \in V$, each player a chooses an action $c_a \in Ac_a$, and the state proceeds to the position $\delta(v, c)$ where c is an *action profile* $(c_a)_{a \in Ag}$.

Given a coalition of agents $C \subseteq Ag$, we write o_C for a tuple of objects $(o_a)_{a \in C}$, one for each agent in C , and such tuples are called *profiles*. We also write o for o_{Ag} . Given a profile o and $a \in Ag$, we let o_a be agent a 's component, and o_{-a} is used to denote the profile $(o_b)_{b \neq a}$.

A *play* $\pi = v_1 v_2 \dots$ is an infinite sequence of positions such that for every $i \geq 1$ there exists an action profile c such that $\delta(v_i, c) = v_{i+1}$. We write $\pi_i = v_i$ for the position at index i in play π and $\pi_{\geq i} = v_i v_{i+1} \dots$ for the suffix of π starting from the position at index i . A *history* h is a finite prefix of a play, $last(h)$ is the last position of history h , $|h|$ is the length of h and $Hist$ is the set of histories.

We will find it useful to generalise the weight function so that it may vary over the course of a play. Given a set AP of atomic propositions, an infinite sequence of weight functions $\vec{\ell} = \ell_1 \ell_2 \dots$, and a play π , the $(AP, \vec{\ell})$ -labelling of π , written $L(\pi; AP, \vec{\ell})$, is given by the infinite sequence of weight assignments

$$L(\pi; AP, \vec{\ell}) := (\ell_1(v_1, p))_{p \in AP} (\ell_2(v_2, p))_{p \in AP} \dots$$

that is obtained by evaluating each weight function ℓ_i in the sequence $\vec{\ell}$ with at the corresponding position v_i over the set AP . In a standard wCGS, the sequence $\vec{\ell}$ is given by the repeating sequence $\vec{\ell} = \ell \ell \ell \dots$. Later, we explore the possibility for an incentive designer to dynamically modify the weight functions, leading to a non-trivial sequence $\vec{\ell}$.

A *strategy* for an agent a is a function $\sigma_a: Hist \rightarrow Ac_a$ that maps each history to an action. We let Str_a be the set of strategies for a and let $Str = \prod_{a \in Ag} Str_a$ be the set of strategy profiles. An *assignment* $\mathcal{A}: Ag \cup Var \rightarrow Str$ is a function from players and variables to strategies. For an assignment \mathcal{A} , an agent a , and a strategy σ for a , $\mathcal{A}[a \mapsto \sigma]$ is the assignment that maps a to σ and is otherwise equal to \mathcal{A} , and $\mathcal{A}[s \mapsto \sigma]$ is defined similarly, where s is a variable. For a strategy profile $\sigma = (\sigma_a)_{a \in Ag}$, we write $\mathcal{A}[Ag \mapsto \sigma]$ for the assignment where $\mathcal{A}[a \mapsto \sigma_a]$ for each $a \in Ag$.

For an assignment \mathcal{A} and a history h , we let $Out(\mathcal{A}, h)$ be the unique play that continues h following the strategies assigned by \mathcal{A} . Formally, $Out(\mathcal{A}, h)$ is the play $h v_0 v_1 \dots$ such

that for all $i \geq 0$, $v_i = \delta(v_{i-1}, \mathbf{c})$ where for all $a \in \text{Ag}$, $\mathbf{c}_a = \mathcal{A}(a)(hv_0 \dots v_{i-1})$, and $v_{-1} = \text{last}(h)$.

Example 1. Consider a scenario involving two manufacturing firms who share the usage of a river and a village community who live downstream of the firms on the same river. We represent this by the set $\text{Ag} = \{1, 2, 3\}$, where agents 1 and 2 are the firms. At every time step, each firm has two choices: to discharge waste water directly into the river, or to first treat the waste water to clean it before discharging it into the river. We represent this by letting $\text{Ac}_i = \{c_a, d_a\}$, $a \in \{1, 2\}$ to denote the ‘clean’ and ‘discharge’ actions for the two firms respectively. The c_a action incurs a fixed cost for firm a , whereas the d_a action incurs no immediate cost for the firm. At every time step, the village community also has two choices: to use the river water for its day-to-day activities, or to use another water source. We denote this by the actions $\text{Ac}_3 = \{r, o\}$, which stand for using the ‘river’ or an ‘other’ source of water respectively. The use of the river by the village incurs a cost proportional to the level of contamination of the river, up to a threshold of k units, in which case the river becomes unusable and incurs a very high cost. The use of the other source of water incurs a high but fixed cost for the community.

In this setting, suppose that the river carries water away (whether clean or dirty) at a constant rate of 1 unit per timestep. We can define states to represent varying degrees of contamination of the river water. Let $V = \{v_0\} \cup \{v_i^c : i \in \{0, \dots, k\}, \mathbf{c} \in \text{Ac}\}$, where $k \in \mathbb{N}_+$ is a threshold such that if the river contamination level ever reaches k units, the river ecosystem is destroyed and the water becomes permanently unusable. The duplication of states for each joint action is used to record the previously taken joint action. The river starts in a clean state ($v_i = \{v_0\}$), and the transition function between intermediate states $v_i^c, i \in \{1, \dots, k-1\}$ can be represented by $\delta(v_i^c, \mathbf{c}') = v_j^{c'}$, where $j = \max(0, \min(k, i + (d-c)/2))$ and D represents the number of agents who selected the d action in \mathbf{c} . Transitions out of the initial state are defined in the same manner, and for the ruin states v_k^c , we let $\delta(v_k^c, \mathbf{c}') = v_k^{c'}$. In other words, if both firms choose the discharge action, the water quality deteriorates by 1 unit, up to the point of ruin. If only one firm chooses the discharge action, the overall quality remains as it is, and if both firms choose the clean action, the river quality improves by 1 unit, up to a maximum level.

We use an atomic proposition q to represent water quality. We can assign $\ell(v_i^c, q) = 1 - 2i/k$ for all $\mathbf{c} \in \text{Ac}$, so that the water quality is 1 when the river is perfectly clean and -1 when the river is unusable.

In order to model the effects of agent actions on their objectives, we additionally assume that the game has atomic propositions $\{u_1, u_2, u_3\}$, which represent the ‘utility’ that each player achieves at a particular state. For a firm $a \in \{1, 2\}$, the weight function for these variables at a state v_i^c is given by $\ell(v_i^c, u_a) = 0.2$ if $c_a = c_a$ and $\ell(v_i^c, u_a) = 0.4$ if $c_a = d_a$. For the community, we let $\ell(v_i^c, u_3) = \ell(v_i^c, q)$ if $\mathbf{c}_3 = r$, $\ell(v_i^c, u_3) = 0$ if $\mathbf{c}_3 = o$. Finally, we let $\ell(v_0, q) = 1$ and $\ell(v_0, p) = 0$ for all $p \in \text{AP} \setminus \{q\}$.

Quantitative Linear Temporal Logic. To specify players’ objectives and define the concept of a rational outcome, we make use of Linear Temporal Logic with quantitative semantics (LTL[\mathcal{F}]) (Almagor, Boker, and Kupferman 2016). We work with this version over LTL primarily because assuming that formula satisfaction values (over which utilities are typically defined) only take on the values of 0 or 1 limits the kinds of incentive schemes that an incentive designer can devise and implement. With quantitative logics such as LTL[\mathcal{F}], there is much more flexibility in capturing the more common conception of an incentive. A classic example of this is in income taxes where the utility of employees depends on their net income, which cannot be reduced to a Boolean outcome.

Definition 2. The syntax of LTL[\mathcal{F}] is defined by the BNF

$$\varphi ::= p \mid f[\varphi, \dots, \varphi] \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi$$

where $p \in \text{AP}$ and $f \in \mathcal{F}$.

Here, \mathbf{X} and \mathbf{U} are the usual temporal operators ‘next’ and ‘until’. The meaning of $f[\varphi_1, \dots, \varphi_n]$ depends on the function f . We use \top , \vee , and \neg to denote, respectively, function 1, function $x, y \mapsto \max(x, y)$ and function $x \mapsto -x$.

We assume here that an agent a ’s decision-making aims to maximise the satisfaction value of an LTL[\mathcal{F}] formula γ_a . Once goals for each agent are specified, we have a *weighted concurrent game*, defined formally as follows:

Definition 3. Let $G = (\text{AP}, \text{Ag}, (\text{Ac}_a)_{a \in \text{Ag}}, V, \delta, \ell, V_i)$ be a wCGS, and \mathcal{A} an assignment. The satisfaction value $\llbracket \varphi \rrbracket_{\mathcal{A}}^G(h) \in [-1, 1]$ of an LTL[\mathcal{F}] formula φ in a history h is defined as follows, where π denotes $\text{Out}(\mathcal{A}, h)$:

$$\begin{aligned} \llbracket p \rrbracket_{\mathcal{A}}^G(h) &= \ell(\text{last}(h), p) \\ \llbracket f[\varphi_1, \dots, \varphi_m] \rrbracket_{\mathcal{A}}^G(h) &= f(\llbracket \varphi_1 \rrbracket_{\mathcal{A}}^G(h), \dots, \llbracket \varphi_m \rrbracket_{\mathcal{A}}^G(h)) \\ \llbracket \mathbf{X}\varphi \rrbracket_{\mathcal{A}}^G(h) &= \llbracket \varphi \rrbracket_{\mathcal{A}}^G(\pi|_{h|+1}) \\ \llbracket \varphi_1 \mathbf{U}\varphi_2 \rrbracket_{\mathcal{A}}^G(h) &= \sup_{i \geq 0} \min(\llbracket \varphi_2 \rrbracket_{\mathcal{A}}^G(\pi|_{h|+i}), \\ &\quad \min_{0 \leq j < i} \llbracket \varphi_1 \rrbracket_{\mathcal{A}}^G(\pi|_{h|+j})) \end{aligned}$$

When the satisfaction value does not depend on the assignment, we write $\llbracket \varphi \rrbracket^G(h)$ for $\llbracket \varphi \rrbracket_{\mathcal{A}}^G(h)$ where \mathcal{A} is any assignment. Additionally, because arena transitions are deterministic with respect to actions in this setting, every assignment gives rise to a unique play. Therefore, we also define the semantics of LTL[\mathcal{F}] formulae over a play π with respect to a valuation function ℓ over a set AP , written $\llbracket \varphi \rrbracket_{\pi}^{\text{AP}, \ell}$, in the same manner as above.

We define some abbreviations, corresponding to classical logic and LTL operators: $\perp := \neg \top$, $\varphi \wedge \varphi' := \neg(\neg\varphi \vee \neg\varphi')$, $\varphi \rightarrow \varphi' := \neg\varphi \vee \varphi'$, $\mathbf{F}\psi := \top \mathbf{U}\psi$, $\mathbf{G}\psi := \neg \mathbf{F}\neg\psi$.

We assume that \mathcal{F} contains the comparison function

$$\leq : (x, y) \mapsto \begin{cases} 1 & \text{if } x \leq y, \\ -1 & \text{otherwise,} \end{cases}$$

similarly, the equality function is defined as

$$= : (x, y) \mapsto \begin{cases} 1 & \text{if } x = y, \\ -1 & \text{otherwise} \end{cases}$$

Definition 4. A *weighted concurrent game* (wCG) is a tuple $\mathcal{G} = (G, (\gamma_a)_{a \in \text{Ag}})$, where G is a wCGS and γ_a is an LTL $[\mathcal{F}]$ formula over AP, for all $a \in \text{Ag}$.

Solution concepts. In this paper we consider agents with LTL $[\mathcal{F}]$ goals, where the formula γ_a denotes the goal of agent a . The utility of agent a in a wCGS G given the strategy profile $\sigma = (\sigma_a)_{a \in \text{Ag}}$ is denoted $\text{util}_a^G(\sigma)$ and is defined as the satisfaction value of γ_a in the initial state of G when agents are assigned to their strategies in σ , that is, $\text{util}_a^G(\sigma) = \llbracket \gamma_a \rrbracket_{\mathcal{A}[\text{Ag} \rightarrow \sigma]}^G(v_\iota)$.

We now recall some key game theoretic solution concepts. Let $\mathcal{G} = (G, (\gamma_a)_{a \in \text{Ag}})$ be a wCG and $\sigma = (\sigma_a)_{a \in \text{Ag}}$ be a strategy profile. We say that σ is a dominant strategy equilibrium if the strategy associated with each agent weakly maximizes her utility, for all possible strategies of other agents. Formally, $\sigma = (\sigma_a)_{a \in \text{Ag}}$ is a *dominant strategy equilibrium* in \mathcal{G} if for all agent a and for all strategy profiles $\sigma' \in \prod_{a \in \text{Ag}} \text{Str}_a$, $\text{util}_a^G(\sigma') \leq \text{util}_a^G(\sigma'_a, \sigma_a)$.

We say σ is a Nash equilibrium if no agent can increase her utility with a unilateral change of strategy. Formally, $\sigma = (\sigma_a)_{a \in \text{Ag}}$ is a *Nash equilibrium* in \mathcal{G} if for all agents a , and all strategy $\sigma'_a \in \text{Str}_a$, we have that $\text{util}_a^G(\sigma_{-a}, \sigma'_a) \leq \text{util}_a^G(\sigma)$.

In an m -resilient equilibrium, we consider deviations by coalitions of agents rather than individuals: this solution concept tolerates deviations of up to m agents (Abraham, Dolev, and Halpern 2008). Formally, $\sigma = (\sigma_a)_{a \in \text{Ag}}$ is an *m -resilient equilibrium* in \mathcal{G} if for all coalitions $C \subseteq \text{Ag}$ with size at most m , and all partial strategy profiles $\sigma'_{\text{Ag} \setminus C} \in \prod_{a \in \text{Ag} \setminus C} \sigma_a$, we have that $\text{util}_a^G(\sigma_{\text{Ag} \setminus C}, \sigma'_C) \leq \text{util}_a^G(\sigma)$.

We write $\text{DSE}(\mathcal{G}), \text{NE}(\mathcal{G}), \text{RE}_m(\mathcal{G})$ to denote the sets of dominant strategy equilibria, Nash equilibria, and m -resilient equilibria of a game \mathcal{G} , respectively.

Example 2. Continuing from Example 1, we assume that the objective of all players is to maintain a high level of utility at all times. Suppose in addition that there is a regulator who is able to impose taxes on the firms. We can model this by introducing two new variables t_1, t_2 to represent the taxes imposed on firms 1 and 2 respectively. Taxes are initially set to 0, i.e., $\ell(v, t_a) = 0$ for all $v \in V, a \in \{1, 2\}$. Given this, we define the goals for the players on a given play π to be $\gamma_a = \mathbf{G}(u_a - t_a)$ for $a \in \{1, 2\}$ and $\gamma_3 = \mathbf{G}u_3$, where the ‘ $-$ ’ function is defined in the obvious manner.

In this scenario, the two firms do not internalize the effect of their actions on the quality of the river. Since discharging wastewater directly into the river is cheaper for them, there is a Nash equilibrium where both firms choose the d action every round, leading to the spoiling of the river after k time steps. An incentive designer may wish to avoid this scenario by putting in place appropriate incentives to discourage indiscriminate pollution. To realize this aim, we now define our model of incentives.

3 Incentives

Given a wCG $\mathcal{G} = (\text{AP}, \text{Ag}, (\text{Ac}_a)_{a \in \text{Ag}}, V, v_\iota, \delta, \ell, (\gamma_a)_{a \in \text{Ag}})$, we interpret incentives as assignments of satisfaction values

to a particular set of propositional variables.

Let $U \subseteq \text{AP}$ be a set of propositional variables in \mathcal{G} . An *incentive scheme* over U is a weight function $\theta : U \rightarrow [-1, 1]$ which induces a modified weight function ℓ_θ such that for all $v \in V$ and $p \in \text{AP}$, we have

$$\ell_\theta(p) = \begin{cases} \theta(p) & p \in U \\ \ell(v, p) & p \notin U. \end{cases}$$

Defined in this way, an incentive scheme applied at a particular point in time modifies the satisfaction values of the variables in U . Depending on the interpretation of the variables, this definition of incentives may take on different meanings. For example, if a variable p represents the monetary cost of being in a particular state, then an incentive on p may be interpreted as a tax or subsidy on the visitation of a state. Moreover, the variables can be used to represent incentives that are agent-specific or ones that are common to multiple agents, depending on the dependencies of utility functions on these variables. Thus, the set U contains the possible variables in the players’ shared environment, which the regulator can intervene upon to influence their decision-making, and are referred to as *incentivized variables*.

Let Θ denote the set of incentive schemes over a set $U \subseteq \text{AP}$. We will assume that incentive schemes have a fixed level of granularity. More specifically, we assume that Θ is a finite set of incentive schemes, and we say that an incentive scheme θ has a granularity $g \in \mathbb{Q}^+$ if for all $v \in V, p \in U$, we have $\theta(p) - \ell(v, p) = k \cdot g$, for some $k \in \mathbb{Z}$.

A *dynamic incentive scheme* is a function $T : \text{Hist} \rightarrow \Theta$ that maps each history to an incentive scheme that is applied on the same round, i.e., the incentive scheme $\theta_k = T(v_1 \dots v_k)$ is applied in round k of the game. A dynamic incentive scheme thus defines a dynamic modification to the satisfaction values of a fixed subset of the propositional variables AP defined in a wCG \mathcal{G} . In contrast, a *static incentive scheme* is a dynamic incentive scheme that outputs the same incentive scheme for all histories ending with the same state. In other words, a static incentive scheme is a *memoryless* dynamic one, which outputs the same incentive at a given state, no matter what history led to the arrival at that state.

It is relatively straightforward to interpret what the application of a static incentive scheme T over a set U means in a given wCG, and we write $\mathcal{G}_T := \mathcal{G}_{\theta^T} = (\text{AP}, \text{Ag}, (\text{Ac}_a)_{a \in \text{Ag}}, V, v_\iota, \delta, \ell_{\theta^T}, (\gamma_a)_{a \in \text{Ag}})$, where θ^T is the incentive scheme that T outputs at each point in the game. A dynamic incentive scheme, on the other hand, can be thought of as inducing a dynamic (history-dependent) weight function on the game \mathcal{G} . In particular, given a play π of a game \mathcal{G} , a dynamic incentive scheme T induces an infinite sequence $\vec{\ell}_T(\pi) = \ell_{\theta^1} \ell_{\theta^2} \dots$ of weight functions, such that for every history $h = v_1 \dots v_k$ of π , we have $\theta^k = T(h)$.

More generally, dynamic weight functions be accommodated in our model by defining new semantics for the satisfaction value $\llbracket \varphi \rrbracket_{\mathcal{A}}^{G, \vec{\ell}}(h)$ of an LTL $[\mathcal{F}]$ formula in a wCGS G , given an infinite sequence of weight functions $\vec{\ell} = \ell_1 \ell_2 \dots$ over a finite set $\mathcal{L} = \{\ell_1, \ell_2, \dots, \ell_{|\mathcal{L}|}\}$ of possible weight functions. This is achieved by letting

$$\llbracket p \rrbracket_{\mathcal{A}}^{G, \vec{\ell}}(h) := \ell_{|h|}(\text{last}(h), p),$$

while the semantics for f , the \mathbf{X} operator, and the \mathbf{U} operator remain the same. Thus, in the specific context of a dynamic incentive scheme T over a set $U \subseteq \text{AP}$, we then have

$$\llbracket p \rrbracket_{\mathcal{A}}^{G_T}(h) := \llbracket p \rrbracket_{\mathcal{A}}^{G, \vec{\ell}_T}(h) = \begin{cases} T(h)(p) & p \in U \\ \ell(\text{last}(h), p) & p \notin U. \end{cases}$$

We also define the semantics of $\text{LTL}[\mathcal{F}]$ over $(\text{AP}, \vec{\ell})$ -labelled plays $L(\pi; \text{AP}, \vec{\ell})$, written $\llbracket \varphi \rrbracket_{\pi}^{\text{AP}, \vec{\ell}}$, in the same way.

Due to the fixed granularity assumption, a dynamic incentive scheme's set of possible outputs is always finite. As we demonstrate later, this assumption is important as it allows one to model check $\text{SL}[\mathcal{F}]$ formulae with this modified semantics, albeit at some additional computational cost. However, this does not affect the worst-case complexity results we derive here.

We argue that for most practical purposes, this is a reasonable assumption to make, as it will usually be the case that beyond a certain level of granularity, the benefits of introducing an even finer partition of the incentive space will be outweighed by the additional complexity introduced in computing the optimal incentive scheme and then implementing it in practice. Examples of granularity appearing in real-world incentive policies include taxation rates or the US Federal Reserve's interest rate on reserve balances.

Example 3. With incentives defined, we can now consider possible incentive implementations by a regulator in our running example. In most instances, a regulator typically has an *objective* that they would like to see accomplished as the result of the implementation of incentive schemes.

In our case, suppose that the regulator wants to impose taxes on the firms to maintain a stable and acceptable water quality level. To achieve this aim, the regulator can intervene on the t_a variables (i.e., $U = \{t_1, t_2\}$), setting them to some positive taxation rate with a granularity of $g = 0.2$. This objective and the taxation constraints may be expressed by the formula $\varphi = \mathbf{G}(g = 1 \wedge t_1 \geq 0 \wedge t_2 \geq 0)$.

With static incentive schemes, the only way to achieve this is to set the values of t_1 and t_2 so that at least one of the firms is worse off by performing the d action. If only one firm is taxed in this way, the policy may be criticised as being unfair. If both firms are taxed in this way, this results in an unnecessary loss of profits to both firms. However, if a dynamic incentive scheme is used, which alternates between taxing the firms a sufficient amount for selecting the d action, the regulator can achieve their objective in a fair and efficient manner.

Problem statement. Using the definitions thus far, we are able to state the core problems that we study in the remainder of this text. We begin with the problems in the family of *incentive verification* problems, which ask whether a given static/dynamic incentive scheme guarantees a certain best- or worst-case equilibrium satisfaction value for a given $\text{LTL}[\mathcal{F}]$ formula φ . For a game \mathcal{G} , formula φ , static incentive scheme T , solution concept ζ , and threshold $c \in [-1, 1]$.

ζ -S-E-INCENTIVE-VERIFICATION: is there a strategy profile $\sigma \in \zeta(\mathcal{G})$ such that $\llbracket \varphi \rrbracket_{\mathcal{A}[\text{Ag} \rightarrow \sigma]}^{G_T} \geq c$?

ζ -S-A-INCENTIVE-VERIFICATION: does it hold that for all strategy profiles $\sigma \in \zeta(\mathcal{G})$, we have $\llbracket \varphi \rrbracket_{\mathcal{A}[\text{Ag} \rightarrow \sigma]}^{G_T} \geq c$?

Similarly, we can define the incentive verification problems for dynamic incentive schemes. For a game \mathcal{G} , formula φ , dynamic incentive scheme T , solution concept ζ , threshold $c \in [-1, 1]$.

ζ -D-E-INCENTIVE-VERIFICATION: is there a strategy profile $\sigma \in \zeta(\mathcal{G})$ such that $\llbracket \varphi \rrbracket_{\mathcal{A}[\text{Ag} \rightarrow \sigma]}^{G_T} \geq c$?

ζ -D-A-INCENTIVE-VERIFICATION: does it hold that for all strategy profiles $\sigma \in \zeta(\mathcal{G})$, we have $\llbracket \varphi \rrbracket_{\mathcal{A}[\text{Ag} \rightarrow \sigma]}^{G_T} \geq c$?

While the ability to verify the effectiveness of a proposed incentive scheme is arguably a useful tool to have, policy-makers are often faced with the task of *designing* effective incentive schemes for a given purpose. This motivates the study of what we refer to as the class of *incentive synthesis* problems, which are defined as follows.

For a given game \mathcal{G} , formula φ , solution concept ζ , incentivised variables $U \subseteq \text{AP}$, threshold $c \in [-1, 1]$.

ζ -S-E-INCENTIVE-SYNTHESIS: is there a static incentive scheme T over U such that for some strategy profile $\sigma \in \zeta(\mathcal{G})$, we have $\llbracket \varphi \rrbracket_{\mathcal{A}[\text{Ag} \rightarrow \sigma]}^{G_T} \geq c$?

ζ -S-A-INCENTIVE-SYNTHESIS: is there a static incentive scheme T over U such that for all strategy profiles $\sigma \in \zeta(\mathcal{G})$, we have $\llbracket \varphi \rrbracket_{\mathcal{A}[\text{Ag} \rightarrow \sigma]}^{G_T} \geq c$?

ζ -D-E-INCENTIVE-SYNTHESIS: is there a dynamic incentive scheme T over U such that for some strategy profile $\sigma \in \zeta(\mathcal{G})$, we have $\llbracket \varphi \rrbracket_{\mathcal{A}[\text{Ag} \rightarrow \sigma]}^{G_T} \geq c$?

ζ -D-A-INCENTIVE-SYNTHESIS: is there a dynamic incentive scheme T over U such that for all strategy profiles $\sigma \in \zeta(\mathcal{G})$, we have $\llbracket \varphi \rrbracket_{\mathcal{A}[\text{Ag} \rightarrow \sigma]}^{G_T} \geq c$?

We remark that although these problems are stated in terms of thresholds, our method of solving these problems allows one to find the optimal satisfaction values of the given formula under some/all equilibria of a game.

4 Static Incentives

In this section, we show how to solve emptiness and synthesis for the static incentive case. We do this by employing procedures for the model checking of formulae in $\text{SL}[\mathcal{F}]$, an extension of $\text{LTL}[\mathcal{F}]$ that includes strategy quantifiers that are useful to quantify over the strategic abilities of agents (Bouyer et al. 2023). We first recall the syntax and semantics of $\text{SL}[\mathcal{F}]$, together with a statement about its model checking against a wCGS. Subsequently, we show how $\text{SL}[\mathcal{F}]$ can be used to express all the various solution concepts considered in this paper. Finally, we show how to use such a representation to solve both the verification and synthesis problems for the case of static incentives.

Definition 5. The syntax of $\text{SL}[\mathcal{F}]$ is defined by the BNF

$$\varphi ::= \psi \mid \exists s. \varphi \mid (a, s)\varphi$$

where ψ is an $\text{LTL}[\mathcal{F}]$ formula, $s \in \text{Var}$, and $a \in \text{Ag}$.

Intuitively, the new operator $\exists s. \varphi$ selects a strategy that maximizes the satisfaction value of φ ; $(a, s)\varphi$ means that the strategy s is assigned to agent a in the evaluation of φ .

Definition 6. Let $G = (\text{AP}, \text{Ag}, (\text{Ac}_a)_{a \in \text{Ag}}, V, \delta, \ell, V_i)$ be a wCGS, and \mathcal{A} an assignment. The satisfaction value $\llbracket \varphi \rrbracket_{\mathcal{A}}^G(h) \in [-1, 1]$ of an $\text{SL}[\mathcal{F}]$ formula φ in a history h is defined as follows:

$$\begin{aligned} \llbracket \exists s. \varphi \rrbracket_{\mathcal{A}}^G(h) &= \max_{\sigma \in \text{Str}} \llbracket \varphi \rrbracket_{\mathcal{A}[s \mapsto \sigma]}^G(h) \\ \llbracket (a, s)\varphi \rrbracket_{\mathcal{A}}^G(h) &= \llbracket \varphi \rrbracket_{\mathcal{A}[a \mapsto \mathcal{A}(s)]}^G(h). \end{aligned}$$

The other cases are similar to the semantics of $\text{LTL}[\mathcal{F}]$.

We also use the standard universal quantifiers $\forall s. \varphi := \neg \exists s. \neg \varphi$. For a variable profile $s = (s_1, \dots, s_n)$ with $1 \leq n \leq |\text{Ag}|$, we use the abbreviations $\exists s \varphi := \exists s_1, \dots, \exists s_n \varphi$ and $\forall s \varphi := \forall s_1, \dots, \forall s_n \varphi$.

Definition 7 ($\text{SL}[\mathcal{F}]$ model-checking (Bouyer et al. 2023)). Given an $\text{SL}[\mathcal{F}]$ formula φ , a wCGS G , an assignment \mathcal{A} , a history h , and a predicate $P \subseteq [-1, 1]$, decide whether $\llbracket \varphi \rrbracket_{\mathcal{A}}^G(h) \in P$.

In (Bouyer et al. 2023), it is shown that the model-checking problem for $\text{SL}[\mathcal{F}]$ is decidable and its complexity is $(k+1)$ -EXPTIME-complete, where k is the *block-nesting* depth of the checked formula. Informally, the block-nesting depth counts how many times an $\text{SL}[\mathcal{F}]$ formula in prenex-normal form alternates between an existential and a universal quantifier¹.

Expressing solution concepts in $\text{SL}[\mathcal{F}]$. Let $s = (s_a)_{a \in \text{Ag}}$ be a profile of strategy variables and γ_a be a $\text{LTL}[\mathcal{F}]$ formula denoting agent a 's goal. The following $\text{SL}[\mathcal{F}]$ formulae express the notion that the strategy profile assigned by σ is a strategic equilibrium:

$$\text{DSE}(s) := \bigwedge_{a \in \text{Ag}} \forall t. [(\text{Ag}, t)\gamma_a \leq (a, s_a)(\text{Ag}_{-a}, t_{-a})\gamma_a]$$

$$\text{NE}(s) := \bigwedge_{a \in \text{Ag}} \forall t. [(\text{Ag}_{-a}, s_{-a})(a, t)\gamma_a \leq (\text{Ag}, s)\gamma_a]$$

$$\text{RE}_m(s) := \bigwedge_{\substack{C \subseteq \text{Ag}: |C| \leq m}} \forall t_C \left[\bigwedge_{a \in C} ((\text{Ag}_{-C}, s_{-C})(C, t_C)\gamma_a \leq (\text{Ag}, s)\gamma_a) \right]$$

Proposition 1. For a wCG $\mathcal{G} = (G, (\gamma_a)_{a \in \text{Ag}})$, a strategy profile σ and a variable profile $s = (s_a)_{a \in \text{Ag}}$, it holds that:

- σ is a dominant strategy equilibrium iff $\llbracket \text{DSE}(s) \rrbracket_{\mathcal{A}[a \mapsto \mathcal{A}(s)]}^G(h) = 1$;
- σ is a Nash equilibrium iff $\llbracket \text{NE}(s) \rrbracket_{\mathcal{A}[a \mapsto \mathcal{A}(s)]}^G(h) = 1$;
- σ is an m -resilient equilibrium iff $\llbracket \text{RE}_m(s) \rrbracket_{\mathcal{A}[a \mapsto \mathcal{A}(s)]}^G(h) = 1$.

¹By prenex normal form, we mean formulae in which all the quantifiers appear in the outermost part of the formula.

Proof. These follow from the definitions of $\text{util}_a^G(\sigma)$ and the solution concepts. \square

Using these constructions, we can now address the verification and synthesis problems for static incentives.

Theorem 1. For $\zeta \in \{\text{DSE}, \text{NE}, \text{RE}_m\}$, $m \in \{1, \dots, n\}$, ζ -S-E-INCENTIVE-VERIFICATION and ζ -S-A-INCENTIVE-VERIFICATION are 2EXPTIME-complete.

Proof. We begin with ζ -S-E-INCENTIVE-VERIFICATION, noting that the approach for the dynamic counterpart is similar. To establish the upper bound, observe that model checking the following $\text{SL}[\mathcal{F}]$ formula in the game \mathcal{G}_T induced by T decides the ζ -S-E-INCENTIVE-VERIFICATION problem:

$$\exists \sigma. [\zeta(\sigma) \wedge (\text{Ag}, \sigma)\varphi].$$

By Proposition 1 and the semantics of the $\text{SL}[\mathcal{F}]$ operators $\exists \sigma$ and \wedge , we see that the satisfaction value of the above formula either represents the highest satisfaction value of φ that can be obtained by some ζ -equilibrium of \mathcal{G}_T , or it takes on the value of -1. Using this, we can compare the obtained satisfaction value with the threshold c to decide the answer to ζ -S-E-INCENTIVE-VERIFICATION. Membership in 2EXPTIME then straightforwardly follows from the fact that model checking an $\text{SL}[\mathcal{F}]$ formula can be done in $(k+1)$ -EXPTIME for formulas φ where the number of alternations in strategy quantifiers is k (Bouyer et al. 2023).

For ζ -S-A-INCENTIVE-VERIFICATION, we can apply the same reasoning as above, only this time with the following $\text{SL}[\mathcal{F}]$ formula:

$$\forall \sigma. [\zeta(\sigma) \rightarrow (\text{Ag}, \sigma)\varphi].$$

For hardness, we reduce from qualitative (strong) rational synthesis, which is the problem of deciding whether a strategy profile exists in a concurrent game such that the goal φ_0 of a player 0 is satisfied in some rational outcome of the game while holding this player's strategy fixed (Kupferman, Perelli, and Vardi 2016). Here, the term 'rational' includes the solution concepts we consider, namely dominant strategy equilibria, Nash equilibria, and h -resilient equilibria. The qualitative component means that player goals are given by formulae expressed in LTL, which is a special case of $\text{LTL}[\mathcal{F}]$ (Bouyer et al. 2023), where the following conditions apply:

- For all $v \in V, p \in \text{AP}, \ell(v, p) \in \{-1, 1\}$;
- For all players $a \in \text{Ag}$, γ_a is an LTL formula, i.e., an $\text{LTL}[\mathcal{F}]$ formula where $\mathcal{F} = \{\neg, \vee, \wedge\}$.

The strong version of the rational synthesis problem quantifies universally over rational outcomes, i.e., it asks whether there exists a strategy for player 0 such that φ_0 is guaranteed to be satisfied for *all* rational outcomes for the remaining players. Since the (strong) rational synthesis problem is 2EXPTIME-complete for all solution concepts we consider (Kupferman, Perelli, and Vardi 2016), we can straightforwardly obtain the lower bound by reducing from standard rational synthesis to our variants of the E-INCENTIVE-VERIFICATION problem and from strong rational synthesis to the A-INCENTIVE-VERIFICATION problem by sim-

ply verifying an incentive scheme which applies no modifications to the game, where the formula to be verified is given by φ_0 and the player whose strategy is fixed is an additional inconsequential dummy player who is added to the game. \square

For the synthesis of a static incentive, we can proceed similarly to the emptiness problem. Essentially, we first (non-deterministically) guess an appropriate static incentive, then check that it solves the problem by employing the emptiness procedure introduced above.

Theorem 2. For $\zeta \in \{DSE, NE, RE_m\}, m \in \{1, \dots, n\}$, ζ -S-E-INCENTIVE-SYNTHESIS and ζ -S-A-INCENTIVE-SYNTHESIS are 2EXPTIME-complete.

Proof. For the upper bound, we proceed by simply nondeterministically guessing an incentive scheme, of which there are $|\Theta| = |\text{AP}|^{1/g}$, where g is the granularity of incentive schemes. Since, g is assumed to be a rational number, the number of incentive schemes is thus exponential in the representation of the granularity.

Static incentive schemes can be represented in space at most equal to that required for representing the valuation function ℓ , which is hence polynomial in the size of the input. Once this incentive scheme is guessed, we can run ζ -S-E-INCENTIVE-VERIFICATION for the existential version of incentive synthesis or ζ -S-A-INCENTIVE-VERIFICATION for the universal version to obtain a solution to our problems. Since guessing a static incentive scheme is in PSPACE and running an incentive verification procedure is in 2EXPTIME for both cases, the overall complexity of the procedure is in 2EXPTIME.

For the lower bound, we can apply the same reduction as used in Proposition 1, but with an additional modification to the game to account for the incentive design component. This can be handled by simply introducing a new propositional variable $q \notin \text{AP}$ to the game, which is guaranteed not to influence the decision-making of agents in any way, and letting the incentive designer intervene only on this new variable, i.e., setting $U = \{q\}$. In this way, there exists an incentive scheme such that the satisfaction value of φ_0 is 1 on some/all ζ -equilibria of the game if and only if the answer to the rational synthesis problem is “yes”. \square

Given these results, we observe that in the case of static incentives, the incentive verification and synthesis problems have the same worst-case complexity. As our results in the following section will demonstrate, this is not necessarily the case when considering dynamic incentive schemes. This contrast, along with the fact that static incentive schemes are typically simpler to implement and communicate in practice, suggests that it may be better to use static incentives in cases where they are sufficient to implement a goal.

5 Dynamic Incentives

Especially in the context of dynamic incentive schemes, a key aspect of our approach is to embed the incentive designer into the game as a player, whose actions correspond

to the application of incentives. To achieve this, we introduce a transformation which converts a given wCGS \mathcal{G} into a modified game \mathcal{G}' , which naturally encodes the abilities of the incentive designer to modify the values of propositional variables as described above. Intuitively, the transformation works by interleaving actions of the incentive designer (which correspond to incentive schemes) and the actions of the players in the original game. This is done because the incentive designer requires access to the history of the game in order to decide what incentives to implement on a given round, but this decision must be made *before* the next actions of the agents so that it applies on the round in which they take their actions. To perform this interleaving while preserving the satisfaction values of players' goals on the modified runs, we will first need to introduce the notion of run inflations and formula translations (Kučera and Strejček 2005).

Given an $(\text{AP}, \vec{\ell})$ -labelled play $L_1 = L(\pi; \text{AP}, \vec{\ell})$ and an $(\text{AP}', \vec{\ell}')$ -labelled play $L_2 = L(\pi'; \text{AP}', \vec{\ell}')$ where $\text{AP} \subseteq \text{AP}'$, we say that L_2 is a d -fold inflation of L_1 if it is the case that for all $p \in \text{AP}$, we have $\ell'_{di}(v'_{di}, p) = \ell_i(v_i, p)$ for every $i \in \mathbb{Z}^+$.

We say that a d -fold inflation L_2 of L_1 is r -labelled if $r \in \text{AP}'$ and for all $i, j \geq 0$, we have that $\ell'_j(v'_j, r) = 1$ if $j = di$ and $\ell'_j(v'_j, r) = -1$ otherwise. Intuitively, r -labelling is an indicator taking the value of 1 only at the positions in the inflated play π' that correspond to positions in the original play π , which will be useful for translating $\text{LTL}[\mathcal{F}]$ formulae over the latter into formulae over the former.

Next, we define a translation function tr^d which transforms an $\text{LTL}[\mathcal{F}]$ formula φ into another $\text{LTL}[\mathcal{F}]$ formula $\text{tr}^d(\varphi)$ that preserves the satisfaction value of φ over a d -fold inflation of any labelled play over which φ is evaluated. More precisely, let φ be an $\text{LTL}[\mathcal{F}]$ formula over a set AP of atomic propositions. The d -fold translation $\text{tr}^d(\varphi)$ of φ is the $\text{LTL}[\mathcal{F}]$ formula defined over the set $\text{AP} \cup \{r\}$, where r is a new atomic proposition, that is given by the following semantics:

- $\text{tr}^d(p) = \mathbf{X}^{d-1}p$ for every $p \in \text{AP}$;
- $\text{tr}^d(f[\varphi_1, \dots, \varphi_m]) = f[\text{tr}^d(\varphi_1), \dots, \text{tr}^d(\varphi_m)]$;
- $\text{tr}^d(\mathbf{X}\varphi) = \mathbf{X}^d\text{tr}^d(\varphi)$;
- $\text{tr}^d(\varphi_1 \mathbf{U} \varphi_2) = (r \rightarrow \text{tr}^d(\varphi_1)) \mathbf{U} (r \rightarrow \text{tr}^d(\varphi_2))$

where \mathbf{X}^e represents an application of the \mathbf{X} operator e times. The following is an adaptation of Lemma 1 in (Gutierrez, Perelli, and Wooldridge 2018) to the setting of $\text{LTL}[\mathcal{F}]$ formulae that we consider.

Lemma 1. Let AP and AP' be two disjoint sets of atomic propositions with $r \in \text{AP}'$. Let $L_1 = L(\pi; \text{AP}, \vec{\ell})$ be an $(\text{AP}, \vec{\ell})$ -labelled play and suppose that $L_2 = L(\pi'; \text{AP}', \vec{\ell}')$ is an r -labelled, d -fold inflation of L_1 with $\text{AP}' = \text{AP} \cup \text{AP}'$. Then, for all $\text{LTL}[\mathcal{F}]$ formulae φ over L_1 , it holds that $\llbracket \varphi \rrbracket_{\pi}^{\text{AP}, \vec{\ell}} = \llbracket \text{tr}^d(\varphi) \rrbracket_{\pi'}^{\text{AP}', \vec{\ell}'}$.

Proof. The proof goes by structural induction on the formula φ , that is, we show that for all $\text{LTL}[\mathcal{F}]$ formula φ and

$i \geq 0$, it holds that $\llbracket \varphi \rrbracket_{\pi_{\geq i}}^{\text{AP}, \vec{\ell}} = \llbracket \text{tr}^d(\varphi) \rrbracket_{\pi'_{\geq di}}^{\text{AP}', \vec{\ell}'}$. Considering the base case, it is straightforward to see that $\llbracket p \rrbracket_{\pi_{\geq 1}}^{\text{AP}, \vec{\ell}} = \llbracket \text{tr}^d(p) \rrbracket_{\pi'_{\geq d-1}}^{\text{AP}', \vec{\ell}'}$ by the definition of the d -fold inflation. Now, by the induction hypothesis and the definition of the semantics for $\llbracket f[\varphi_1, \dots, \varphi_m] \rrbracket_{\pi_{\geq i}}^{\text{AP}, \vec{\ell}}$, it follows that

$$\begin{aligned} \llbracket \text{tr}^d(f[\varphi_1, \dots, \varphi_m]) \rrbracket_{\pi'_{\geq di}}^{\text{AP}', \vec{\ell}'} &= \llbracket f[\text{tr}^d(\varphi_1), \dots, \text{tr}^d(\varphi_m)] \rrbracket_{\pi'_{\geq di}}^{\text{AP}', \vec{\ell}'} \\ &= f(\llbracket \text{tr}^d(\varphi_1) \rrbracket_{\pi'_{\geq di}}^{\text{AP}', \vec{\ell}'}, \dots, \llbracket \text{tr}^d(\varphi_m) \rrbracket_{\pi'_{\geq di}}^{\text{AP}', \vec{\ell}'}) \\ &= f(\llbracket \varphi_1 \rrbracket_{\pi_{\geq i}}^{\text{AP}, \vec{\ell}}, \dots, \llbracket \varphi_m \rrbracket_{\pi_{\geq i}}^{\text{AP}, \vec{\ell}}) = \llbracket f[\varphi_1, \dots, \varphi_m] \rrbracket_{\pi_{\geq i}}^{\text{AP}, \vec{\ell}}. \end{aligned}$$

Moving on to the next operator, suppose that $\varphi = \mathbf{X}\psi$ for some LTL $[\mathcal{F}]$ formula ψ and recall that by definition, $\llbracket \mathbf{X}\psi \rrbracket_{\pi_{\geq i}}^{\text{AP}, \vec{\ell}} = \llbracket \psi \rrbracket_{\pi_{\geq i+1}}^{\text{AP}, \vec{\ell}}$. By the semantics of the \mathbf{X} operator and the induction hypothesis, we have

$$\begin{aligned} \llbracket \mathbf{X}\psi \rrbracket_{\pi_{\geq i}}^{\text{AP}, \vec{\ell}} &= \llbracket \psi \rrbracket_{\pi_{\geq i+1}}^{\text{AP}, \vec{\ell}} = \llbracket \text{tr}^d(\psi) \rrbracket_{\pi'_{\geq d(i+1)}}^{\text{AP}', \vec{\ell}'} \\ &= \llbracket \mathbf{X}^d \text{tr}^d(\psi) \rrbracket_{\pi'_{\geq di}}^{\text{AP}', \vec{\ell}'} = \llbracket \text{tr}^d(\mathbf{X}\psi) \rrbracket_{\pi'_{\geq di}}^{\text{AP}', \vec{\ell}'}. \end{aligned}$$

Finally, for the until operator, we apply the same reasoning. Suppose that $\varphi = \varphi_1 \mathbf{U} \varphi_2$. Then, by the semantics of the \mathbf{U} operator and the induction hypothesis, we have

$$\begin{aligned} \llbracket \varphi_1 \mathbf{U} \varphi_2 \rrbracket_{\pi_{\geq i}}^{\text{AP}, \vec{\ell}} &= \sup_{j \geq 0} \min(\llbracket \varphi_2 \rrbracket_{\pi_{\geq i+j}}^{\text{AP}, \vec{\ell}}, \min_{0 \leq k < j} \llbracket \varphi_1 \rrbracket_{\pi_{\geq i+k}}^{\text{AP}, \vec{\ell}}) \\ &= \sup_{j \geq 0} \min(\llbracket \text{tr}^d(\varphi_2) \rrbracket_{\pi'_{\geq d(i+j)}}^{\text{AP}', \vec{\ell}'}, \\ &\quad \min_{0 \leq k < j} \llbracket \text{tr}^d(\varphi_1) \rrbracket_{\pi'_{\geq d(i+k)}}^{\text{AP}', \vec{\ell}'}) \\ &= \sup_{j \geq 0} \min(\llbracket \max(-r, \text{tr}^d(\varphi_2)) \rrbracket_{\pi'_{\geq d(i+j)}}^{\text{AP}', \vec{\ell}'}, \\ &\quad \min_{0 \leq k < j} \llbracket \max(-r, \text{tr}^d(\varphi_1)) \rrbracket_{\pi'_{\geq d(i+k)}}^{\text{AP}', \vec{\ell}'}) \\ &= \sup_{j \geq 0} \min(\llbracket \max(-r, \text{tr}^d(\varphi_2)) \rrbracket_{\pi'_{\geq di+j}}^{\text{AP}', \vec{\ell}'}, \\ &\quad \min_{0 \leq k < j} \llbracket \max(-r, \text{tr}^d(\varphi_1)) \rrbracket_{\pi'_{\geq di+k}}^{\text{AP}', \vec{\ell}'}) \\ &= \sup_{j \geq 0} \min(\llbracket r \rightarrow \text{tr}^d(\varphi_2) \rrbracket_{\pi'_{\geq di+j}}^{\text{AP}', \vec{\ell}'}, \\ &\quad \min_{0 \leq k < j} \llbracket r \rightarrow \text{tr}^d(\varphi_1) \rrbracket_{\pi'_{\geq di+k}}^{\text{AP}', \vec{\ell}'}) \\ &= \llbracket \text{tr}^d(\varphi_1 \mathbf{U} \varphi_2) \rrbracket_{\pi'_{\geq di}}^{\text{AP}', \vec{\ell}'}. \end{aligned}$$

□

Together, the concepts of inflation and translation allow us to define the *incentive-augmented game* \mathcal{G}' , which is a modified version of a game \mathcal{G} which includes the incentive designer as a player whose actions represent the imposition of different incentive schemes. Given a wCG $\mathcal{G} = (\text{AP}, \text{Ag}, (\text{Ac}_a)_{a \in \text{Ag}}, V, v_i, \delta, \ell, (\gamma_a)_{a \in \text{Ag}})$, the *incentive-augmented version* of \mathcal{G} is defined as the wCG $\mathcal{G}' = (\text{AP}', \text{Ag}', (\text{Ac}_a)_{a \in \text{Ag}'}, V', v_i, \delta', \ell', (\gamma'_a)_{a \in \text{Ag}'})$, where

- $\text{AP}' = \text{AP} \cup \{r\}$, where $r \notin \text{AP}$.
- $\text{Ag}' = \text{Ag} \cup \{n+1\}$, where player $n+1$ corresponds to the incentive designer;
- Ac_a remains unchanged for all $a \in \text{Ag}$, and $\text{Ac}_{n+1} = \Theta$;
- $V' = V \cup V^\Theta$, where $V^\Theta = \{v^\theta \mid v \in V, \theta \in \Theta\}$;
- $\delta' : V' \times \prod_{a \in \text{Ag}'} \text{Ac}'_a \rightarrow V'$ is defined in the following way. For all $v \in V, \mathbf{c}' \in \text{Ac}'$ with $\mathbf{c}'_{n+1} = \theta \in \Theta$, we have $\delta'(v, \mathbf{c}') = v^\theta$ and for all $v \in V^\Theta, \mathbf{c}' \in \text{Ac}'$ with $\mathbf{c}'_{-n+1} = (c_a)_{a \in \text{Ag}}$, we have $\delta'(v^\theta, \mathbf{c}') = \delta(v, (c_a)_{a \in \text{Ag}})$;
- $\ell' : V' \times \text{AP} \rightarrow [-1, 1]$ is defined such that for all $v \in V, p \in \text{AP}$, we have $\ell'(v, p) = 0$ and for all $v^\theta \in V^\Theta$, we have

$$\ell'(v^\theta, p) = \begin{cases} \theta(p), & p \in U \\ \ell(v, p), & p \notin U. \end{cases}$$

Moreover, for all $v \in V'$, we have

$$\ell'(v, r) = \begin{cases} -1, & v \in V \\ 1, & v \in V^\Theta; \end{cases}$$

- $\gamma'_a = \text{tr}^2(\gamma_a)$.

This translation of goals is well-defined due to the addition of the r variable which is used to r -label plays in \mathcal{G}' . Intuitively, \mathcal{G}' introduces a new player – the ‘incentive player’ – whose actions correspond to the application of incentives at each round of the original game. The incentive player’s actions are taken in an alternating manner with the rest of the players in the positions V , which are interleaved with the new positions V^Θ such that at every odd timestep $t = 2k + 1, k \in \mathbb{N}$, the game will be in a position $v_t \in V$, whereas at every positive even timestep $t' = 2k, k \in \mathbb{Z}^+$, the game will be in a position $v_{t'}^\theta \in V^\Theta$. In this way, at positions $v \in V$, only the incentive player’s action determines the next position and at positions $v^\theta \in V^\Theta$, only the original players’ actions (i.e., players 1, ..., n) will determine the next position.

Due to the definition of the modified valuation function ℓ' and the 2-fold translation of the players’ objectives, the incentive schemes chosen by the incentive player at each of their turns modify the satisfaction values of the players’ objectives, influencing the players’ strategic considerations.

Moreover, it is straightforward to see that once the incentive player’s strategy has been fixed, this induces a new game with potentially different stable outcomes from the original game. Under this construction, the task of the incentive player can then be formulated as the problem of choosing a strategy such that some or all of the induced equilibria of the resulting game for the remaining players maximises the satisfaction value of the incentive designer’s objective. The following result establishes the relationship between strategies for the incentive player in an incentive-augmented game and the application of dynamic incentive schemes in the original game.

Proposition 2. *Suppose that \mathcal{G} is a weighted concurrent game, \mathcal{G}' is its incentive-augmented version, φ is an LTL $[\mathcal{F}]$ formula, and (σ', T) is a strategy profile in \mathcal{G}' . Then, there*

is a strategy profile σ in \mathcal{G} and a dynamic incentive scheme T such that

$$\llbracket \varphi \rrbracket_{A[\text{Ag} \mapsto (\sigma)]}^{\mathcal{G}_T} = \llbracket \text{tr}^2(\varphi) \rrbracket_{A[\text{Ag}' \mapsto (\sigma', \mathcal{T})]}^{\mathcal{G}'}$$

Proof. We prove the claim by constructing the strategy profile σ and the dynamic incentive scheme T . Firstly, let σ be the strategy profile in \mathcal{G} such that for all even-length histories $h' = v_1 v_1^{\theta_1} v_2 v_2^{\theta_2} \dots v_k v_k^{\theta_k}$, $k \in \mathbb{Z}^+$ that are generated by σ' in \mathcal{G}' , it holds that $\sigma(h) = \sigma'(h')$, where $h = v_1 v_2, \dots, v_k$. This strategy profile will be referred to as the *2-fold deflation* of the strategy profile σ' , which extracts the actions that players take at every second round of the game \mathcal{G}' . Next, let T be the dynamic incentive scheme such that for any given history $h' = v_1 v_1^{\theta_1} \dots v_{k-1}^{\theta_{k-1}} v_k$ of a play π' in \mathcal{G}' and the corresponding history $h = v_1 v_2 \dots v_k$ of a play π in \mathcal{G} resulting from the execution of σ , we have that $\mathcal{T}(h') = T(h)$. By the construction of the game \mathcal{G}' , we see that the incentive scheme θ_k corresponding to the chosen action $\mathcal{T}(h')$ is implemented in position $v_k^{\theta_k}$, that is, for all $p \in \text{AP}$, we have $\ell'(v_k^{\theta_k}, p) = \vec{\ell}_T(h) = \ell_{\theta_k}$. In other words, the satisfaction values for all variables on round $2k$, $k \in \mathbb{Z}^+$ of the game \mathcal{G}' coincide with those on round k of the game \mathcal{G} .

Now, let π' be the play generated in \mathcal{G}' by (σ', \mathcal{T}) , $\vec{\ell}' = \ell' \ell' \dots$, π be the play generated in \mathcal{G} by (σ) , and $\vec{\ell}_T(\pi) = \ell_{\theta_1} \ell_{\theta_2} \dots$. Then, we can conclude that the $(\text{AP}', \vec{\ell}')$ -labelled play $L_2 = L(\pi'; \text{AP}', \vec{\ell}')$ in \mathcal{G}' is an r -labelled d -fold inflation of $L_1 = L(\pi; \text{AP}, \vec{\ell}_T(\pi))$ in \mathcal{G} and hence, the result follows by the application of Lemma 1. \square

Finally, we turn our attention to the dynamic incentive verification and synthesis problems. Again, we are able to make use of the expressive power of $\text{SL}[\mathcal{F}]$ to define a suitable formula which can be model-checked to determine the solution to our problems.

Theorem 3. For $\zeta \in \{DSE, NE, RE_m\}$, $m \in \{1, \dots, n\}$, ζ -D-E-INCENTIVE-VERIFICATION and ζ -D-A-INCENTIVE-VERIFICATION are 2EXPTIME-complete.

Proof. Regarding ζ -D-E-INCENTIVE-VERIFICATION, we use the construction \mathcal{G}' , which introduces the incentive designer as a player in the game, whose actions are interleaved with those of the original players Ag . Given this, we can solve the incentive verification problem by model checking the following $\text{SL}[\mathcal{F}]$ formula in the game \mathcal{G}' :

$$(n+1, \mathcal{T}) [\exists \sigma'. [\zeta(\sigma') \wedge (\text{Ag}, \sigma') \text{tr}^2(\varphi)]] ,$$

where the strategy \mathcal{T} is defined such that for all histories $h = v_1 v_2 \dots v_k \in \text{Hist}$ of the original game \mathcal{G} and a corresponding history $h' = v_1 v_1^{\theta_1} v_2 v_2^{\theta_2} \dots v_k^{\theta_k}$, we have that $\mathcal{T}(h') = T(h)$. By the construction of \mathcal{G}' , strategies for the original players in Ag are inconsequential on odd timesteps, and moreover, the formula $\zeta(\sigma')$ is now defined in terms of the 2-fold translated versions of players' goals. Applying the definition of the 2-fold translation of

Thus, by Proposition 2, $\zeta(\sigma')$ has a satisfaction value of 1 in the above formula if and only if the 2-fold deflation of σ' in the original game \mathcal{G}_T is a ζ -equilibrium.

Moreover, by Proposition 2 again, the satisfaction value of the subformula $(\text{Ag}, \sigma') \text{tr}^2(\varphi)$ in \mathcal{G}' is equal to the satisfaction value of the formula $(\text{Ag}, \sigma) \varphi$ in the game \mathcal{G}_T . Hence, model-checking this formula indeed solves the ζ -D-E-INCENTIVE-VERIFICATION problem. For all of our solution concepts, this formula has an alternation depth of 1, and hence, the formula can be model-checked in 2EXPTIME.

Similarly, for ζ -D-A-INCENTIVE-VERIFICATION, we model check the formula:

$$(n+1, \mathcal{T}) [\forall \sigma'. [\zeta(\sigma') \rightarrow (\text{Ag}, \sigma') \text{tr}^2(\varphi)]] ,$$

which again has an alternation depth of 1 when converted into prenex normal form.

For hardness, we can apply the same reduction as used in the proof of Proposition 1, i.e., we reduce from qualitative rational synthesis and its strong variant to the ζ -D-E-INCENTIVE-VERIFICATION and ζ -D-A-INCENTIVE-VERIFICATION problems respectively. Again, we check a dynamic incentive scheme that does not modify any of the variables and use player 0's goal as the global formula to be checked, with a threshold $c = 1$. \square

As with static incentive schemes, the incentive verification problem for dynamic incentive schemes remains in 2EXPTIME. However, the following result indicates that we do not necessarily obtain the same complexity when moving from static to dynamic incentive synthesis.

Theorem 4. For $\zeta \in \{DSE, NE, RE_m\}$, $m \in \{1, \dots, n\}$, ζ -D-E-INCENTIVE-SYNTHESIS is 2EXPTIME-complete and ζ -D-A-INCENTIVE-SYNTHESIS is in 3EXPTIME and is 2EXPTIME-hard.

Proof. Starting with the existential version of the problem, we again make use of the construction \mathcal{G}' . The ζ -D-E-INCENTIVE-SYNTHESIS problem can be solved by model-checking the following $\text{SL}[\mathcal{F}]$ formula:

$$\exists \mathcal{T}. [(n+1, \mathcal{T}) \exists \sigma'. [\zeta(\sigma') \wedge (\text{Ag}, \sigma') \text{tr}^2(\varphi)]] .$$

By a similar argument as in the proof of Theorem 3, model-checking this formula solves the ζ -D-E-INCENTIVE-SYNTHESIS, and we can extract the dynamic incentive scheme T from a witness \mathcal{T} to the outer existential quantifier of the above formula using the method outlined in the proof of Proposition 2. Noting that the quantifiers over \mathcal{T} and σ' do not introduce another level of alternation, model checking this formula is also in 2EXPTIME.

For ζ -D-A-INCENTIVE-SYNTHESIS, we cannot group the quantifiers for incentive schemes and the candidate ζ -equilibria, which introduces another level of alternation in the representative $\text{SL}[\mathcal{F}]$ formula:

$$\exists \mathcal{T}. [(n+1, \mathcal{T}) \forall \sigma'. [\zeta(\sigma') \rightarrow (\text{Ag}, \sigma') \text{tr}^2(\varphi)]] .$$

This leads to a 3EXPTIME procedure for the universal incentive synthesis problem. For the lower bounds, one can use the same reduction as in Theorem 2, where we set up the game so that the incentive designer has no effect on the rational outcomes of the original rational synthesis problem. The addition of memory for the incentive designer does not impact their capabilities in this setup, so the arguments for the reduction carry through. \square

6 Conclusion

We have introduced an approach to designing incentives for rational agents with $LTL[\mathcal{F}]$ goals in order to achieve a societal objective. We formalized the problems of verifying and synthesizing static and dynamic incentive schemes for a range of solution concepts. We then proved that, for all solution concepts that we consider, the complexity of incentive verification is 2EXPTIME-complete, which is not harder than the corresponding qualitative rational verification problems. In practice, verification is double exponential only in the sizes of the formulae representing the societal objectives, which are typically small for most applications. In the case of incentive synthesis, we showed it is 2EXPTIME-complete in the static case and in 3EXPTIME in the dynamic case.

Similar to the rational verification problem, our static incentive synthesis has double exponential cost only in the size of the goals. However, there is also an additional exponential cost in the size of the representation of the granularity, which can be chosen by the incentive designer for their purposes. In practical applications such as setting interest rates or fixing minimum bidding increments in auctions, incentive designers typically use a rather coarse level of granularity, which would not incur a significant blowup in the algorithm's worst-case runtime. For dynamic incentive synthesis, there is a triple exponential cost in the size of the goals, and we again have an exponential cost in the size of the granularity, but the cost is the same as in verification for the remaining parts of the input. Thus, even with a potentially exponential gap between static and dynamic incentive synthesis, for relatively succinct goal specifications, there may not be a significant difference in runtimes between the two approaches.

For future work, we intend to study the synthesis of incentive schemes with minimal modifications to the original game and to analyze Incentive Design in the cooperative setting. Due to the expressivity of $SL[\mathcal{F}]$, many additional requirements or constraints that one may wish to impose on the incentive schemes can be naturally incorporated into the global goal in our framework, which enables a wide variety of incentive design problems to be solved using the approaches outlined here.

Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101105549, the PNRR MUR project PE0000013-FAIR, and the PRIN 2020 projects PINPOINT and RIPER (No. 20203FFYLK). Giuseppe Perelli was also supported by Sapienza University of Rome under the "Progetti Grandi di Ateneo" programme, grant RG123188B3F7414A (AS-GARD - Autonomous and Self-Governing Agent-Based Rule Design). Michael Wooldridge was supported by the UKRI under a Turing AI World Leading Research Fellowship (EP/W002949/1).

References

- Abraham, I.; Dolev, D.; and Halpern, J. Y. 2008. Lower bounds on implementing robust and resilient mediators. In *Theory of Cryptography: Fifth Theory of Cryptography Conference*, 302–319. Springer.
- Alechina, N.; De Giacomo, G.; Logan, B.; and Perelli, G. 2022. Automatic synthesis of dynamic norms for multi-agent systems. In *Proc. of the Int. Conference on Principles of Knowledge Representation and Reasoning*, volume 19, 12–21.
- Almagor, S.; Avni, G.; and Kupferman, O. 2015. Repairing multi-player games. In *Proc. of the Int. Conference on Concurrency Theory, CONCUR 2015*, volume 42 of *LIPICs*, 325–339. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Almagor, S.; Boker, U.; and Kupferman, O. 2016. Formally reasoning about quality. *Journal of the ACM* 63(3):24:1–24:56.
- Almagor, S.; Kupferman, O.; and Perelli, G. 2018. Synthesis of controllable nash equilibria in quantitative objective game. In *Proc. of the Int. Joint Conference on Artificial Intelligence, IJCAI 2018*, 35–41. ijcai.org.
- Alur, R.; Henzinger, T.; and Kupferman, O. 2002. Alternating-time temporal logic. *Journal of the ACM* 49(5):672–713.
- Baier, C., and Katoen, J. 2008. *Principles of model checking*. MIT Press.
- Belardinelli, F.; Jamroga, W.; Malvone, V.; Mittelman, M.; Murano, A.; and Perrussel, L. 2022. Reasoning about human-friendly strategies in repeated keyword auctions. In *Proc. of the Int. Conference on Autonomous Agents and Multi-Agent Systems, AAMAS 2022*, 62–71. IFAAMAS.
- Bouyer, P.; Kupferman, O.; Markey, N.; Maubert, B.; Murano, A.; and Perelli, G. 2023. Reasoning about Quality and Fuzziness of Strategic Behaviors. *ACM Trans. Comput. Log.* 24(3):21:1–21:38.
- Brice, L.; Raskin, J.; and van den Bogaard, M. 2023. Rational Verification for Nash and Subgame-Perfect Equilibria in Graph Games. In Leroux, J.; Lombardy, S.; and Peleg, D., eds., *Proc. of the Int. Symposium on Mathematical Foundations of Computer Science, MFCS 2023*.
- Centeno, R., and Billhardt, H. 2011. Using incentive mechanisms for an adaptive regulation of open multi-agent systems. In *Proc. of the Int. Joint Conference on Artificial Intelligence, IJCAI 2011*.
- Condurache, R.; Filiot, E.; Gentilini, R.; and Raskin, J. 2016. The Complexity of Rational Synthesis. In *Int. Colloquium on Automata, Languages, and Programming, ICALP 2016*.
- Filiot, E.; Gentilini, R.; and Raskin, J. 2018. Rational Synthesis Under Imperfect Information. In Dawar, A., and Grädel, E., eds., *Proc. of the Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018*.
- Fisman, D.; Kupferman, O.; and Lustig, Y. 2010. Rational synthesis. In Esparza, J., and Majumdar, R., eds., *Tools and Algorithms for the Construction and Analysis of Systems*.

- Gutierrez, J.; Perelli, G.; and Wooldridge, M. J. 2018. Imperfect information in reactive modules games. *Information and Computation* 261:650–675.
- Hyland, D.; Gutierrez, J.; and Wooldridge, M. 2023a. Principal-agent boolean games. In *Proc. of the Int. Joint Conference on Artificial Intelligence, IJCAI 2023*, 144–152. ijcai.org.
- Hyland, D.; Gutierrez, J.; and Wooldridge, M. J. 2023b. Incentive engineering for concurrent games. In *Proc. of the Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2023*, volume 379 of *EPTCS*, 344–358.
- Kupferman, O., and Shenwald, N. 2022. The Complexity of LTL Rational Synthesis. In Fisman, D., and Rosu, G., eds., *Int. Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS 2022*. Springer.
- Kupferman, O.; Perelli, G.; and Vardi, M. Y. 2016. Synthesis with rational environments. *Annals of Mathematics and Artificial Intelligence* 78(1):3–20.
- Kučera, A., and Strejček, J. 2005. The stuttering principle revisited. *Acta Informatica* 41(7–8):415–434.
- Laffont, J.-J., and Martimort, D. 2009. The theory of incentives: the principal-agent model. In *The theory of incentives*. Princeton University Press.
- Maubert, B.; Mittelman, M.; Murano, A.; and Perrussel, L. 2021. Strategic reasoning in automated mechanism design. In *Proc. of the Int. Conference on Knowledge Representation and Reasoning, KR 2021*.
- Mittelman, M.; Maubert, B.; Murano, A.; and Perrussel, L. 2022. Automated synthesis of mechanisms. In *Proc. of the International Joint Conference in Artificial Intelligence, IJCAI 2022*, 426–432. ijcai.org.
- Narayanasamy, S. K. 2022. Automating mechanism design with program synthesis. In *Proc. Automated Learning Agents Workshop*.
- Nisan, N.; Roughgarden, T.; Tardos, É.; and Vazirani, V. 2007. *Algorithmic Game Theory*. Cambridge University Press.
- Olsder, G. 2009. Phenomena in inverse stackelberg games, part 1: Static problems. *Journal of optimization theory and applications* 143:589–600.
- Parkes, D. C.; Cavallo, R.; Constantin, F.; and Singh, S. 2010. Dynamic incentive mechanisms. *AI Magazine* 31(4):79–94.
- Pnueli, A. 1977. The temporal logic of programs. In *Proc. of the Annual Symposium on Foundations of Computer Science, SFCS 1977*, 46–57. IEEE.
- Ratliff, L. J., and Fiez, T. 2020. Adaptive incentive design. *IEEE Transactions on Automatic Control* 66(8):3871–3878.
- Ratliff, L. J.; Dong, R.; Sekar, S.; and Fiez, T. 2019. A perspective on incentive design: Challenges and opportunities. *Annual Review of Control, Robotics, and Autonomous Systems* 2:305–338.
- Weber, T. A. 2011. *Optimal control theory with applications in economics*. MIT press.
- Wooldridge, M. J.; Endriss, U.; Kraus, S.; and Lang, J. 2013. Incentive engineering for boolean games. *Artificial Intelligence* 195:418–439.
- Yang, J.; Wang, E.; Trivedi, R.; Zhao, T.; and Zha, H. 2022. Adaptive incentive design with multi-agent meta-gradient reinforcement learning. In *Proc. of the Int. Conference on Autonomous Agents and Multiagent Systems, AAMAS 2022*.