

Adaptive Incentive Design with Multi-Agent Meta-Gradient Reinforcement Learning

Jiachen Yang¹, Ethan Wang¹, Rakshit Trivedi², Tuo Zhao¹ & Hongyuan Zha³

¹Georgia Institute of Technology, ²Harvard University, ³Chinese University of Hong Kong, Shenzhen
 jyang462, ewang78, tourzhao@gatech.edu, rstrivedi@seas.harvard.edu, zhahy@cuhk.edu.cn

ABSTRACT

Critical sectors of human society are progressing toward the adoption of powerful artificial intelligence (AI) agents, which are trained individually on behalf of self-interested principals but deployed in a shared environment. Short of direct centralized regulation of AI, which is as difficult an issue as regulation of human actions, one must design institutional mechanisms that indirectly guide agents' behaviors to safeguard and improve social welfare in the shared environment. Our paper focuses on one important class of such mechanisms: the problem of adaptive incentive design, whereby a central planner intervenes on the payoffs of an agent population via incentives in order to optimize a system objective. To tackle this problem in high-dimensional environments whose dynamics may be unknown or too complex to model, we propose a model-free meta-gradient method to learn an adaptive incentive function in the context of multi-agent reinforcement learning. Via the principle of online cross-validation, the incentive designer explicitly accounts for its impact on agents' learning and, through them, the impact on future social welfare. Experiments on didactic benchmark problems show that the proposed method can induce selfish agents to learn near-optimal cooperative behavior and significantly outperform learning-oblivious baselines. When applied to a complex simulated economy, the proposed method finds tax policies that achieve better trade-off between economic productivity and equality than baselines, a result that we interpret via a detailed behavioral analysis.

KEYWORDS

Incentive Design; Multi-Agent Reinforcement Learning

ACM Reference Format:

Jiachen Yang¹, Ethan Wang¹, Rakshit Trivedi², Tuo Zhao¹ & Hongyuan Zha³. 2022. Adaptive Incentive Design with Multi-Agent Meta-Gradient Reinforcement Learning. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 10 pages.

Madness is rare in individuals—but in groups, parties, nations, and ages, it is the rule.

Nietzsche [32]

1 INTRODUCTION

As advances in artificial intelligence research drive the growing presence of AI in critical infrastructures—such as transportation [6, 7], information communication [16], financial markets [17] and

agriculture [26]—it is increasingly important for AI research to complement the solipsistic view of models and agents acting in isolation with a broader viewpoint: these models and agents may be developed by independent and self-interested principals but are eventually deployed in a shared multi-agent ecosystem. Apart from the special cases of pure common or conflicting interest (i.e., team or zero-sum payoffs), the majority of possible scenarios involve mixed motives [42] whereby cooperation for optimal group payoff is attainable if selfish behavior out of greed or fear—which results in low individual and group payoff—can be overcome. Research on endowing individual agents with social capabilities and designing central institutional mechanisms to safeguard and improve social welfare, recently named “Cooperative AI” [9], is a long-term necessity that requires present-day research efforts.

We focus on one specific pillar of this research agenda: the problem of *adaptive incentive design* [40], whereby a central institution shapes the behavior of self-interested agents to improve social welfare, by introducing an incentive function to modify their individual payoffs. The trivial solution of setting individual payoffs equal to the average system payoff is known to be suboptimal [37]: e.g., uniform redistribution of income leads to low productivity. This problem can be formalized as a reverse Stackelberg game [15, 47], in which the leader first proposes a function (e.g., the incentive function) that maps from the follower's action space to the leader's decision space, while the follower chooses a best response. This is a difficult bi-level optimization problem even in the linear case [14]. To tackle this problem in high-dimensional multi-agent systems, we propose a model-free method based on meta-gradient reinforcement learning [59] and the principle of online cross-validation [49], for the incentive designer (ID) to explicitly account for the learning dynamics of agents in response to incentives. Potential applications in the long term include: e.g. shaping consumption patterns to improve the efficiency of smart power grids [54] and to mitigate climate crisis [3]; reducing wait times or traffic congestion in taxi dispatch and traffic tolling systems [11, 31]; solving the social dilemma of autonomous vehicle adoption [7]; improving the trade-off between economic productivity and income equality via taxation [61].

In the spirit of complexity economics [2], whereby the tractability of linear models with analytic equilibria [13, 37, 41] is eschewed for the greater realism and richer dynamics of high-dimensional agent-based simulation, we work in the framework of Markov games [27] with reinforcement learning (RL) agents [50]. Within this context, we interpret “incentive design” in the broad sense of influencing agents' behaviors via modifying their individual payoffs. The issue of incentive compatibility, despite being central to analytically tractable applications such as auctions where the goal is to elicit truthful reporting of private valuations [40], do not pertain in general to complex simulations involving RL agents for the following

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

reasons: 1) the concept of private individual preferences may not make sense in the application (e.g., a social dilemma whose payoff is known to all parties); 2) there may be no *a priori* or *fixed* private valuations, because an agent’s preference is completely represented by its reward function, which itself depends on the incentive function and hence changes dynamically along with the ID’s online optimization process; and 3) a complex simulation involving nonlinear processes and discrete rules is used to investigate dynamical behavior rather than to converge to equilibria.

Our use of simulation is motivated by two considerations. Firstly, there may arise a future ecosystem of AI in diverse spaces, such as multiple firms in a financial market [1] and multiple recommendation systems in the same consumer sector [29]. With the increasing success of RL on progressively more difficult tasks [5, 46, 56] and growing efforts to apply RL to real-world problems [19, 22], such real-world *in silico* AI are likely to involve RL for optimization of long-term objectives and will be ontologically equivalent to the entities in our work. Secondly, agent-based simulation is also relevant to incentive design for a group of self-interested humans, firms, or states, by viewing the reward-maximizing behavior of RL agents as an approximation of the bounded rationality of such real world entities [33]. Therefore, we validate our approach in existing simulation benchmarks: *Escape Room* [60] and *Cleanup* [18], which are deceptively hard but easily interpretable benchmark problems, and the complex economic simulation called *Gather-Trade-Build* [61], where agent behavior such as specialization and correlations between taxation and productivity are qualitatively in accord with results from economic theory and reality.

The centralized incentive designer (ID) in our work is related but orthogonal to the centralization commonly seen in the literature on cooperative MARL [36, 38]: centralized training in MARL permits a global entity to update each agent’s individual policy parameters using shared global information and thereby directly optimize a single team reward, while the ID in our work can only set an incentive function to affect each agent’s total individual reward. Given a fixed incentive function, our agents live in a fully-decentralized POMDP [35]. The ID’s intervention on the agents’ reward function also differs from the insertion of controlled agents (e.g., centrally-controlled autonomous vehicles) whose actions indirectly regulate the behavior of other individuals (e.g. human drivers) [57].

In short, the algorithmic and experimental contributions of this paper are: (1) we propose a meta-gradient approach for the reverse Stackelberg game formulation of incentive design in high-dimensional multi-agent reinforcement learning settings; (2) we explain and show experimentally that using standard RL for the incentive designer faces significant difficulties even small finite state finite action Markov games; (3) our proposed method can converge to known global optima in standard benchmark problems, and it generates significantly higher social welfare than the previous state-of-the-art in a complex high-dimensional economic simulation of market dynamics with taxation.

2 RELATED WORK

A large body of previous work on incentive, utility, and mechanism design belongs to the analytic paradigm, which faces limitations such as linear agent cost and planner incentive functions [37, 41], finite single-round games [23], state-based potential games [25],

or pertain to special problems such as welfare distribution [13] or seller-buyer auctions [45]. These simplifications result in lower applicability to complex, nonlinear, and temporally-extended environments such as dynamic economies [61]. In contrast, we adopt the paradigm of agent-based simulation [53] and take state-of-the-art agent learning methods (i.e., deep reinforcement learning, at present) as the starting point to inform a method for incentive design, at the cost of discarding analytical tractability.

Previous work on incentive or mechanism design with RL differ from ours in the choice of the algorithm for the incentive designer or the model of agents. Brero et al. [8], Tang [52] apply RL to the upper-level planner for non-RL agents. Pardoe et al. [39] use perturbation-based gradient ascent to search for hyperparameters of a k -armed bandit algorithm that determines the parameters of an auction. Mguni et al. [30] employ Bayesian optimization and treat the lower-level multi-agent RL as a black-box. The central planner in Baumann et al. [4] optimizes social welfare in 2-player matrix games by anticipating the players’ one-step updates. Li et al. [24] assume that agents’ total payoff is a continuous and differentiable function of the joint strategy—which does not hold in general if agents’ original reward can be any combination of discrete and nondifferentiable rules—and differentiate through the variational inequality reformulation of Nash equilibria. The closest work to ours are Yang et al. [60], where fully-decentralized RL agents learn mutual pairwise incentivization, and Danassis et al. [10], Zheng et al. [61], where a central RL planner optimizes an adaptive tax or price policy at the same time-scale as RL agents’ policy optimization.

The technical aspect of our method builds on single-agent meta-gradient RL [59] and discovering intrinsic rewards [62], which we extend to the multi-agent setting and refine with the principle of online cross-validation [49]. Related to but different from the variety of existing single-agent meta-learning methods enumerated in Xu et al. [58, Table 1], our method learns a general neural network representation of an incentive function within a single lifetime, as opposed to methods that optimize hyperparameters [28, 48, 59], learn target functions [58], or use multiple lifetimes over different environments to find general update functions [20, 34].

3 METHOD

We propose a method, called “MetaGrad”, for an Incentive Designer (ID) to optimize a measure of social welfare by explicitly accounting for the impact of incentives on the behavior of a population of n independent agents. Each agent $i \in \{1, \dots, n\}$ has an individual reward function $R^i: \mathcal{S} \times \mathcal{A}^n \times \mathcal{U} \rightarrow \mathbb{R}$, which depends not only on the global state $s \in \mathcal{S}$ and the joint action $\mathbf{a} \in \mathcal{A}^n$, as in standard Markov games [27] with transition function $P(s'|s, \mathbf{a})$, but also on an incentive that is parameterized by $u \in \mathcal{U}$ for some bounded set $\mathcal{U} \subset \mathbb{R}^l$ (e.g., a vector of marginal tax rates). We assume that R^i is differentiable with respect to the argument u —this holds in the common case of additive incentives such as highway tolling, as well as in complex mechanisms such as a bracketed tax schedule (7). This incentive u is generated by the ID via a learned incentive function $\mu_\eta: \mathcal{S} \times \mathcal{A}^n \rightarrow \mathcal{U}$, parameterized by η , which adaptively responds to the current system state and joint action of the agents. Each agent i independently learns a policy π_{θ^i} , parameterized by θ^i , to optimize its own individual expected return, while the ID’s

performance is measured by a social welfare reward $R^{\text{ID}}(s, \mathbf{a})$. Let π and θ denote the agents' joint policy and policy parameters, respectively. For brevity, we use $R_\eta^i(s, \mathbf{a})$ to denote $R^i(s, \mathbf{a}, \mu_\eta(s, \mathbf{a}))$, and we identify θ with the policy π_θ where no confusion arises.

The ID aims to solve the bilevel optimization problem

$$\text{(ID)} \quad \max_{\eta} J^{\text{ID}}(\eta; \hat{\theta}) := \mathbb{E}_{\pi_{\hat{\theta}}, P(s^t | s_t, \mathbf{a})} \left[\sum_{t=0}^H \gamma^t R_t^{\text{ID}} - \psi(s_t) \right] \quad (1)$$

$$\text{(Agents)} \quad \hat{\theta} = \left\{ \operatorname{argmax}_{\theta^i} J^i(\theta; \eta) := \mathbb{E}_{\pi_{\theta}} \left[\sum_{t=0}^H \gamma^t R_\eta^i(s_t, \mathbf{a}_t) \right] \right\}_{i=1}^n \quad (2)$$

over episode horizon H , and ψ is a known cost for incentivization.

We assume that the n agents apply a gradient-based reinforcement learning procedure $\text{RL}(\cdot)$ to update their policies in response to the reward determined by η , i.e., that (2) is achieved by $\hat{\theta}^i = \text{RL}(\theta_0^i; \eta)$, where θ_0^i is an initial policy. In the ideal case, one should measure the population behavior under the final joint policy $\hat{\theta}$, after convergence of the RL process under a fixed η , to evaluate the effectiveness of η in optimizing social welfare and conduct a single η update. However, the high sample count required for convergence of RL in practice is prohibitively expensive, especially if one wishes to apply gradient descent to optimize η . To tackle this challenge, we build on the effectiveness of meta-gradient RL in optimizing hyperparameters and parameterized objectives concurrently with an agent's policy optimization [34, 58, 59]. We apply iterative gradient descent to the upper objective (1) on the same timescale as the agents' policy optimization (2), by explicit differentiating through the agents' policy updates. We emphasize that even though the ID does not wait for convergence of the final $\hat{\theta}$, we follow the principle of online cross-validation [49] and extend it to the optimal control setting: the data used for the η -update is still generated by the agents' *updated* policies, not by any arbitrary policy, in order to measure accurately the impact of η on the ID's objective through the agents' learning process. This differs from previous single-agent meta-gradient RL [58, 59], where the trajectories used for the outer update were not generated by the updated agent policy.

Specifically, we implement the following algorithm (Algorithm 1). Given the current policy θ_0 and incentive function μ_η , agents collect trajectories $\{\{\tau_j^i\}_{i=1}^n\}_{j=0}^{M-1} \sim \pi_{\theta_0}$ (Algorithm 1, line 3) and conduct M policy update steps (Algorithm 1, line 5):

$$\hat{\theta}(\eta) := \theta_M = \theta_0 + \sum_{j=0}^{M-1} \Delta \theta_j(\eta). \quad (3)$$

Each agent's update $\Delta \theta_j^i(\eta) \propto \nabla_{\theta^i} J^i(\theta_j^i; \eta, \tau_j^i)$ depends on the fixed η . For example, if agents learn by policy gradient methods, we have $\Delta \theta_j^i \propto \sum_{t=0}^T \nabla_{\theta^i} \log \pi_{\theta^i}(a_t^i | o_t^i) A_t^i(\tau_j^i; \eta)$, where A_t^i is an advantage function that depends on η via $R_\eta^i(s, \mathbf{a})$. Now let $\hat{\tau}$ denote the subsequent trajectory generated by the agents' updated policies $\hat{\theta}$ (Algorithm 1, line 6), which serves as the *validation trajectory* that measures the indirect impact of η on the ID's return through the agents' learning. The ID computes and ascends the gradient of objective (1) w.r.t. η via the chain rule (Algorithm 1, line 7)

$$\nabla_{\eta} J^{\text{ID}}(\eta; \hat{\theta}, \hat{\tau}) = \sum_{i=1}^n \left(\nabla_{\eta} \hat{\theta}^i(\eta) \right)^{\top} \left(\nabla_{\hat{\theta}^i} J^{\text{ID}}(\eta; \hat{\theta}, \hat{\tau}) \right) - \nabla_{\eta} \psi(\tau) \quad (4)$$

Algorithm 1 Meta-Gradient Incentive Design with pipelining

```

1: procedure
2:   Initialize all agents' policy parameters  $\theta^i$ , incentive function
   parameters  $\eta$ 
3:   Generate trajectory  $\tau$  using  $\theta$  and  $\eta$ 
4:   for each iteration do
5:     For all agents, update  $\hat{\theta}^i$  with  $\tau$  using (3)
6:     Generate a new trajectory  $\hat{\tau}$  using new  $\hat{\theta}$ 
7:     Update  $\hat{\eta}$  by gradient ascent along (5) using  $\tau$  and  $\hat{\tau}$ 
8:      $\tau \leftarrow \hat{\tau}$ ,  $\eta \leftarrow \hat{\eta}$ ,  $\theta^i \leftarrow \hat{\theta}^i$  for all  $i \in [n]$ .
9:   end for
10: end procedure

```

where $\nabla_{\eta} \hat{\theta}^i$ is computed using a replica of the θ^i update step.

Proximal meta-gradient optimization. Instead of computing both factors of (4), we can view (1) as a standard objective in policy-based RL, with the difference that we optimize with respect to η instead of the policy parameters $\hat{\theta}$. Hence, one can apply the policy gradient algorithm [51] by replacing $\nabla_{\hat{\theta}}$ with ∇_{η} , as shown in Yang et al. [60, Appendix C] and used implicitly by Xu et al. [58, 59]. We extend this viewpoint by showing (in Appendix A.1) that trust-region arguments [43] hold for meta-gradients, which justifies the use of a proximal policy optimization (PPO)-type gradient [44] for the outer optimization:

$$\nabla_{\eta} J^{\text{ID}}(\eta; \hat{\theta}, \hat{\tau}) = \mathbb{E}_{\mathbf{a}_t, s_t \sim \pi_{\hat{\theta}(\eta)}} \left[\min \left(r(\hat{\theta}; \eta) A_t, \operatorname{clip}(r(\hat{\theta}; \eta), 1 - \epsilon, 1 + \epsilon) A_t \right) \right] \quad (5)$$

$$r(\hat{\theta}; \eta) := \frac{\nabla_{\eta} \pi_{\hat{\theta}(\eta)}(\mathbf{a}|s)}{\pi_{\hat{\theta}(\eta)}(\mathbf{a}|s)}, \quad (6)$$

where $A_t := \sum_{l=t}^{T-1} (\gamma \lambda)^{l-t} \delta_l$ is a generalized advantage estimator computed using $\delta_t := R^{\text{ID}}(s_t, \mathbf{a}_t) + \gamma V(s_{t+1}) - V(s_t)$, critic V for the ID, discount γ and λ -returns.

3.1 Technical Relation to Prior Multi-Agent Learning Methods for Incentivization

The ‘‘AI Economist’’ [61] treats agents' learning as a black-box: it applies standard RL to a central planner who learns an adaptive tax policy concurrently with the agents' policy learning within a fully-decentralized multi-agent economy. Rather than addressing the bi-level optimization problem, this expands the multi-agent system and exacerbates the already existing problem of non-stationarity in decentralized MARL, which required heuristics such as curriculum learning and tax annealing that are difficult to tune. In contrast, our method to train the incentive function fundamentally differs from standard RL: the gradient (5) is taken with respect to the η variables of the incentive function, *through the policy updates* $\hat{\theta} \leftarrow \theta + \Delta \theta$ of all the regular RL agents, where $\hat{\theta}^i$ preserves the dependence of each agent's update on η in the computational graph.

Our technical method is the centralized analogue of the fully-decentralized pairwise incentivization in LIO [60]. That work begins with the premise that all, or some, agents in the environment are equipped with the LIO learning mechanism, but this may not hold

in general environments where no principal opts to use a LIO agent. In contrast, our work only assumes that agents learn from reward functions that depend on and can be differentiated with respect to incentives, which pertains to more general potential applications.

4 EXPERIMENTAL SETUP

We evaluated our approach in three environments: 1) *Escape Room* (ER) [60], a small but deceptively hard pedagogical example that accentuates the core challenges of incentive design for RL agents; 2) *Cleanup* [18], a high-dimensional instance of a sequential social dilemma; and 3) the *Gather-Trade-Build* simulation of a market economy with taxation, trading, and competition for limited resource. We conducted eight independent runs per method for *Escape Room*, 10 for *Cleanup*, and four for *Gather-Trade-Build*. Section 4.1 summarizes the high-level features of these environments; Appendix B.1 provides complete specifications. Section 4.2 describes the implementation of the method and baselines.

4.1 Environments

Escape Room [60]. The *Escape Room* game $ER(n, m)$ is a discrete n -player Markov game with individual extrinsic rewards and parameter $m < n$. An agent gets +10 extrinsic reward for exiting a door, but this requires m other agents to sacrifice their own self-interest by pulling a lever at a cost of -1 each. We fix all agents to be standard independent RL agents without “give-reward” capabilities, and we introduce a central incentive designer who can modify each agent’s reward by adding a scalar bounded in $(0, 2)$, since an incentive value of $1 + \epsilon$ for $\epsilon > 0$ is sufficient for an agent to overcome the -1 penalty for pulling the lever. The ID’s reward R^{ID} is the sum of all agents’ rewards, and the cost of incentivization ψ is the sum of all scalar incentives. Hence, the global optimum reward for the incentive designer is $10(n - m) - m - m(1 + \epsilon)$.

Cleanup [18]. The *Cleanup* scenario is a high-dimensional grid-world sequential social dilemma that serves as a difficult benchmark for independent learning agents. Agents get +1 individual reward by collecting “apples”, whose spawn rate decreases in proportion to the amount of “waste”. Each agent can contribute to the public good by firing a cleaning beam to clear waste, but doing so would enable other agents to defect and selfishly collect apples, hence posing a difficult social dilemma. We extended the open-source implementation [55] to provide a global observation image for the incentive designer.

Gather-Trade-Build (GTB) [61]. The GTB simulation is a 2D grid world in which agents with varying skill levels collect resources that replenish stochastically, spend resources to build houses for coins (i.e., income), and trade coins for resources in an auction system, at the expense of labor costs for each action. It is not known to be a sequential social dilemma. GTB has a much larger state and action space than *Escape Room* and *Cleanup* due to the auction system, which provides 44 trading actions and supplements each agent’s spatial observation with the current and historical market information (e.g., counts of bids and asks for various price levels for each resource). At time t , the system productivity prod_t is defined as the sum of all agents’ coins, and equality eq_t is defined such that $\text{eq} = 1$ corresponds to uniform coin over agents and $\text{eq} = 0$ means only one agent has non-zero coin. The incentive designer is a central

tax planner who optimizes a trade-off between productivity and equality, defined as $\sum_{t=1}^H \text{prod}_t \cdot \text{eq}_t$ over horizon H , by imposing taxes according the following mechanism. Each episode lasts for $H = 100$ time steps and consists of 10 tax periods. At the start of each tax period, the ID sets a tax schedule $T: \mathbb{R} \rightarrow \mathbb{R}$ that determines the tax $T(z)$ applied to an agent’s income z earned within the period. T is a bracketed tax schedule based on the US federal taxation scheme: given a set of income thresholds $\{m_b\}_{b=0}^{B=7}$ with $m_0 = 0 < m_1 < \dots < m_B = \infty$, the ID defines T by setting a vector of *marginal tax rates* $[\tau_b]_{b=1}^B$, where $\tau_b \in [0, 1]$ applies to bracket (m_b, m_{b+1}) , such that the total tax on income z is given by

$$T(z) = \sum_{b=0}^{B-1} \tau_b ((m_{b+1} - m_b) \mathbf{1}_{z > m_{b+1}} + (z - m_b) \mathbf{1}_{m_b < z \leq m_{b+1}}) \quad (7)$$

where $\mathbf{1}_p = 1$ if p is true and 0 otherwise. At the end of each tax period, the total collected tax is evenly distributed back to all agents: if agent i gets total income z_i^p within period p , then the agent’s final adjusted income at the end of the tax period is given by:

$$\hat{z}_i^p = z_i^p - T(z_i^p) + \frac{1}{N} \sum_{j=1}^n T(z_j^p). \quad (8)$$

We used the 15×15 map called “env-pure_and_mixed-15x15” [61], which features similar spatial distribution of resource spawn points as the original 25×25 map. Each agent’s observation consists of an 11×11 egocentric spatial window, their resource inventories, collection and building skills, personal and other agents’ bids and asks, and quantities derived from the current period’s tax rate. The ID observes the complete spatial world state, agents’ inventories and incomes, cumulative bids and asks, and all derived tax quantities, but does not know agents’ private skill and utility functions. Agents have the same discrete action space consisting of movements, building, and trading actions. Appendix B.1.3 provides more information on observation/action spaces and agent utilities.

4.2 Implementation and Baselines

We describe the key implementation of all methods here and include all remaining details in Appendix B.2. We use $M = 1$ for MetaGrad across all experiments, such that an ID update occurs after each policy update by agents. We employ pipelining to improve the efficiency of MetaGrad: the validation trajectory $\hat{\tau}$ generated by agents’ updated policies (Algorithm 1, line 6), which is required for the ID’s update step, is used for the agents’ policy update in the next iteration (Algorithm 1, line 5). To differentiate through the agents’ learning step, MetaGrad has access to agents’ policy parameters and gradients. This assumption can be removed by using behavioral cloning to obtain surrogate models of agents, which has been demonstrated by existing methods that rely on knowledge of agent parameters [12, 60].

Our main baseline is termed **dual-RL**, in which the incentive designer itself is a standard RL agent who optimizes the system-level objective at the same time-scale as the original RL agents. Dual-RL is the centralized analogue of the decentralized agents with “give-reward” actions, introduced as a baseline in Yang et al. [60]. It is also formally equivalent to the method called “AI economist” in Zheng et al. [61]. In *Escape Room*, we compare with discrete-action and continuous-action variants of dual-RL, labeled “dual-RL (d)”

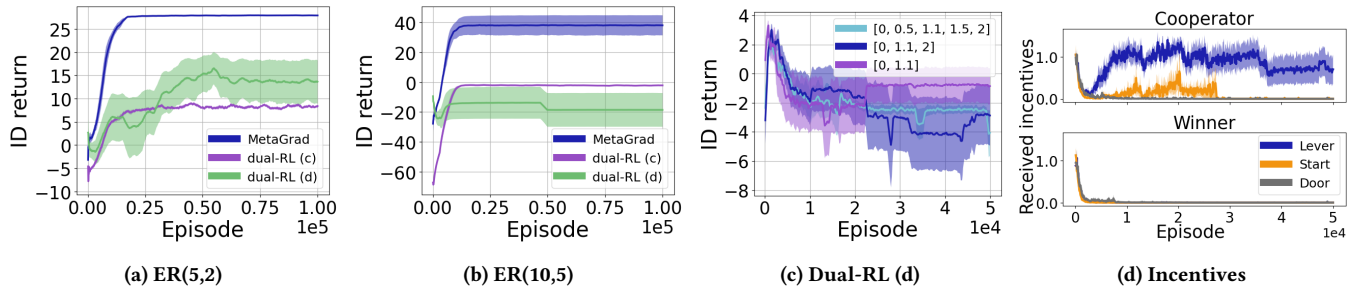


Figure 1: Escape Room. (a) In ER(5, 2), MetaGrad converges to the global optimum where two agents are incentivized by $1 + \epsilon$ (for some small $\epsilon > 0$) to pull the lever with -1 penalty, three other agents exit the door with +10 reward, which results in a total ID reward near 26. (b) MetaGrad also converges to the global optimum in ER(5, 2) with ID reward near 40. (c) Dual-RL (discrete) was unstable and had high variance across seeds, for various choices of discrete action spaces for the incentive values. (d) Incentives given by MetaGrad for each action by each agent in ER(2, 1).

and “dual-RL (c)”. In *Cleanup*, we compare with “dual-RL (c)”. In GTB, we implemented the core aspects of the “AI Economist” based on available information in Zheng et al. [61] (reabeled as “dual-RL” here), and also compare with the static US federal tax rates. We tuned hyperparameters for all methods equally using a successive elimination method, detailed in Appendix B.3.

Escape Room. We used policy gradient [51] without parameter sharing as the base agent implementation for all methods. The incentive function μ_η in MetaGrad is a neural network that maps the global state and agents’ joint action to a scaled sigmoid output layer of size $|\mathcal{A}| = 3$ (the number of possible agent actions), such that the value of each output node i lies in $(0, 2)$ and is interpreted as the incentive for action i taken by any agent. This parameterization enables MetaGrad to scale to larger number of agents, e.g. ER(10, 5). The cost for incentivization is the sum of all incentives given to agents, and is accounted by $\psi_t(\eta)$ in MetaGrad’s loss function.¹ For **dual-RL (d)**, we tried three different sets of discrete incentives $S_r = \{0, 1.1\}$, $S_r = \{0, 1.1, 2.0\}$, and $S_r = \{0, 0.5, 1.0, 1.5, 2.0\}$. The designer’s action space is $\text{Discrete}(|S_r|^{|\mathcal{A}|})$. Hence, the designer’s action is an assignment of a scalar incentive value to each possible agent action, and the policy output is a categorical distribution. In **dual-RL (c)**, the designer’s action space is $(0, 2)^n$, and its policy $\pi(a_{\text{ID}}|o)$ is defined by sampling $u \sim \mathcal{N}(\mu_\eta(s, a), 1)$ with neural network $\mu_\eta: \mathcal{S} \times \mathcal{A}^n \rightarrow \mathbb{R}^n$, then applying the same sigmoid output layer σ as MetaGrad to get $a_{\text{ID}} = \sigma(u)$. The total incentives given to agents are subtracted from R^{ID} . To compare with the method of Baumann et al. [4], we implemented separate experiments with actor-critic for agents’ policy updates.

Cleanup. We used actor-critic agents with TD(0) critic updates [50] and parameter-sharing as the base agent for all methods. MetaGrad and dual-RL (c) have the same architecture as for Escape Room, except that: 1) the observation input for agents and the ID is an RGB image; 2) the ID has a vector observation indicating whether or not each agent performed a cleaning action; 3) each output node of the incentive function is interpreted as the incentive for an action type in the set {fire cleaning beam, collect apples, else}.

¹At *train time*, one cannot for cost in R^{ID} of the current episode because MetaGrad only learns from R^{ID} in the next episode, not the episode where incentives are given. At test time, incentives are subtracted from R^{ID} so that comparison to baselines is fair.

Gather-Trade-Build. We used PPO agents [44] with parameter sharing for all methods. The incentive function in MetaGrad has $B = 7$ output nodes, where the value τ_b at each node b is capped by sigmoid activation to lie in $(0, 1)$ and is interpreted as the tax rate τ_b for bracket (m_b, m_{b+1}) . By (7), (8), and (25), each agent’s policy update is a differentiable function of the incentive function parameters η . The ID has seven action subspaces (one for each of the $B = 7$ tax brackets), each with 21 discrete actions that choose the marginal tax rate in $\{0, 0.05, \dots, 1.0\}$. Dual-RL applies standard RL to the ID, whose action space is a direct product of seven action subspaces (one for each of the $B = 7$ tax brackets), each with 21 discrete choices of the marginal tax rate in $\{0, 0.05, \dots, 1.0\}$. Zheng et al. [61] reported the need for a two-phase curriculum with tax annealing to stabilize training for Dual-RL, which may face difficulties with non-stationarity of the expanded multi-agent system. Hence, in the curriculum version of GTB experiments, we first train agents in a free-market (zero tax) scenario in Phase 1, while the RL-based tax policy is introduced in Phase 2 with a gradual annealing schedule on the maximum tax rate. We repeated experiments without curriculum to investigate the stability of MetaGrad. For measurements of economic activity and tax rates produced by trained models, we report the mean and standard error over the four trained models of the mean over 100 test episodes per model.

5 RESULTS

Escape Room. MetaGrad converged to the known global optimum value of approximately 26 in ER(5, 2) and 40 in ER(10, 5) (Figures 1a and 1b, respectively). Figure 1d shows the dynamics of incentivization during training in the case of ER(2, 1), where we labeled each agent *at the end of training* as either a “cooperator” or a “winner” based on whether the agent primarily pulls the lever or exits the door, respectively. We see that the cooperator consistently receives incentives of $1 + \epsilon$ during the majority of episodes, which explains the emergence of its cooperative behavior, whereas the winner receives zero incentives asymptotically, which shows the designer learned to avoid unnecessary costs. In contrast, both dual-RL (d) and dual-RL (c) did not solve ER, even including various choices of the discrete action space (Figure 1c). This is because a standard RL incentive designer optimizes the expected return of *one* episode,




Figure 2: GTB without curriculum: MetaGrad finds tax policies that induce higher social welfare than baselines, by promoting higher productivity at similar levels of equality.




Figure 4: GTB without curriculum: economic activity after 200k training episodes.

but the impact of its "give-reward" action only appears after agents have conducted learning updates over *many* episodes. In any given episode, the ID's reward contains no information about the impact of the actions it chose during that episode. The only conceivable way that dual-RL learns is by serendipity: the ID's action a in a *previous* episode i led to a change in agents' behavior that results in positive reward for the ID during the *current* episode j , and the ID happened to take a again in episode j , which results in correct credit assignment when learning from episode j . In fact, among eight independent runs for dual-RL, the only run that succeeded had the same random seed as in hyperparameter search. Appendix C, Figure 14, shows that MetaGrad outperforms the incentive design method of Baumann et al. [4], which faced numerical instabilities when extended from matrix games to Markov games.




Figure 6: 7x7 map

performance. Under dual-RL, agents did not receive appropriate incentives and hence behaved selfishly—occasionally using the cleaning beam but immediately competing for any apples that spawned—and therefore converged to low social welfare.

Cleanup. As shown in Figure 6, MetaGrad achieved a high level of social welfare, which is only possible because one agent received incentives to take cleaning actions while another agent collects apples. The method labeled "Cen" trains a single policy that acts for both agents; it serves as an empirical upper bound on




Figure 3: MetaGrad tax rates for each independent run.

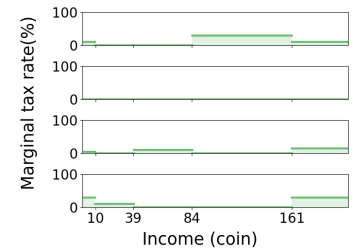


Figure 5: Dual-RL tax rates for each independent run.

5.1 Gather-Trade-Build

Social welfare. MetaGrad outperforms both dual-RL and US federal without requiring heuristics such as curriculum learning and tax annealing (Figure 2a). It discovers tax rates which differ from the static US federal tax rates in two notable aspects (compare Figure 3 and Figure 9). Firstly, MetaGrad imposes much higher taxes than US federal on the lowest income bracket (e.g., 0-10 coin), but chooses relatively lower taxes for the next income bracket (10-39). Hence, compared to US federal, there is less incentive for agents to fall in the lowest bracket, which may explain the higher income of agents under MetaGrad versus US federal (Figure 7). Secondly, the highest income bracket does not face significantly higher tax rates than other brackets, and even gets the lowest rate in one instance. While this results in lower equality than US federal (Figure 2b), the increase in economic activity—such as resource collection, building, and trading (Figure 4)—improves system productivity (Figure 2c) and ultimately produces significantly higher social welfare.




Figure 9: US federal

When curriculum learning is applied to all methods—i.e., for all methods, initializing agents with the same policy that was pre-trained in a free market context—MetaGrad also exceeds the performance of dual-RL. Except for one out of four runs where social welfare abruptly dropped around 150k episodes, MetaGrad also outperforms US federal. Similar to the training dynamics observed in Zheng et al. [61], we also observe a transient period where dual-RL passes through unstable local optima (Figure 10a,

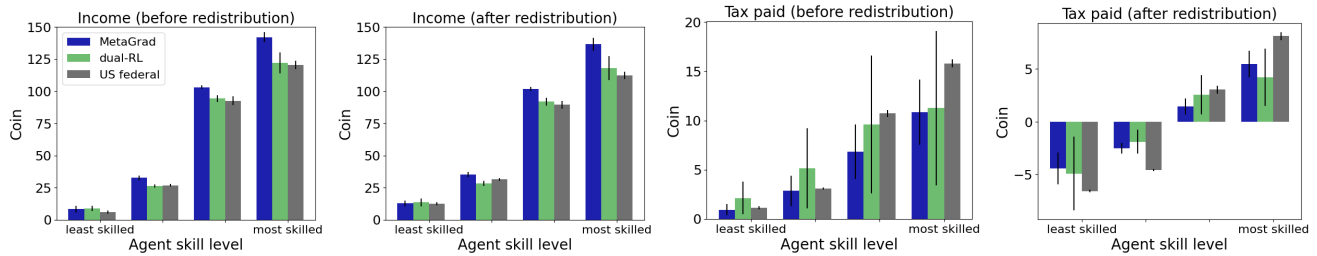


Figure 7: GTB without curriculum: income and tax before and after redistribution, after 200k training episodes.

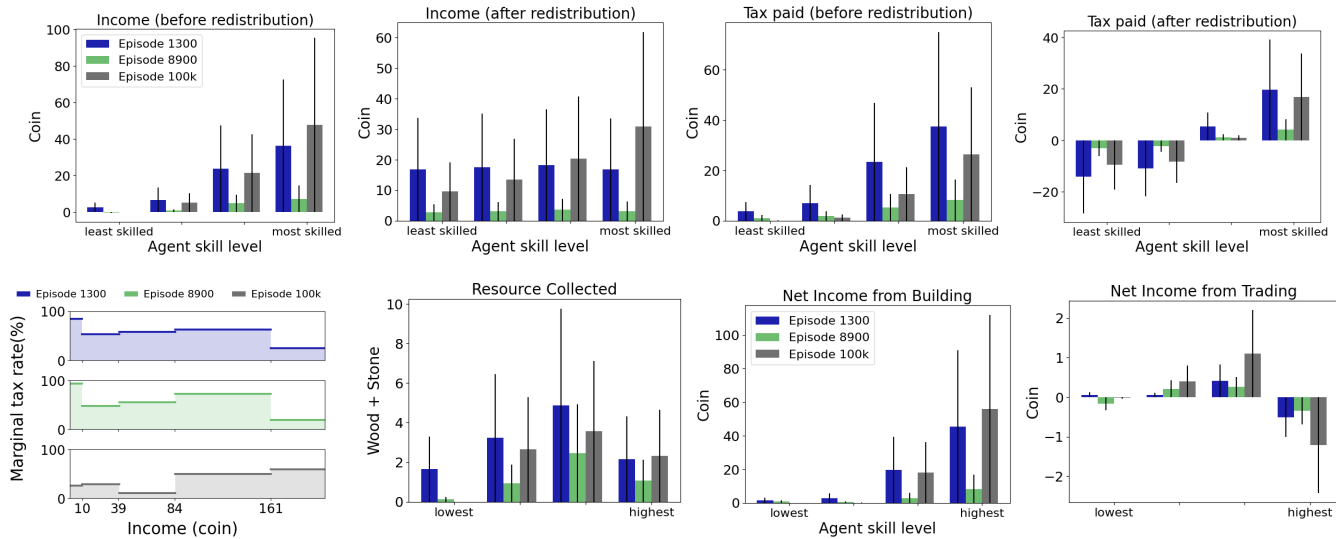


Figure 8: Training dynamics. Status of agent and designer behavior at 1300, 8900, and 100k episodes. Top row: income and tax before and after redistribution. Bottom row: tax rates and economic activity.

near 25k episodes); however, in contrast to their results, dual-RL did not manage to surpass US federal in asymptotic performance. This is likely because the sudden introduction of a tax planner in Phase 2 may do more harm than good for stability, especially when the extra hyperparameters introduced by curriculum learning and tax annealing are hard to tune. Both MetaGrad and dual-RL enact higher taxes at lower income brackets than US federal.

Economic activity. Because skill levels determine the amount of coins per house built, differences in agents’ skill levels generate income inequality and behavioral specialization. For example, agents with second-highest skill tend to collect the most resources and sell them for income, whereas agents with the highest skill spend less effort on resource collection but generate income by building houses from purchased resources (Figures 4a and 4c). Notably, under all tax policies, all except the highest skill agents receive net positive income from trading (Figure 4c). Even though resource collection is comparable across methods, tax policies found by MetaGrad encouraged highest trading activity and hence highest overall income from building, compared to US federal and dual-RL. In the curriculum case, dual-RL tax policies impose high taxation (above 50%) on the three lowest income brackets Figure 13. This may explain the fact that the lowest-skilled agent collects zero resources (Figure 12a), which lowers overall system productivity.

Taxation and income. Agents with the two highest skill levels pay significantly less tax for tax policies found by MetaGrad than they do for the US federal tax rates, whereas agents with the two lowest skill pay comparably equal tax (Figure 7). This means that MetaGrad does better than US federal at encouraging higher skilled agents to increase economic activity such as building and trading, without affecting resource collection by lower skill agents (as can be seen in Figure 4a). While this comes at the expense of lower equality (Figure 2b), the incomes of lower-skilled agents both before and after redistribution are actually higher for MetaGrad than US federal, because lower-skilled agents benefit from increased trading activity (Figure 4c). Dual-RL tax policies caused agents of all skill levels to pay more taxes than MetaGrad, which is correlated in overall lower building and trading activity.

Training dynamics. In GTB without curriculum (Figure 2a), the fact that social welfare rises while equality drops in the early 50k episodes is due to *tabula rasa* learning with heterogeneous skills. In the curriculum setting, the changing tax rates under MetaGrad produced a more dynamical social welfare curve than the fixed US federal rates during training (Figure 10a). This can be useful for extracting potential causal relations between taxation and agents’ economic behavior. For one run of MetaGrad, we measured taxation, income, and economic activity over 100 test episodes at 1300,

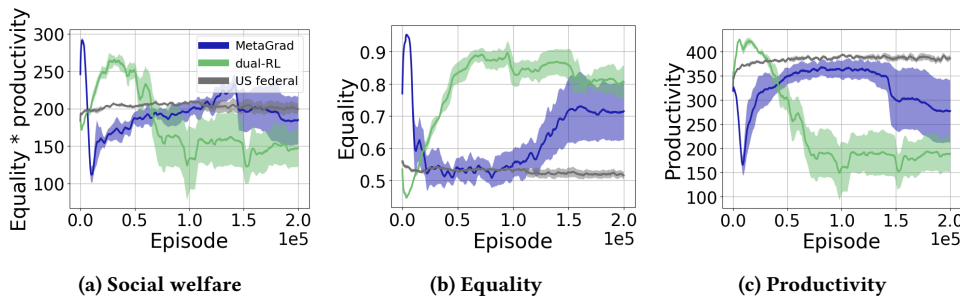


Figure 10: GTB with curriculum. Both MetaGrad and dual-RL find transient states of high social welfare but are less stable than US federal.




Figure 11: MetaGrad tax rates with curriculum.

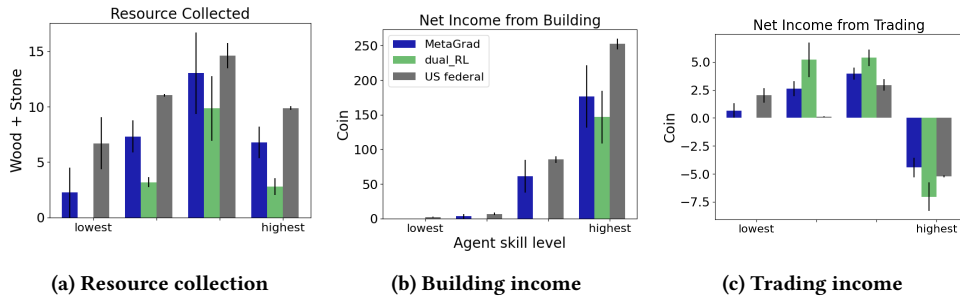


Figure 12: GTB with curriculum: economic activity after 200k training episodes in Phase 2.




Figure 13: Dual-RL tax with curriculum.

8900, and 100k episodes during training, corresponding to the early peak, valley, and steady rising region in the social welfare curve in Figure 10a. These measurements are shown in Figure 8, and we make the following observations. Social welfare is highest at episode 1300, but agents actually have lower income (before redistribution) at episode 1300 than at episode 100k. Because taxes do not have instantaneous effect on productivity, MetaGrad quickly learned that social welfare can be artificially inflated by finding a tax scheme to increase the equality index—hence the sharp early peak in Figure 10b. As shown in Figure 8, MetaGrad’s tax rates at episode 1300 produced nearly uniform income (after redistribution) over all skill levels, by enacting high taxes on the higher skilled agents. This tax policy disincentivizes agents from earning high income, since episode 1300 is followed by a precipitous drop in productivity (Figure 10c) that reaches a global minimum near episode 8900, where agent activity levels are lowest (Figure 8). Nonetheless, MetaGrad adapted its tax policy to produce a recovery of productivity at relatively constant equality, from episodes 25k to 100k (Figures 10b and 10c). The tax policy at episode 100k resembles the progressive schedule of the US federal policy, albeit with significantly lower rates for the 39-84 income bracket that applies to the high-skilled agents except for the very top earners.

Difficulty of GTB versus *Escape Room* and *Cleanup*. One may wonder why the performance gap between MetaGrad and dual-RL is much higher in *Escape Room* and *Cleanup* than in GTB. This is because the incentivization problem in GTB is conceptually easier for an incentive designer who uses conventional RL. In GTB, agents can learn from positive rewards regardless of the tax rate, which implies that the ID always receives a learning signal that tracks the agents’ behaviors. For dual-RL, the ID receives such

feedback in the *next* episode, but at least it is non-zero. However, in the other two problems, if incentives do not pass a threshold ($1 + \epsilon$ to overcome the environment penalty in *Escape Room*; a more complex opportunity cost for sacrificing self-interest in *Cleanup*), then the agents do not receive any predefined environment rewards, which means an RL-based designer does not receive any feedback at all. In this case, MetaGrad’s knowledge of the way that agents’ policy parameters change in response to the incentive function, and the use of online cross-validation to learn from the ID’s returns in subsequent episodes, is crucial.

6 CONCLUSION

We proposed the use of complex simulations involving reinforcement learning agents as an *in silico* experimental approach to problems of incentive design. To tackle the issue of delayed impact of incentives, which poses difficulties for directly applying standard RL to the incentive designer, we proposed a meta-gradient approach for the incentive designer to account exactly for the agents’ learning response to incentives. The new method significantly outperforms baselines on benchmark problems and also improves the trade-off between productivity and equality in a complex simulated economy.

Beyond incentive design, one may consider the extension of ideas in this work to the context of *mechanism* design, interpreted in the general sense of modifying the underlying dynamics of the environment [21] to shape agents’ behavior and optimize a system-level objective. More generally, we hope this work shows the feasibility of a path toward a data and simulation-driven approach for improving complex systems in society.

REFERENCES

- [1] Beena Ammanath, David Jarvis, and Susanne Hupfer. 2020. *Thriving in the era of pervasive AI: Deloitte's State of AI in the Enterprise, 3rd Edition*. Technical Report. Deloitte.
- [2] W Brian Arthur. 2021. Foundations of complexity economics. *Nature Reviews Physics* (2021), 1–10.
- [3] Wolfram Barfuss, Jonathan F Donges, Vitor V Vasconcelos, Jürgen Kurths, and Simon A Levin. 2020. Caring for the future can turn tragedy into comedy for long-term collective action under risk of collapse. *Proceedings of the National Academy of Sciences* 117, 23 (2020), 12915–12922.
- [4] Tobias Baumann, Thore Graepel, and John Shawe-Taylor. 2020. Adaptive Mechanism Design: Learning to Promote Cooperation. In *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–7.
- [5] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. 2019. Dota 2 with Large Scale Deep Reinforcement Learning. *arXiv preprint arXiv:1912.06680* (2019).
- [6] David Bissell, Thomas Birtchnell, Anthony Elliott, and Eric L Hsu. 2020. Autonomous automobilities: The social impacts of driverless vehicles. *Current Sociology* 68, 1 (2020), 116–134.
- [7] Jean-François Bonnefon, Azim Shariff, and Iyad Rahwan. 2016. The social dilemma of autonomous vehicles. *Science* 352, 6293 (2016), 1573–1576.
- [8] Gianluca Brero, Alon Eden, Matthias Gerstgrasser, David C Parkes, and Duncan Rheingans-Yoo. 2020. Reinforcement Learning of Simple Indirect Mechanisms. *arXiv preprint arXiv:2010.01180* (2020).
- [9] Allan Dafoe, Edward Hughes, Yoram Bachrach, Tantum Collins, Kevin R McKee, Joel Z Leibo, Kate Larson, and Thore Graepel. 2020. Open Problems in Cooperative AI. *arXiv preprint arXiv:2012.08630* (2020).
- [10] Panayiotis Danassis, Aris Filos-Ratsikas, and Boi Faltings. 2021. Achieving Diverse Objectives with AI-driven Prices in Deep Reinforcement Learning Multi-agent Markets. *arXiv preprint arXiv:2106.06060* (2021).
- [11] André de Palma and Robin Lindsey. 2011. Traffic congestion pricing methodologies and technologies. *Transportation Research Part C: Emerging Technologies* 19, 6 (2011), 1377–1399.
- [12] Jakob Foerster, Richard Y Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. 2018. Learning with Opponent-Learning Awareness. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. 122–130.
- [13] Ragavendran Gopalakrishnan, Jason R Marden, and Adam Wierman. 2014. Potential games are necessary to ensure pure nash equilibria in cost sharing games. *Mathematics of Operations Research* 39, 4 (2014), 1252–1296.
- [14] Noortje Groot, Bart De Schutter, and Hans Hellendoorn. 2012. Reverse stackelberg games, part i: Basic framework. In *2012 IEEE International Conference on Control Applications*. IEEE, 421–426.
- [15] Yu-Chi Ho, P Luh, and Ramal Muralidharan. 1981. Information structure, Stackelberg games, and incentive controllability. *IEEE Trans. Automat. Control* 26, 2 (1981), 454–460.
- [16] Philip N Howard, Samuel Woolley, and Ryan Calo. 2018. Algorithms, bots, and political communication in the US 2016 election: The challenge of automated political communication for election law and administration. *Journal of information technology & politics* 15, 2 (2018), 81–93.
- [17] Ming-Wei Hsu, Stefan Lessmann, Ming-Chien Sung, Tiejun Ma, and Johnnie EV Johnson. 2016. Bridging the divide in financial market forecasting: machine learners vs. financial economists. *Expert Systems with Applications* 61 (2016), 215–234.
- [18] Edward Hughes, Joel Z Leibo, Matthew Phillips, Karl Tuyls, Edgar Dueñez-Guzman, Antonio García Castañeda, Iain Dunning, Tina Zhu, Kevin McKee, Raphael Koster, et al. 2018. Inequity aversion improves cooperation in intertemporal social dilemmas. In *Advances in neural information processing systems*. 3326–3336.
- [19] Eugene Ie, Chih-wei Hsu, Martin Mladenov, Vihan Jain, Sanmit Narvekar, Jing Wang, Rui Wu, and Craig Boutilier. 2019. Recsim: A configurable simulation platform for recommender systems. *arXiv preprint arXiv:1909.04847* (2019).
- [20] Louis Kirsch, Sjoerd van Steenkiste, and Juergen Schmidhuber. 2019. Improving Generalization in Meta Reinforcement Learning using Learned Objectives. In *International Conference on Learning Representations*.
- [21] János Kramár, Neil Rabinowitz, Tom Eccles, and Andrea Tacchetti. 2020. Should I tear down this wall? Optimizing social metrics by evaluating novel actions. *arXiv preprint arXiv:2004.07625* (2020).
- [22] Nevena Lazić, Craig Boutilier, Tyler Lu, Ehern Wong, Binz Roy, MK Ryu, and Greg Imwalle. 2018. Data center cooling using model-predictive control. In *NeurIPS*.
- [23] Changxi Li, Fenghua He, Hongsheng Qi, Daizhan Cheng, Longbiao Ma, Yonghong Wu, and Shuo Chen. 2018. Potential games design using local information. In *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 1911–1916.
- [24] Jiayang Li, Jing Yu, Yu Nie, and Zhaoran Wang. 2020. End-to-end learning and intervention in games. *Advances in Neural Information Processing Systems* 33 (2020).
- [25] Na Li and Jason R Marden. 2013. Designing games for distributed optimization. *IEEE Journal of Selected Topics in Signal Processing* 7, 2 (2013), 230–242.
- [26] Konstantinos G Liakos, Patrizia Busato, Dimitrios Moshou, Simon Pearson, and Dionysis Bochtis. 2018. Machine learning in agriculture: A review. *Sensors* 18, 8 (2018), 2674.
- [27] Michael L Littman. 1994. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*. Elsevier, 157–163.
- [28] Jonathan Lorraine, Paul Vicol, and David Duvenaud. 2020. Optimizing millions of hyperparameters by implicit differentiation. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 1540–1552.
- [29] Jie Lu, Dianshuang Wu, Mingsong Mao, Wei Wang, and Guangquan Zhang. 2015. Recommender system application developments: a survey. *Decision Support Systems* 74 (2015), 12–32.
- [30] David Mguni, Joel Jennings, Emilio Sison, Sergio Valcarcel Macua, Sofia Ceppi, and Enrique Munoz de Cote. 2019. Coordinating the Crowd: Inducing Desirable Equilibria in Non-Cooperative Systems. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. 386–394.
- [31] Fei Miao, Shuo Han, Shan Lin, John A Stankovic, Desheng Zhang, Sirajum Munir, Hua Huang, Tian He, and George J Pappas. 2016. Taxi dispatch with real-time sensing data in metropolitan areas: A receding horizon control approach. *IEEE Transactions on Automation Science and Engineering* 13, 2 (2016), 463–478.
- [32] Friedrich Wilhelm Nietzsche. 1866. *Beyond good and evil: Prelude to a philosophy of the future*. Vintage.
- [33] Yael Niv. 2009. Reinforcement learning in the brain. *Journal of Mathematical Psychology* 53, 3 (2009), 139–154.
- [34] Junhyuk Oh, Matteo Hessel, Wojciech M Czarnecki, Zhongwen Xu, Hado P van Hasselt, Satinder Singh, and David Silver. 2020. Discovering Reinforcement Learning Algorithms. *Advances in Neural Information Processing Systems* 33 (2020).
- [35] Frans A Oliehoek and Christopher Amato. 2016. *A concise introduction to decentralized POMDPs*. Springer.
- [36] Afshin Oroojlooyjadid and Davood Hajinezhad. 2019. A review of cooperative multi-agent deep reinforcement learning. *arXiv preprint arXiv:1908.03963* (2019).
- [37] Dario Paccagnan, Rahul Chandan, and Jason R Marden. 2019. Utility Design for Distributed Resource Allocation-Part I: Characterizing and Optimizing the Exact Price of Anarchy. *IEEE Trans. Automat. Control* (2019).
- [38] Liviu Panait and Sean Luke. 2005. Cooperative multi-agent learning: The state of the art. *Autonomous agents and multi-agent systems* 11, 3 (2005), 387–434.
- [39] David Pardoe, Peter Stone, Maytal Saar-Tscheskany, and Kerem Tomak. 2006. Adaptive mechanism design: a metalearning approach. In *Proceedings of the 8th international conference on Electronic commerce*. 92–102.
- [40] Lillian J Ratliff, Roy Dong, Shreyas Sekar, and Tanner Fiez. 2019. A Perspective on Incentive Design: Challenges and Opportunities. *Annual Review of Control, Robotics, and Autonomous Systems* 2 (2019), 305–338.
- [41] Lillian J Ratliff and Tanner Fiez. 2020. Adaptive incentive design. *IEEE Trans. Automat. Control* (2020).
- [42] David Robinson and David Goforth. 2005. *The topology of the 2x2 games: a new periodic table*. Vol. 3. Psychology Press.
- [43] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015. Trust region policy optimization. In *International conference on machine learning*. PMLR, 1889–1897.
- [44] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [45] Weiran Shen, Pingzhong Tang, and Song Zuo. 2019. Automated Mechanism Design via Neural Networks. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. 215–223.
- [46] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. 2017. Mastering the game of go without human knowledge. *Nature* 550, 7676 (2017), 354–359.
- [47] H. von Stackelberg. 1934. Marktform und Gleichgewicht. *Julius Springer* (1934).
- [48] Bradley Stadie, Lunjun Zhang, and Jimmy Ba. 2020. Learning Intrinsic Rewards as a Bi-Level Optimization Problem. In *Conference on Uncertainty in Artificial Intelligence*. PMLR, 111–120.
- [49] Richard S Sutton. 1992. Adapting bias by gradient descent: An incremental version of delta-bar-delta. In *AAAI*. 171–176.
- [50] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [51] Richard S Sutton, David A McAllester, Satinder P Singh, Yishay Mansour, et al. 1999. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, Vol. 99. Citeseer, 1057–1063.
- [52] Pingzhong Tang. 2017. Reinforcement mechanism design. In *IJCAI*. 5146–5150.
- [53] Leigh Tesfatsion and Kenneth L Judd. 2006. *Handbook of computational economics: agent-based computational economics*. Elsevier.
- [54] John S Vardakas, Nizar Zorba, and Christos V Verikoukis. 2014. A survey on demand response programs in smart grids: Pricing methods and optimization algorithms. *IEEE Communications Surveys & Tutorials* 17, 1 (2014), 152–178.

- [55] Eugene Vinyals. 2020. *Sequential Social Dilemma Games*. https://github.com/eugenevinyals/sequential_social_dilemma_games.
- [56] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 575, 7782 (2019), 350–354.
- [57] Cathy Wu, Alexandre M Bayen, and Ankur Mehta. 2018. Stabilizing traffic with autonomous vehicles. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 6012–6018.
- [58] Zhongwen Xu, Hado van Hasselt, Matteo Hessel, Junhyuk Oh, Satinder Singh, and David Silver. 2020. Meta-Gradient Reinforcement Learning with an Objective Discovered Online. *arXiv preprint arXiv:2007.08433* (2020).
- [59] Zhongwen Xu, Hado P van Hasselt, and David Silver. 2018. Meta-gradient reinforcement learning. In *Advances in neural information processing systems*. 2396–2407.
- [60] Jiachen Yang, Ang Li, Mehrdad Farajtabar, Peter Sunehag, Edward Hughes, and Hongyuan Zha. 2020. Learning to Incentivize Other Learning Agents. *Advances in Neural Information Processing Systems* 33 (2020).
- [61] Stephan Zheng, Alexander Trott, Sunil Srinivasa, Nikhil Naik, Melvin Gruesbeck, David C Parkes, and Richard Socher. 2020. The AI Economist: Improving Equality and Productivity with AI-Driven Tax Policies. *arXiv preprint arXiv:2004.13332* (2020).
- [62] Zeyu Zheng, Junhyuk Oh, and Satinder Singh. 2018. On learning intrinsic rewards for policy gradient methods. In *Advances in Neural Information Processing Systems*. 4644–4654.