



CIRANO
Knowledge into action

ALGORITHMIC COLLUSION AND THE MINIMUM PRICE MARKOV GAME

IGOR SADOUNE
MARCELIN JOANIS
ANDREA LODI

CS

2025s-07
WORKING PAPER

The purpose of the **Working Papers** is to disseminate the results of research conducted by CIRANO research members in order to solicit exchanges and comments. These reports are written in the style of scientific publications. The ideas and opinions expressed in these documents are solely those of the authors.

Les cahiers de la série scientifique visent à rendre accessibles les résultats des recherches effectuées par des chercheurs membres du CIRANO afin de susciter échanges et commentaires. Ces cahiers sont rédigés dans le style des publications scientifiques et n'engagent que leurs auteurs.

CIRANO is a private non-profit organization incorporated under the Quebec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the government of Quebec, and grants and research mandates obtained by its research teams.

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du gouvernement du Québec, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO Partners – Les partenaires du CIRANO

Corporate Partners – Partenaires Corporatifs

Autorité des marchés financiers
Banque de développement du Canada
Banque du Canada
Banque Nationale du Canada
Bell Canada
BMO Groupe financier
Caisse de dépôt et placement du Québec
Énergir
Hydro-Québec
Intact Corporation Financière
Investissements PSP
Manuvie
Mouvement Desjardins
Power Corporation du Canada
VIA Rail Canada

Governmental partners - Partenaires gouvernementaux

Ministère des Finances du Québec
Ministère de l'Économie, de
l'Innovation et de l'Énergie
Innovation, Sciences et Développement
Économique Canada
Ville de Montréal

University Partners – Partenaires universitaires

École de technologie supérieure
École nationale d'administration
publique
de Montréal
HEC Montreal
Institut national de la recherche
scientifique
Polytechnique Montréal
Université Concordia
Université de Montréal
Université de Sherbrooke
Université du Québec
Université du Québec à Montréal
Université Laval
Université McGill

CIRANO collaborates with many centers and university research chairs; list available on its website. *Le CIRANO collabore avec de nombreux centres et chaires de recherche universitaires dont on peut consulter la liste sur son site web.*

© April 2025. Igor Sadoune, Marcelin Joanis, Andrea Lodi. All rights reserved. *Tous droits réservés.* Short sections may be quoted without explicit permission, if full credit, including © notice, is given to the source. *Reproduction partielle permise avec citation du document source, incluant la notice ©.*

The observations and viewpoints expressed in this publication are the sole responsibility of the authors; they do not represent the positions of CIRANO or its partners. *Les idées et les opinions émises dans cette publication sont sous l'unique responsabilité des auteurs et ne représentent pas les positions du CIRANO ou de ses partenaires.*

ISSN 2292-0838 (online version)

Algorithmic collusion and the minimum price Markov game

Igor Sadoune^{*}, Marcelin Joanis[†], Andrea Lodi[‡]

Abstract/Résumé

This paper introduces the Minimum Price Markov Game (MPMG), a theoretical model that reasonably approximates real-world first-price markets following the minimum price rule, such as public auctions. The goal is to provide researchers and practitioners with a framework to study market fairness and regulation in both digitized and non-digitized public procurement processes, amid growing concerns about algorithmic collusion in online markets. Using multi-agent reinforcement learning-driven artificial agents, we demonstrate that (i) the MPMG is a reliable model for first-price market dynamics, (ii) the minimum price rule is generally resilient to non-engineered tacit coordination among rational actors, and (iii) when tacit coordination occurs, it relies heavily on self-reinforcing trends. These findings contribute to the ongoing debate about algorithmic pricing and its implications.

Cet article présente le jeu du prix minimum de Markov (MPMG), un modèle théorique qui se rapproche raisonnablement des marchés réels qui suivent la règle du prix minimum, tels que les enchères publiques. L'objectif est de fournir aux chercheurs et aux praticiens un cadre pour étudier l'équité du marché et la réglementation dans les processus de marchés publics numériques et non numériques, dans un contexte de préoccupations croissantes concernant la collusion algorithmique sur les marchés en ligne. En utilisant des agents artificiels basés sur l'apprentissage par renforcement multi-agents, nous démontrons que (i) le MPMG est un modèle fiable pour la dynamique du marché au premier prix, (ii) la règle du prix minimum est généralement résistante à la coordination tacite non technique entre les acteurs rationnels, et (iii) lorsque la coordination tacite se produit, elle s'appuie fortement sur des tendances qui se renforcent d'elles-mêmes. Ces résultats contribuent au débat en cours sur la tarification algorithmique et ses implications.

Keywords/Mots-clés: Algorithmic Game Theory, Multiagent Reinforcement Learning, Algorithmic Coordination, Algorithmic Pricing / Théorie algorithmique des jeux, apprentissage par renforcement multi-agents, coordination algorithmique, tarification algorithmique

Pour citer ce document / To quote this document

Sadoune, I., Joanis, M., & Lodi, A. (2025). Algorithmic collusion and the minimum price Markov game (2025s-07, Cahiers scientifiques, CIRANO.) <https://doi.org/10.54932/MMFD7174>

To replicate the study, follow [replication repo \(GitHub\)](#).

MPMG is implemented in a Python package; see [PyPI](#) or [GitHub](#).

^{*} Department of Mathematics and Industrial Engineering, Polytechnique Montreal. CIRANO. igor.sadoune@polymtl.ca

[†] Department of Mathematics and Industrial Engineering, Polytechnique Montreal. CIRANO. marcelin.joanis@polymtl.ca

[‡] Jacobs Technion-Cornell Institute, Cornell Tech and Technion – IIT. andrea.lodi@cornell.edu

1 Introduction

Concerns about algorithmic pricing have grown as technological advancements in machine learning, widespread data availability, and the digitization of the economy have created fertile grounds for such practices [1, 2, 3, 4]. Algorithmic pricing can potentially lead to collusion, as observed in digital markets such as online retail [5, 6] and electricity supply markets [7, 8, 9, 10, 11]. These sectors demonstrate the potential for artificial decision-makers to learn collusive behaviors through dynamic pricing algorithms and optimal bidding strategies. This echoes the concerns about algorithmic collusion in public procurement, which are increasingly relevant, especially as digital platforms and automated decision-making tools become more prevalent in these settings. Since public procurement is a crucial mechanism for collective welfare, the vulnerability of this sector to collusive outcomes is particularly alarming. For example, it has been shown that bidding patterns were incompatible with a competitive equilibrium in Ukrainian E-procurement data [12]. Given these developments, questions arise. Is public procurement directly threatened by algorithmic pricing? And if so, under which conditions?

In contrast to digital markets, public procurement in most Western economies typically adheres to first-price sealed-bid auction rules, and more specifically, to the minimum price rule that determines the winner by selecting the lowest bid. This rule is supported by a longstanding body of literature that demonstrates the theoretical efficiency of first-price auctions and, by extension, the efficacy of the minimum price rule in maximizing collective welfare [13, 14, 15]. Public procurement processes, generally not digitized, have made more data available to the public due to recent transparency laws, a response to a long history of traditional cartel collusion [16, 17, 18]. This initiative has been effective as it has helped regulators and researchers detect collusion in public markets [19, 20]. However, the data can also be used to train artificial decision-makers to provide optimal bidding strategies. The potential for these strategies to lead to algorithmic collusion, whether accidental or deliberately learned, raises concerns for legal authorities because such collusion would not trigger antitrust laws [21]. Indeed, accidental algorithmic collusion stems from profit-maximization motives, and proving the deliberate formation of algorithmic coordination remains challenging.

In this paper, we introduce the Minimum Price Markov Game (MPMG), a model designed to provide a framework for studying emergent behavior in the context of minimum price auctions, commonly used in real-world public procurement. We draw on Robert Axelrod’s concept of “the emergence of cooperation among egoist” [22], exploring tacit coordination among isolated agents in the MPMG. Departing from the traditional game-theoretic Bayesian framework of first-price auctions, we consider a game of complete information based on the assumption of data availability. Indeed, public auction data typically includes the features of auctions and agents from past to ongoing contracts, as well as the associated bids [23]. This allows us to base the MPMG on the Prisoner’s Dilemma (PD), leveraging a structure that captures the essence of public procurement as a social dilemma—competitive by law but cooperative by nature. Indeed, the PD is well known for encapsulating the paradox of rationality central to our discussion on the study of emergent coordination among isolated agents in public markets. Furthermore, we test the MPMG using various state-of-the-art Multi-Agent Reinforcement Learning (MARL) methods and thereby show that the minimum price rule, in our setting, is generally robust but not impervious to tacit coordination among isolated decision-makers.

Our methodological choice is motivated by extensive literature that shows the effectiveness of combining game theory and reinforcement learning for modeling strategic interactions and optimizing behavior through dynamic adaptation among agents. We explore such literature in Section 2 and show that our methodological framework is ideal for exploring emergent behaviors and potential tacit collusion in public procurement markets. In Section 3, we provide the mathematical model for the MPMG and discuss the benefits of the Markov property in our setting. Section 4 presents the results of our computer experiments and concludes on their implications for tacit coordination in minimum price-driven auctions. Finally, Section 5 provides interpretation, practical considerations and perspectives.

2 Related Work

The interplay among Markov Decision Process (MDP), auction design, and multi-agent learning has gained significant attention in recent research. Although the literature addressing algorithmic collusion within this framework is relatively sparse, our study draws inspiration from the growing body of work on integrating Stackelberg game models within multi-agent learning frameworks, MDP for auction designs, and tacit coordination in social dilemmas. By reviewing these interconnected areas, we aim to highlight the advancements in understanding strategic interactions, optimal bidding strategies, and tacit coordination among agents. This overview highlights the pivotal role of MDP and related methodologies in shaping contemporary economic and strategic decision-making processes.

MDP and auction design. The MDP framework facilitates the exploration, understanding, and prediction of strategic interactions and behaviors within market designs. Examples include computing Nash equilibria in auction games [24, 25] and optimizing bidding strategies in combinatorial auctions [26, 27]. Although the Markov game framework

is influential, bandit algorithms have also proven to be exceptionally suitable for modeling optimal bidding strategies in online auction designs [28, 29, 30, 31]. This is due to the simplicity and effectiveness of bandit algorithms, which operate in a stateless or single-state environment and make decisions independently of past actions without the need for an underlying state transition model.

MDP and Stackelberg game. Stackelberg models, in which players move sequentially rather than simultaneously and therefore were historically used to study hierarchical or sequential leadership, are being increasingly used in various applications within the Markov game and multiagent learning framework [32]. The Stackelberg game is a staple in the analysis of various economic and strategic situations where the ability to move first or commit to a decision before one’s competitors can confer a strategic advantage. This includes industries with clear market leaders who can influence the market environment [33]. For example, Reinforcement Learning (RL) agents have been trained to learn the Stackelberg equilibrium in a Markov decision process representing economic mechanisms for assigning items to individualistic agents in sequential-pricing stages [34, 35, 36].

Multiagent learning, social dilemmas, and tacit coordination. In the exploration of tacit collusion and cooperative strategies within repeated game frameworks, a multitude of studies have contributed to a nuanced understanding of how agents can implicitly coordinate their actions under restricted communication conditions to optimize collective outcomes. The effectiveness of collusion in auction settings without explicit communication is a recurring theme [37], with similar phenomena observed in oligopoly models through Q-learning mechanisms where collusion emerges without direct interaction among agents [38, 39]. These findings are complemented by insights into multi-agent control using deep reinforcement learning, which further illustrates how complex cooperative behaviors can be engineered even in large-scale agent systems [40].

In scenarios where direct communication is either limited or non-existent, learning algorithms play a pivotal role in facilitating implicit cooperation and understanding competitive behaviors. For instance, learning methods that incorporate the anticipation of other agents’ learning processes have shown promise in fostering cooperative strategies that are robust against exploitation [41]. This is particularly evident in the Iterated Prisoner’s Dilemma, where strategies developed through evolutionary algorithms have shown to outperform traditional approaches [42]. Indeed, the contradictory nature of these games makes their study with RL agents quite challenging, as demonstrated in [43]. Moreover, the development of frameworks that enforce cooperation and resist collusion highlights the potential for creating stable cooperative environments in public goods games and other multi-agent settings [44, 45].

In this paper, we build on this literature and explore the potential for tacit coordination among isolated artificial agents in our social dilemma-based MPMG using multiagent learning.

3 The Minimum Price Markov Game

In real-world multiagent systems, a common observation is that heterogeneous agents coexist in non-cooperative environments. For example, in heterogeneous markets, stronger agents typically enjoy a competitive advantage, making growth, or even survival, more difficult for weaker ones. This disparity in market power—whether due to factors like market share, technology, or cost structures—creates a stratification of prices, influencing price equilibrium. In the context of first-price public auctions, stronger agents tend to submit the lowest bids, leading to the expectation that they will win more frequently. This assumption is reasonable, as their ability to cut costs is positively correlated with their operational capacity.

In a game-theoretic setting, it is necessary to define a payoff function that applies to both heterogeneous and homogeneous markets without losing generality. The payoff structure should smoothly transition from homogeneous to heterogeneous markets, allowing for a mathematically consistent analysis of strategic behaviors across different numbers of players. Market heterogeneity should be reflected in differences in payoffs, win frequencies, and incentives to avoid price collusion, with these variations proportionate to the gaps in market strength.

Furthermore, in public procurement, there is a tension between the pursuit of personal gain and collective welfare. Firms strive to maximize their individual profits, while the state aims to allocate public funds efficiently and maintain market competition and fairness. From the firms’ perspective, individually rational choices made without cooperation can lead to suboptimal outcomes for all participants, as maximal profits are only achievable through cooperative behavior. This conflict between individual and collective rationality is the core characteristic of social dilemmas. The minimum price rule, with its mathematical elegance, does not diminish this inherent conflict; rather, it preserves this tension fully, maintaining the balance between self-interest and group benefit.

It follows that a minimum price-ruled auction game formulation will fall in accordance with the principles of social dilemma games. Namely, it must display non-zero-sum interactions for which players can benefit from cooperating or

suffer from a single defection. The dominant strategy (Nash Equilibrium), being collectively irrational, is thus Pareto inefficient, meaning that an efficient coordination scheme leads to the strategy profile associated with the non-Nash Pareto Optimal outcome. In other words, the role of coordination is to transition from the Nash Equilibrium to the profit-maximizing strategy profile.

In addition, the single-stage formulation must capture the core elements of market dynamics, which inherently unfold through repeated interactions. These characteristics must be maintained in its dynamic extension, such as in the iterated Minimum Price Game (MPG) and the Minimum Price Markov Game (MPMG).

In this section, we provide a formulation for the MPMG, that we believe is a reasonable approximation of the typical real-world minimum price-ruled public procurement process. Our formulation is rooted in assumptions that capture the core of this process, creating a conducive environment to observe how artificial learners dynamically shape their strategies in such dilemma. The MPMG is founded on a static normal-form game, the MPG, which we will define first.

3.1 The Minimum Price Game: A Single-Stage Formulation

The MPG involves n agents, represented by a set $\mathcal{N} = \{1, \dots, n\}$, competing to win a contract valued at v in a first-price procurement auction.

Assumption 1 (Single Stage Auction). *The bidding process unfolds in a single round.*

Assumption 2 (Simultaneous Play). *All agents submit their bids simultaneously.*

Although bids are, in practice, collected over a certain period, they remain sealed and confidential. This ensures that all bids are (virtually) submitted at the same time, thus justifying Assumption 2. Assumptions 1 and 2 complement each other and allow for a normal-form game formulation.

Assumption 3 (Common Value). *The contract value is common for all agents.*

Assumption 3 simplifies the commonly used Bayesian game formulation of first-price auction to reflect the conditions of Bertrand competition [46], where agents compete on prices in a transparent market. This assumption is supported by the uniformity of contract requirements and regulatory costs across bidders, market transparency, and the homogeneity in the evaluation of contract values. These factors standardize the perceived value of the contract, thereby aligning all agents' bids based on a common economic assessment rather than individual speculative valuations.

Strategies. All players are symmetric in their strategy set. Indeed, each agent i chooses a strategy $s_i \in S$ where $S = \{FP, CP\}$. The strategy *Fair Price (FP)* is associated with bidding the fair price b_i , and *Collusive Price (CP)* with bidding a higher price, defined by $\alpha \cdot b_i$, where $\tau \geq \alpha > 1$, with the upper bound τ ensuring realistic bids. We denote the strategy profile for a given occurrence as the tuple (s_1, s_2, \dots, s_n) .

Heterogeneity. Each agent i has a market power defined by the parameter $\beta_i \in (0, 1)$, with $\sum_{i \in \mathcal{N}} \beta_i = 1$, indicating that the average market strength, $\mu(\beta)$, is always $\frac{1}{n}$. This market power might represent the size of the agent or its market share, influencing its ability to reduce costs and leverage economies of scale. All agents share the same cost function, which determines their bids

$$b_i = (1 - \beta_i)v \quad \forall i \in \mathcal{N}, \quad (1)$$

The notion of strong and weak agents is relative to the distribution of β , but according to (1), the closer β is to 1, the lower the agent can bid, thus the stronger it is. Market heterogeneity is quantified by $\sigma(\beta)$, the standard deviation of the distribution of power parameter values. In a perfectly homogeneous market, $\sigma(\beta) = 0$ and $\beta_i = \frac{1}{n}$ for all i .

Remark 1. *As discussed earlier, we need to model heterogeneous agents within a game formulation that incorporates a social dilemma. Additionally, the payoff structure must reflect a reward function that enables adaptive agents to learn their strategies in repeated settings. Example 1 illustrates why a strict application of the minimum price rule is unsuitable in this context. The core issue is that a strong player could dominate any auction regardless of its strategy, undermining the coexistence of heterogeneous agents.*

Example 1 (Hard minimum price rule). Let b_{weak} and b_{strong} be the bids in an auction within a heterogeneous duopoly. Naturally, and according to (1), we have $b_{strong} < b_{weak}$. Assuming the payoffs u are determined by a minimum price rule, i.e., $u_i = b_i$ if $b_i = \min(b_{weak}, b_{strong})$ and $u_i = 0$ otherwise, we obtain the following bimatrix game:

		weak agent	
		<i>FP</i>	<i>CP</i>
strong agent	<i>FP</i>	0 <i>b_{strong}</i>	0 <i>b_{strong}</i>
	<i>CP</i>	? ?	0 αb_{strong}

We observe that the outcomes of the asymmetric strategy profiles depend on the level of heterogeneity. The strong agent could potentially win regardless of the strategy profile, while the weak agent might struggle to learn a long-term strategy since it may only receive a payoff of 0 as a reward signal in a repeated setting.

Payoffs. Let $u_i(FP, k)$ represent the payoff of agent i when it bids the fair price and k opponents also bid the fair price, and let $u_i(CP, k)$ denote the payoff of agent i when it bids the collusive price while k opponents bid the fair price, where $k \in [0, n - 1]$. The key idea from the minimum price rule is that coordination must be unanimous to achieve collusive profits. Therefore, the individual payoff for playing CP when all players cooperate surpasses the individual payoff from universal defection, i.e.,

$$u_i(CP, k = 0) > u_i(FP, k = n - 1). \quad (2)$$

However, defection must always yield a higher payoff than cooperation when at least one opponent defects, leading to

$$u_i(FP, k > 0) > u_i(CP, k > 0). \quad (3)$$

In fact, because of the minimum price rule, we need to make sure that the cooperating agents receive payoffs of 0 when some other players defect. Also, we set a sharing mechanism where defecting agents earn a share based on their market strengths relative to the defecting opponents. Therefore, let $\Omega \subseteq \mathcal{N}$ denote the set of agents playing FP , and let the total market power of the defecting agents be $\beta_\Omega = \sum_{j \in \Omega} \beta_j$. This means that when agent i plays FP , its payoff,

$$\frac{\beta_i}{\beta_\Omega} \cdot b_i,$$

depends on β_Ω . Note that in this case, $\beta_\Omega = 1$ when $k = n - 1$, and $\beta_\Omega = \beta_i$ when $k = 0$. When agent i plays CP , then its payoff is either 0 if $k > 0$ or

$$\alpha \cdot \beta_i \cdot b_i,$$

if $k = 0$. Table 1 summarizes the individual payoffs for agent i in all scenarios.

Strategy Profile	$k = 0$	$n - 1 > k > 0$	$k = n - 1$
$u_i(FP, k)$	b_i	$\frac{\beta_i}{\beta_\Omega} \cdot b_i$	$\beta_i \cdot b_i$
$u_i(CP, k)$	$\alpha \cdot \beta_i \cdot b_i$	0	0

The payoff function presented in Table 1 respects (2) and (3) since,

$$b_i < \min_{j \neq i} \alpha \cdot b_j \quad \forall i,$$

and

$$\frac{\beta_i}{\beta_\Omega} \cdot b_i > 0.$$

Analysis. The parameter β controls the incentive to defect. As we can derive from (1) and Table 1, and since α is the collusive factor common to all agents, the incentive factor for agent i associated with defecting while every opponent colludes ($FP, k = 0$) is

$$\gamma_i = \frac{1}{\beta_i}. \quad (4)$$

The pairwise relation in terms of incentive factors between agent i and an opponent j is given by

$$\gamma_i = \frac{\beta_j}{\beta_i} \gamma_j. \quad (5)$$

The incentive to defect is proportional to the market power β , and therefore weak agents have more incentive to defect than strong ones, as expected in real-life scenarios [13, 47]. Example 2 proposes a numerical depiction of the incentive to defect in the MPG.

Example 2 (Incentive factor). Consider three heterogeneous agents $\{1, 2, 3\}$, and their associated power parameters $[0.25, 0.25, 0.5]$. For simplicity, we set v to 1 in equation (1). In the collusive scenario, all agents choose to play CP , meaning that $\beta_\Omega = 1$, and according to (1) and Table 1, we have $u_1 = u_2 = 0.19\alpha$ and $u_3 = 0.25\alpha$, with the strongest agent receiving the highest payoff.

Now consider that agent 1, a weak agent, chooses to defect instead. In this case, we have $\beta_\Omega = 0.25$, and thus, $u_1 = 0.75$ and $u_2 = u_3 = 0$. As predicted by (4), the incentive factor for the weak agent to defect is 4.

If the strong agent decides to defect while others cooperate, then $\beta_\Omega = 0.5$, and thus, $u_1 = u_2 = 0$ and $u_3 = 0.5$. Here, the incentive factor for the strong agent to defect is 2. As predicted by (5), the incentive factor for the weak agent to defect is twice as high because its market power ($\beta_1 = \beta_2 = 0.25$) is half that of its stronger opponent.

Our sharing mechanism, using β_ω , allows for heterogeneous agents to coexist while preserving the dynamic stemming from the minimum price. In fact, the power parameter β not only represents the ability to cut costs but also the work capacity. As strong agents are more rewarded than weak ones in symmetrical plays (see Table 1), the payoff structure of the MPG emulates the natural turnover in auction participation of heterogeneous agents observed in markets, making the coexistence between weak and strong agents possible. Indeed, a higher payoff mathematically represents the higher frequency of wins by stronger agents in a strict minimum price-ruled process. When the setting simplifies to a homogeneous game, the sharing mechanism in symmetrical plays smooths out a random turnover between agents of equal power, as shown in Example 3.

Example 3 (Homogeneous special case). Consider a homogeneous duopoly ($\sigma(\beta) = 0$ and $n = 2$). According to (1) and Table 1, we get the following bimatrix game:

		agent 2	
		<i>FP</i>	<i>CP</i>
agent 1	<i>FP</i>	$b/2$	b
	<i>CP</i>	b	$\alpha b/2$

In symmetrical plays, the payoffs are the same for both agents. This is mathematically equivalent to a random selection mechanism for breaking ties between the lowest bidders under the hard minimum price rule. Since homogeneous players produce the same bid, b , a selection mechanism would involve a random process to break ties in symmetrical play. This could be done by either using the uniform distribution, $i^* \sim \mathcal{U}(\mathcal{N})$, or by adding a stochastic perturbation to the bids, $b_i + \epsilon_i$.

Proposition 1. *The n -player homogeneous MPG is a Prisoner's Dilemma.*

Proof. Consider a reference player $i \in \mathcal{N}$. Also, consider $\alpha \leq 2$, a realistic upper bound for the collusive factor. The following payoff structure unfolds by definition:

- T : Temptation payoff (player i bids the fair price while all opponents bid their collusive price),
- R : Reward payoff (everybody cooperates),
- P : Punishment payoff (everybody defects),
- S : Sucker's payoff (player i bids the collusive price while at least one opponent defects).

i. P is the dominant strategy profile.

In the homogeneous MPG, $\beta_i = \frac{1}{n}$ for all $i \in \mathcal{N}$, meaning that $b_i = b$ for all $i \in \mathcal{N}$ according to (1). Hence,

$$P = u_i(FP, k = n - 1) = \frac{b}{n} \quad \text{and} \quad S = u_i(CP, k > 0) = 0,$$

and $P > S$. Since

$$u_i(FP, k = 0) = b > \frac{\alpha b}{n} = u_i(CP, k = 0) \quad \forall i \in \mathcal{N}$$

for any $\alpha < 2$ (since $n \geq 2$), $T > S$, meaning that FP (defection) is the best response for every player. The strategy profile P is therefore the unique Nash equilibrium.

ii. *The strategy profile R is Pareto Optimal.*

A strategy profile is Pareto Optimal if no player can improve their outcome without making at least one other player worse off. Since

$$u_i(CP, k = 0) = \frac{\alpha b}{n} > \frac{b}{n} = u_j(FP, k = n - 1) \quad \forall (i, j) \in \mathcal{N},$$

because $\alpha > 1$, $R > P$, and $(CP, k = 0)$ is therefore Pareto Optimal.

iii. *The payoff structure satisfies the inequality $T > R > P > S$.*

Since

$$u_i(FP, k = 0) = b > \frac{\alpha b}{n} = u_i(CP, k = 0) \quad \forall i \in \mathcal{N},$$

because $\frac{\alpha}{n} < 1$, $T > R$. By combining the results of i. and ii., we conclude $T > R > P > S$.

□

Corollary 1. *The n -player heterogeneous MPG is not a Prisoner's Dilemma.*

Proof. In the heterogeneous MPG, the market powers differ among agents, leading to variations in the payoffs. If the level of heterogeneity is high enough, the strongest agent may receive a higher payoff with R (collusion) than with T (being the only defector). In this case,

$$u_i(CP, k = 0) = \alpha \beta_i b_i \quad \forall i \in \mathcal{N},$$

if $\alpha \beta_i > 1$, then $u_i(CP, k = 0) = R > T = u_i(FP, k = 0)$, and condition iii. in the proof of Proposition 1 no longer holds. □

3.2 Markov Game Formulation

In contrast to iterated games, Markov games evolve over multiple stages, with states dependent on previous actions. This dynamic nature better captures the complexity of real-world markets where strategies evolve over time, especially in the context of behavioral emergence, which may heavily rely on self-strengthening trends. The MPMG extends the framework of the MPG by providing the advantage of the Markov property, that relies on a state space that encodes information about opponents' past plays and characteristics.

Definition 1. Markov game. A Markov game, also known as a stochastic game, is a generalization of a MDP to multiple interacting agents contained in the set N . Each player i has an associated action space \mathcal{A}_i . The game is defined over a common state space \mathcal{S} . The dynamics of the game are captured by a transition function $T : \left(\prod_{i \in N} \mathcal{A}_i\right) \times \mathcal{S} \rightarrow [0, 1]$, which determines the probability of moving from one state to another based on the actions of all players. Each player i has a reward function $R_i : \left(\prod_{i \in N} \mathcal{A}_i\right) \times \mathcal{S} \rightarrow \mathbb{R}$, defining the payoff r_i for player i based on the current state and actions taken by all players. The objective for each player i is to maximize the expected sum of discounted rewards, expressed as $E \left[\sum_{t=0}^{\infty} \gamma^t r_{i,t} \right]$, where γ is the discount factor that values present rewards over future rewards. A key feature of an MDP, and therefore of a Markov game, is the Markov property, which asserts that future states depend only on the current state and the actions taken by the agents, not on the sequence of events that preceded it, ensuring that past states are irrelevant once the current state is known.

Action space. The action space $\mathcal{A} = \{0, 1\}$ is binary, reflecting the two strategies available to players, with bidding the fair price being represented by 0 and bidding the collusive price with 1. This representation simplifies the computation of action and joint action frequencies. Let $a_i^t \in \mathcal{A}$ be the action taken by agent i at time t , and $f^t \in \mathcal{A}^n$ be the joint action (a_1, \dots, a_n) in auction t . The action and joint action frequencies up to iteration t are then respectively given by

$$\bar{a}_i^t = \frac{1}{t} \sum_{m=1}^t a_i^m,$$

and

$$\bar{f}^t = \frac{1}{t} \sum_{m=1}^t f^m.$$

Assumption 4 (Information availability). *All players have access to information about the game status and basic features of their opponents.*

Assumption 4 is supported by the availability of auction data, which results from transparency laws. This assumption is crucial for constructing the state space, as it provides the signals necessary for players to make informed decisions during implementation.

State space. The choice of state space in a Markov game is often at the discretion of the implementer, depending on the specifics of the problem and the goals of the implementation. However, for our purposes, we define the MPMG with core variables that capture essential aspects of the game. Specifically, these sets include all players' information on action frequencies $\bar{\mathbf{a}}^t = \{\bar{a}_i^t\}_{i \in N}$, joint action frequencies $\bar{\mathbf{f}}^t = \{\bar{f}_i^t\}_{i \in N}$, average rewards $\bar{\mathbf{r}}^t = \{\bar{r}_i^t\}_{i \in N}$, and power parameters $\beta = \{\beta_i\}_{i \in N}$. Therefore, the state space at any time t is $\mathcal{S}_t = \{\bar{\mathbf{a}}^t, \bar{\mathbf{f}}^t, \bar{\mathbf{r}}^t, \beta\}$, and is of size $2n + |\mathcal{A}|^n$. This minimal configuration for the state space aligns with Assumption 4. Furthermore, the power parameters also allow for the possibility of modeling a dynamic market in which agents could lose or gain market power.

Reward function. The payoff structure defined for the MPG leads to a reward function in the MPMG, emphasizing the inherent connection between the Markov game and its RL implementation. The reward function, together with the transition function, which defines how the game state evolves in response to actions, fundamentally shapes the dynamics and outcomes of the game. In fact, in the classical RL setting, the normalized reward function for agent i at each iteration t is expressed as

$$G_i^t = (1 - \gamma) \sum_{m=0}^{t-1} \gamma^m r_i^{t-m},$$

where γ is the discount factor and r_i^{t-m} is the immediate reward received by agent i at time $t - m$. The MPMG relies on the tension between Nash and Pareto dominant strategies inherent in the classical Prisoner's Dilemma. However, as with the state space, the reward function can be adjusted for specific implementation purposes (e.g., reward shaping) as long as the fundamental dynamics are preserved. Following Assumption 3, we can further assume that the contract value v remains equivalent as the iterations unfold, normalizing it to 1. Therefore, (1) becomes

$$b_i = 1 - \beta_i.$$

We do not make this assumption formal, as this normalization does not affect the generality of the model but only simplifies the calculations. The immediate rewards for agent i at any time t are given in Table 1.

Assumption 5 (Single Public Contractor). *All public contractors share the same goal.*

In real-world scenarios, the set of public contractors in a given procurement market usually includes many entities. Despite this plurality, it is safe to assume that all public contractors work towards the same overarching goal. Assumption 5 allows us to reduce the set of public contractors to a singleton.

Dynamic. The MPMG unfolds as an episodic game of arbitrary length (number of episodes), with each episode consisting of a single-step auction (Assumptions 1 and 2). Agents compete against each other in an environment representing the same entity in each episode (Assumption 5). Due to the Markov property, memory is encapsulated in the observation space \mathcal{S} at any episode t . Agents use this observation space to learn their strategies, effectively playing a Markov game where the reward function, $R(a_1, \dots, a_n)$, is deterministic and based on the joint action. The transition probability function is given by $P(s'|s, a)$, where s' is the new state and $s \in \mathcal{S}$. When the state space is not used for decision-making, the agents are effectively playing an iterated MPG.

4 Computer Experiments

Cooperation in MARL settings is well studied. However, this algorithmic cooperation mostly relies on explicit designs—such as shared joint action spaces [48], algorithmic communication [49], reputation-based models [50]—which are not suitable for our context, as these approaches emulate scenarios where agents deliberately bias their actions towards cooperation.

In this study, we evaluate the inherent properties of the MPMG, requiring unbiased players. If collusion occurs, it would arise solely from the game’s intrinsic structure. Therefore, we employed basic yet established state-of-the-art MARL algorithms to define artificial agents that are symmetric in their decision-making processes, do not communicate, and do not hold beliefs about their opponents. In other words, we aim to create isolated and equally rational agents.

We conducted a total of twenty-four experiments, combining six agent types with four MPMG configurations (2-player homogeneous, 5-player homogeneous, 2-player heterogeneous, and 5-player heterogeneous MPMG). As discussed in Section 3, $0 \leq \sigma(\beta) \leq 1$ represents the standard deviation of the power parameters (β) distribution, with $\sigma(\beta) = 0$ representing the homogeneous case. For experiments in the heterogeneous setting, we set $\sigma(\beta) = 0.5$. Additionally, we set $\alpha = 1.3$, meaning that the collusive bids are 30% higher than the fair bids. Each agent was trained on 100 auctions over 100 replications for statistical robustness.

4.1 Agent Representation

We based our agents on three foundational approaches of MARL, that is, Multi-Armed Bandits (MAB), Deep Q-learning, and Actor-Critic Policy Gradient. Implementation details are available in Appendix B.

Definition 2. Multiagent Reinforcement Learning. MARL extends single-agent RL techniques to settings with multiple agents learning concurrently in a shared environment. In RL, an agent interacts with its environment to maximize cumulative rewards, where key components include agent, environment, state, action, reward, policy, and value function. The agent, acting as the decision-maker, responds to the environment, which consists of external factors and situations represented as states. Actions are possible decisions the agent can take, evaluated by rewards as feedback on their success. The policy is the agent’s strategy for selecting actions based on the current state, while the value function estimates future rewards to determine the desirability of states or actions. In MARL, multiple agents simultaneously learn and update their policies, causing the environment to appear non-stationary from any single agent’s perspective. The dynamics lead to more complex interactions compared to single-agent RL, requiring robust techniques to address the challenges of Markov games and multiagent systems in general.

The MPG as a bandit problem. MAB algorithms [51] represent a simplified form of RL or MARL, focusing on the exploration-exploitation trade-off without considering future states (*stateless*), that is, without a model of the environment. MAB methods solve the bandit problem, where the decision-maker must sequentially choose among different actions with unknown reward distributions in order to maximize the total reward over time. Each action represents a distinct stochastic process with its own, a priori unknown, probability distribution of rewards. The goal is to minimize *regret*, which is the difference between the total reward the decision-maker could have earned by always choosing the best action and the actual reward obtained by balancing exploration and exploitation. MAB agents actually play an iterated MPG, as they do not utilize the observation space. We deployed ϵ -greedy agents [52] that balance exploitation and exploration by selecting a random action with a probability ϵ decreasing over time, Bayesian Thompson Sampling (TS) agents [53] that maintain a probability distribution for the expected reward of each action, and Upper Confidence Bound (UCB) agents [54] that select the action with the highest upper confidence bound of the estimated return.

Deep Q-learning. Methods based on Q-learning [55] maintain state-action pair values, named Q-values, based on the state-action pair expected reward. Decision making is then conducted following the highest Q-values associated with the state-action pair in question. In our experiments, we used the Double Dueling Deep Q-network (D3QN) algorithm [56]. In the D3QN, a multi-layer perceptron is used to approximate the Q-values, and the Double Q-learning approach is employed to reduce overestimation bias by maintaining two sets of Q-value estimates. The Dueling architecture, on the other hand, decomposes the Q-value into two separate streams representing the state-value function and the advantage function, which helps the model learn better representations of the value of being in a given state, independent of the specific actions.

Actor-critic policy gradient. Unlike Q-learning, policy gradient methods focus on optimizing a parameterized policy $\pi_\theta(a|s)$ directly by gradient ascent on expected rewards [55]. In the actor-critic framework, two function approximators are employed. The *actor*, which is responsible for selecting actions based on the current policy, and the *critic*, which evaluates the action-value function to provide feedback on the quality of the chosen actions. The actor updates the policy in the direction suggested by the critic, while the critic learns to estimate the expected return, enabling a stable and efficient training process by reducing the variance of the policy gradient estimates. Here, we used Multi-Agent Proximal Policy Optimization (MAPPO) agents [57] that extend the Proximal Policy Optimization (PPO), a popular policy gradient method, to multi-agent settings. The core idea of PPO is to strike a balance between policy improvement and preventing drastic updates to the policy, which could harm performance. To achieve this, PPO uses a modified objective function that constrains how much the new policy is allowed to diverge from the old policy during an update.

4.2 Results

Before presenting the results of the experiments, it is essential to outline several important considerations for evaluating artificial agents in the multi-agent context. Indeed, the interpretation of the results depends on both individual and collective convergence. By individual convergence, we mean the convergence of the algorithms themselves. For D3QN and MAPPO, this is assessed by examining training loss function values over time. In the case of MAPPO, both the actor and critic networks must converge. In our experiment, they all converged, as can be seen in Figure 4. Additionally, D3QN and MAPPO offer internal metrics indicating their policies. For MAPPO, we can directly observe the evolution of the policies regarding actions *FP* and *CP*, while D3QN uses Q-values for each action. The evolution of MAPPO policy values and D3QN Q-values during training is shown in Figure 5. In the case of the stateless bandit algorithms (ϵ -greedy, TS, and UCB), cumulative regrets are generally used to establish learning in the single-agent context, though they are less meaningful in Markov games. Nevertheless, cumulative regrets are presented in Figure 6. Figures 4, 5, and 6 are provided in Appendix A.

However, the joint action frequencies, or the frequencies at which each strategy profile is played, are metrics common to all experiments, and they provide a clear picture of collective convergence. Therefore, they offer a reliable way to evaluate and compare our MARL agents in terms of behavioral trends. In the single-stage MPMG (MPG) or in similar static games, three distinguishable collective outcomes emerge: agents play the Nash profile, the Pareto profile, or they diverge in their choices, meaning that only a subgroup of players share a nonzero return associated with the Nash play. In dynamic settings, outcomes are based on convergence toward a particular outcome, with agents represented by stochastic processes typically converging toward mixed strategies. The challenge, then, is fitting a discrete interpretation (do they cooperate or not?) onto a continuous spectrum of behaviors.

The evolution of the average joint action frequencies over training episodes is displayed in Figure 3 in Appendix A. The average joint action frequencies from the last training episode for all experiments and outcomes are displayed in Table 2, showcasing the impact of atomicity and heterogeneity on the agents’ ability to coordinate, as well as the different behavioral trends emerging from each configuration of agents and market parameters.

In addition, Figure 1 ranks each agent-MPMG configuration pair in terms of collusive potential by using scores based on the data in Table 2. More specifically, each value in Figure 1 represents the difference between the average joint action frequency of the last training episode of the Nash and Pareto outcomes. The resulting value is then scaled from 0 (most Nash) to 1 (most Pareto) using the *minmax* formula.


4.3 Analysis

Our results, summarized in Table 2 and Figure 1, highlight the varying effectiveness of different MARL approaches in the MPMG. We can see that, in the 5-player heterogeneous MPMG, and in contrast to the results of the 2-player homogeneous MPMG, agents always converge towards a suboptimal outcome far from the Pareto profile, but closer to

Table 2: Average joint action frequencies across repeats at the last training episode for each outcome in each experiment. Bold values indicate the highest average joint action frequency among the three outcomes.

	$n = 2, \sigma(\beta) = 0.0$			$n = 2, \sigma(\beta) = 0.5$			$n = 5, \sigma(\beta) = 0.0$			$n = 5, \sigma(\beta) = 0.5$		
	<i>FP</i>	<i>CP</i>	Other	<i>FP</i>	<i>CP</i>	Other	<i>FP</i>	<i>CP</i>	Other	<i>FP</i>	<i>CP</i>	Other
D3QN	0.37	0.16	0.47	0.37	0.17	0.47	0.07	0.02	0.91	0.07	0.02	0.91
TS	0.72	0.04	0.24	0.69	0.04	0.27	0.41	0.01	0.59	0.32	0.01	0.67
ϵ -greedy	0.70	0.03	0.27	0.70	0.03	0.27	0.35	0.01	0.64	0.35	0.01	0.64
MAPPO	0.52	0.10	0.38	0.51	0.11	0.39	0.24	0.01	0.75	0.21	0.01	0.78
UCB	0.40	0.58	0.02	0.43	0.55	0.02	0.34	0.43	0.23	0.22	0.08	0.70

Figure 1: One-dimensional heat-map highlighting the relative difference among agents and MPMG configurations in terms of their collusive potential. From the most Nash (left) to the most Pareto (right) configuration.



the Nash Equilibrium. This shows the impact of atomicity and heterogeneity on the capacity of agents to coordinate themselves, let alone cooperate in collusion.

Also, UCB agents seem to favor the Pareto profile, consistently across different configurations. In fact, UCB is the only algorithm that not only fosters higher rates of convergence towards the Pareto profile but also shows robustness against atomicity and heterogeneity, at least when they are increased individually. This is clear in Figure 1, where we can see that the three highest scores belong to UCB agents. However, we can see in Table 2 that in the 5-player heterogeneous MPMG, UCB agents behave similarly to their counterparts.

Moreover, it seems there is no correlation between algorithmic sophistication and collusive potential. Stateless bandits can either converge towards an approximate Nash Equilibrium (TS, ϵ -greedy) or favor the Pareto outcome (UCB), while methods using more information either favor the Nash Equilibrium (MAPPO) or other suboptimal outcomes (D3QN).

5 Discussion

In this section, we discuss, interpret, and put our results into perspective.

5.1 Interpretation


Three key insights emerge from the computer experiments.

First, the MPMG is a reliable approximation of market dynamics, as it exhibits properties consistent with Jean Tirole’s seminal work on market collusion potential [47]. Specifically, as shown in Section 3.1, in the MPMG, the incentive to defect while opponents cooperate is asymmetrical with respect to the power parameter β , thereby undermining the intrinsic value of the Pareto strategy. This effect strengthens as the number of agents increases, further reducing the likelihood of successful collusion.

Second, tacit coordination relies heavily on self-reinforcing trends. Indeed, UCB displays counter-intuitive behaviors for a no-regret algorithm. In such settings, the regret of not defecting is higher than the regret associated with not cooperating (see Table 1). Additionally, other no-regret algorithms, TS and ϵ -greedy, favor the Nash Equilibrium. The particularity of UCB, at least in the original version we implemented [54], is that it is a deterministic method because of the way exploration is handled. Therefore, when fully symmetric UCB agents populate the game, coordination towards the Pareto outcome builds upon self-reinforcing beliefs as every agent acts the same. As we can see from Figure 2a, the confidence intervals for the average (across repeats) joint action frequencies associated with defection and cooperation of the UCB agents are null. In contrast, Figure 2b shows non-null confidence intervals for MAPPO agents, and Figure

3 in Appendix A shows that only UCB-related graphs can have null confidence intervals. Furthermore, it has been shown that UCB agents are the RL agents most prone to cooperate in the Iterated Prisoner’s Dilemma, with cooperation rates similar to ours [58].

Figure 2: Average joint action frequencies of UCB (a) and MAPPO (b) agents in the 2-player homogeneous MPMG.



Finally, more information does not lead to more cooperation. There appears to be no positive correlation between the sophistication of our AI agents and the emergence of tacit coordination, a conclusion we share with other authors [58, 43, 59]. In fact, in the Prisoner’s Dilemma, no amount of information can a priori defuse the tension created by the existence of the Nash Equilibrium, as long as agents do not engage in self-reinforcing trust over cooperation.

We can therefore conclude that, provided the MPMG reasonably approximates minimum price-ruled auctions, and although collusion can occur in the most naive settings, the minimum price rule is resilient to non-engineered tacit coordination among rational actors, as is any rule preserving the structure of a social dilemma.

5.2 Practical Considerations

This study focuses on evaluating the behavior of artificial agents without explicitly quantifying their collusive potential in terms of collusive gains. Although the Pareto outcome provides higher returns compared to the Nash outcome, agents cannot strictly maximize their returns, as this would involve behavior divergent from typical learning processes. Moreover, even during evaluation—when agents no longer explore and strictly follow their policies—they often converge to mixed strategies rather than a pure Nash or Pareto strategy. Additionally, unlike bandits, D3QN and MAPPO map a policy to state observation, meaning that even when they learn the pure Pareto profile, their collusive gains are highly context-dependent. As a result, their returns in evaluation remain stochastic, and assessing them in terms of collusive gains may be both challenging and irrelevant.

It is also important to note that, in the context of the MPMG, other (non-Nash) suboptimal outcomes can be viewed as rotation schemes, that is, turnovers in the winning set of players. This kind of strategy emerges from mixed behavior and can lead to greater returns than a consistent Nash play.

5.3 Algorithmic Pricing and Legal Perspective

Our results align with specialized studies that reveal the challenge of achieving algorithmic cooperation, whether tacit or explicit [60]. However, they do not contradict the legal literature, where scholars often assert that algorithmic collusion is relatively straightforward to achieve [61, 62, 4, 63]. This study offers a framework and benchmark for tacit collusion and helps explain the contradictory reports in the literature by shedding light on the opaque nature of algorithmic tacit coordination.

Nonetheless, it is essential to add nuance for a comprehensive understanding of this complex issue. While legal concerns stem from tangible evidence found in digital open markets and auctions (see Section 1), algorithmic implicit coordination has been shown to be impeded in the specific context of unbiased social dilemmas, as discussed in Section 2. Additionally, the potential for algorithmic and tacit collusion in transparent markets (including public markets) will always exist. Complex cooperative behaviors can be engineered in virtually any large-scale agent system [40], even if such implementations remain challenging. For instance, cooperation in competitive settings can be elusive even when explicit communication is permitted [64, 65]. What is certain, at least from our perspective, is that further research

is required to establish the precise theoretical and algorithmic conditions, as well as the practical considerations, regarding algorithmic collusion in public

5.4 Future Work

Extending the MPMG. Future research should aim to extend the MPMG in both quantitative and qualitative dimensions. Quantitatively, this could involve creating environments that more closely mimic actual simulations, using approaches like empirical games [66]. Qualitatively, it could involve relaxing certain assumptions, such as the common value assumption (Assumption 3), particularly in highly heterogeneous market scenarios. Ultimately, developing a realistic, full-scale simulation of such a system would also necessitate a continuous formulation for the action space and a modular environment allowing for the implementation of various supplier selection models, thus generalizing the MPMG to a general auction Markov game.

Additionally, the public contractor could be modeled as an active participant in the environment, represented by an adaptive learning agent. This scenario has been explored by [34], where reinforcement learning was used in the context of Stackelberg games as an algorithmic defense against algorithmic pricing for e-commerce platforms. Practically, these perspectives could shape the future of more robust E-procurement designs, and theoretically, they offer a computer-based experimental framework that complements laboratory-based behavioral studies.

AI rationality and cyber cartels. In an AI-governed world, the concept of rationality takes on new dimensions, particularly in economic interactions and market behavior. The experiments in this study operate under a set of restrictions to maintain an environment unbiased towards cooperation. However, as mentioned in Section 4, market players could align on algorithmic pricing using centralized learning and execution (i.e., a shared joint action space). This raises the question: could tacit collusion, emerging from a deliberately biased algorithmic structure, be easily detectable?

The potential existence of cyber cartels introduces new research opportunities into how agents might evade detection and punishment by law enforcers through planned algorithmic collusion. Such a consortium could gain advantages in both efficiency and stealth. Explicit coding for collusion would typically require real-life coordination among market players, potentially leaving traces and being susceptible to algorithmic auditing by legal authorities. However, implicit algorithmic coordination can either be programmed or accidental, creating a grey zone that complicates regulatory efforts. It is also worth noting that both digitized and non-digitized markets can be affected, as the only requirement for algorithmic pricing is the availability of data.

References

- [1] J. A. Gerlick and S. M. Liozu, "Ethical and legal considerations of artificial intelligence and algorithmic decision-making in personalized pricing," *Journal of Revenue and Pricing Management* **19** (4, 2020) 85–98. <https://link.springer.com/article/10.1057/s41272-019-00225-2>.
- [2] L. L. Gormsen, "Algorithmic antitrust and consumer choice," *Economic Analysis of Law in European Legal Scholarship* **12** (2022) 65–86.
- [3] A. Ezrachi and M. Stucke, "Sustainable and unchallenged algorithmic tacit collusion," *Northwestern Journal of Technology and Intellectual Property* **17** (2020) .
- [4] Q. Li, N. Philipsen, C. Cauffman, and C. C. Ni, "Ai-enabled price discrimination as an abuse of dominance: a law and economics analysis," *China-EU Law Journal* **2023 9:1 9** (4, 2023) 51–72. <https://link.springer.com/article/10.1007/s12689-023-00099-z>.
- [5] E. Calvano, G. Calzolari, V. Denicolò, and S. Pastorello, "Artificial intelligence, algorithmic pricing, and collusion," *American Economic Review* **110** (10, 2020) 3267–97.
- [6] S. Assad, R. Clark, D. Ershov, and L. Xu, "Algorithmic pricing and competition: Empirical evidence from the german retail gasoline market," *SSRN Electronic Journal* (11, 2020) . <https://papers.ssrn.com/abstract=3682021>.
- [7] M. H. Rothkopf, "Daily repetition: A neglected factor in the analysis of electricity auctions," *The Electricity Journal* **12** (1999) 60–70. <https://sci-hub.st/https://www.sciencedirect.com/science/article/pii/S104061909900010X>.
- [8] Y. Ye, D. Qiu, M. Sun, D. Papadaskalopoulos, and G. Strbac, "Deep reinforcement learning for strategic bidding in electricity markets," *IEEE Transactions on Smart Grid* (8, 2019) 1–1.
- [9] A. C. Tellidou and A. G. Bakirtzis, "Agent-based analysis of capacity withholding and tacit collusion in electricity markets," *IEEE Transactions on Power Systems* **22** (11, 2007) 1735–1742.
- [10] J. Viehmann, S. Lorenczik, and R. Malischek, "Multi-unit multiple bid auctions in balancing markets: An agent-based q-learning approach," *Energy Economics* **93** (2021) .
- [11] I. Matsukawa, "Detecting collusion in retail electricity markets: Results from japan for 2005 to 2010," *Utilities Policy* **57** (4, 2019) 16–23.
- [12] B. Baranek, V. Titl, and L. Musolff, "Detection of collusive networks in e-procurement," *SSRN Electronic Journal* (6, 2021) . <https://papers.ssrn.com/abstract=3864186>.
- [13] J.-J. Laffont and D. Martimort, "Collusion under asymmetric information," *Econometrica* **65** (1997) 875. <https://www.researchgate.net/publication/4898655>.
- [14] B. Lebrun, "Uniqueness of the equilibrium in first-price auctions," *Games and Economic Behavior* **55** (4, 2006) 131–151.
- [15] D. Bergemann, B. Brooks, and S. Morris, "First-price auctions with general information structures: Implications for bidding and revenue," *Econometrica* **85** (1, 2017) 107–143. <https://onlinelibrary.wiley.com/doi/full/10.3982/ECTA13958><https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA13958><https://onlinelibrary.wiley.com/doi/10.3982/ECTA13958>
- [16] Y. Chassin and M. Joanis, "Détecter et prévenir la collusion dans les marchés publics en construction: Meilleures pratiques favorisant la concurrence," 2010. <https://www.researchgate.net/publication/254398401%0AD%C3%A9tecter>.
- [17] J. E. H. Jr., "Detecting cartels," *Handbook of Antitrust Economics* (2008) 213–258.
- [18] R. P. McAfee and J. Mcmillan, "Bidding rings," 1992. <https://www.jstor.org/stable/2117323>.
- [19] B. K. O. Tas, "Collusion detection in public procurement with limited information," *SSRN Electronic Journal* (2017) .
- [20] S. Chassang and J. Ortner, "Collusion in auctions with constrained bids: Theory and evidence from public procurement," *Journal of Political Economy* **127** (2019) 2269–2300.
- [21] E. Calvano, G. Calzolari, V. Denicolò, J. E. Harrington, and S. Pastorello, "Protecting consumers from collusive prices due to ai," *Science* **370** (11, 2020) 1040–1042. <https://www.science.org/doi/10.1126/science.abe3796>.

-
- [22] R. Axelrod, "The emergence of cooperation among egoists," *American Political Science Review* **75** (6, 1981) 306–318. <https://www.cambridge.org/core/journals/american-political-science-review/article/abs/emergence-of-cooperation-among-egoists/EEAB3C6460F5BC63A4DE813E1B010B21>.
- [23] I. Sadoune, M. Joanis, and A. Lodi, "Implementing a hierarchical deep learning approach for simulating multilevel auction data," *Computational Economics* **2024** (5, 2024) 1–28. <https://link.springer.com/article/10.1007/s10614-024-10622-4>.
- [24] C. Graf, V. Zobernig, J. Schmidt, and C. Klöckl, "Computational performance of deep reinforcement learning to find nash equilibria," *Computational Economics* (2023) .
- [25] M. Bichler, M. Fichtl, and M. Oberlechner, "Computing bayes nash equilibrium strategies in auction games via simultaneous online dual averaging," *EC 2023 - Proceedings of the 24th ACM Conference on Economics and Computation* (8, 2022) 294. <https://arxiv.org/abs/2208.02036v2>.
- [26] G. Brero, B. Lubin, and S. Seuken, "Probably approximately efficient combinatorial auctions via machine learning," *Proceedings of the AAAI Conference on Artificial Intelligence* **31** (2, 2017) 397–405. <https://ojs.aaai.org/index.php/AAAI/article/view/10624>.
- [27] M. Beyeler, G. Brero, B. Lubin, and S. Seuken, "imlca: Machine learning-powered iterative combinatorial auctions with interval bidding,". <https://dl.acm.org/doi/10.1145/3465456.3467535>.
- [28] P. Hummel and R. P. McAfee, "Machine learning in an auction environment," *WWW 2014 - Proceedings of the 23rd International Conference on World Wide Web* **17** (2014) 7–17.
- [29] K. T. Nguyen, "A bandit learning algorithm and applications to auction design," *Advances in Neural Information Processing Systems* **33** (2020) 12070–12079.
- [30] G. Gao, H. Huang, M. Xiao, J. Wu, Y. E. Sun, and S. Zhang, "Auction-based combinatorial multi-armed bandit mechanisms with strategic arms," *Proceedings - IEEE INFOCOM 2021-May* (5, 2021) .
- [31] S. Basu and A. Sankararaman, "Double auctions with two-sided bandit feedback,". <https://arxiv.org/abs/2208.06536v1>.
- [32] D. Goktas, S. Zhao, and A. Greenwald, "Zero-sum stochastic stackelberg games," *Advances in Neural Information Processing Systems* **35** (12, 2022) 11658–11672.
- [33] D. Fudenberg and J. Tirole, *Game Theory*. MIT Press, 1991.
- [34] G. Brero, A. Eden, D. Chakrabarti, M. Gerstgrasser, A. Greenwald, V. Li, and D. C. Parkes, "Stackelberg pomdp: A reinforcement learning approach for economic design,". <https://arxiv.org/abs/2210.03852v3>.
- [35] G. Brero, A. Eden, M. Gerstgrasser, D. Parkes, and D. Rheingans-Yoo, "Reinforcement learning of sequential price mechanisms," *Proceedings of the AAAI Conference on Artificial Intelligence* **35** (5, 2021) 5219–5227. <https://ojs.aaai.org/index.php/AAAI/article/view/16659>.
- [36] G. Brero, D. Chakrabarti, A. Eden, M. Gerstgrasser, V. Li, and D. Parkes, "Learning stackelberg equilibria in sequential price mechanisms," *Proc. ICML Workshop for Reinforcement Learning Theory*, 2021.
- [37] A. Skrzypacz and H. Hopenhayn, "Tacit collusion in repeated auctions," *Journal of Economic Theory* **114** (1, 2004) 153–169.
- [38] L. Waltman and U. Kaymak, "Q-learning agents in a cournot oligopoly model," *Journal of Economic Dynamics and Control* **32** (10, 2008) 3275–3293.
- [39] S. O. Kimbrough, M. Lu, and F. Murphy, *Learning and Tacit Collusion by Artificial Agents in Cournot Duopoly Games*, pp. 477–492. Springer-Verlag, 12, 2005.
- [40] J. K. Gupta, M. Egorov, and M. Kochenderfer, "Cooperative multi-agent control using deep reinforcement learning," pp. 66–83. 2017.
- [41] J. Foerster, R. Y. Chen, M. Al-Shedivat, S. Whiteson, P. Abbeel, and I. Mordatch, "Learning with opponent-learning awareness," *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS* **1** (2018) 122–130.
- [42] M. Harper, V. Knight, M. Jones, G. Koutsououlos, N. E. Glynatsi, and O. Campbell, "Reinforcement learning produces dominant strategies for the iterated prisoner's dilemma," *PLOS ONE* **12** (12, 2017) e0188046. <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0188046>.
- [43] V. Vassiliades, A. Cleanthous, and C. Christodoulou, "Multiagent reinforcement learning: Spiking and nonspiking agents in the iterated prisoner's dilemma," *IEEE Transactions on Neural Networks* **22** (4, 2011) 639–653.

-
- [44] K. Li and D. Hao, “Cooperation enforcement and collusion resistance in repeated public goods games,” *Proceedings of the AAAI Conference on Artificial Intelligence* **33** (7, 2019) 2085–2092.
- [45] R. Chaudhuri, K. Mukherjee, R. Narayanam, and R. D. Vallam, “Collaborative reinforcement learning framework to model evolution of cooperation in sequential social dilemmas,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **12712 LNAI** (2021) 15–26. https://link.springer.com/chapter/10.1007/978-3-030-75762-5_2.
- [46] M. R. Baye and D. Kovenock, “Bertrand competition,” *The New Palgrave Dictionary of Economics* (2008) 1–7. <https://www.semanticscholar.org/paper/783221d378c9fe1ff3839b3ba5bf33c10e258705>.
- [47] J. Tirole, “The economics of tacit collusion,” *Development* (2003) .
- [48] Y. Zhou, S. Liu, Y. Qing, K. Chen, T. Zheng, Y. Huang, J. Song, and M. Song, “Is centralized training with decentralized execution framework centralized enough for marl?”. <https://arxiv.org/abs/2305.17352v1>.
- [49] S. Sukhbaatar, A. Szlam, and R. Fergus, “Learning multiagent communication with backpropagation,” *Advances in Neural Information Processing Systems* (2016) 2252–2260.
- [50] K. Zhang, Z. Yang, and T. Başar, “Multi-agent reinforcement learning: A selective overview of theories and algorithms,” *Studies in Systems, Decision and Control* **325** (2021) 321–384.
- [51] T. Lattimore and C. Szepesvári, *Bandit algorithms*. Cambridge University Press, 7, 2020. https://books.google.com/books/about/Bandit_Algorithms.htmlid=bbjpdWaaQBAJ.
- [52] J. Vermorel and M. Mohri, “Multi-armed bandit algorithms and empirical evaluation,”.
- [53] N. Gupta, O.-C. Granmo, and A. Agrawala, “Thompson sampling for dynamic multi-armed bandits,” *2011 10th International Conference on Machine Learning and Applications and Workshops* **1** (2011) 484–489.
- [54] P. Auer, N. Cesa-Bianchi, and P. Fischer, “Finite-time analysis of the multiarmed bandit problem,” *Machine Learning* **47** (5, 2002) 235–256. <https://link.springer.com/article/10.1023/A:1013689704352>.
- [55] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction, second edition*. The MIT Press, 2018.
- [56] H. V. Hasselt, A. Guez, and D. Silver, “Deep reinforcement learning with double q-learning,” pp. 2094–2100. 2016. www.aaai.org.
- [57] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv.org* (2017) .
- [58] B. Lin, D. Bouneffouf, and G. Cecchi, “Online learning in iterated prisoner’s dilemma to mimic human behavior,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **13631 LNCS** (2022) 134–147. https://link.springer.com/chapter/10.1007/978-3-031-20868-3_10.
- [59] F. Marty and T. Warin, “Deciphering algorithmic collusion: Insights from bandit algorithms and implications for antitrust enforcement,” *Journal of Economy and Technology* (10, 2024) . <https://linkinghub.elsevier.com/retrieve/pii/S2949948824000519>.
- [60] J. Miklós-Thal and C. Tucker, “Collusion by algorithm: Does better demand prediction facilitate coordination between sellers?” *Management Science* **65** (4, 2019) 1552–1561.
- [61] A. Ezrachi and M. E. Stucke, “How pricing bots could form cartels and make things more expensive,” *Havard Business Review* (2016) . <https://hbr.org/2016/10/how-pricing-bots-could-form-cartels-and-make-things-more-expensive?autocomplete=true>.
- [62] A. Ezrachi and S. M. E., “Virtual competition: The promise and perils of the algorithm-driven economy,” *Havard Press University* (2016) .
- [63] E. J. Green, R. C. Marshall, and L. M. Marx, “Tacit collusion in oligopoly,” *The Oxford Handbook of International Antitrust Economics* **2** (2015) 1–25. www.oxfordhandbooks.com.
- [64] J. E. Harrington, R. H. Gonzalez, and P. Kujal, “The relative efficacy of price announcements and express communication for collusion: Experimental findings,” *Journal of Economic Behavior and Organization* **128** (8, 2016) 251–264.
- [65] M. A. Fonseca and H.-T. Normann, “Explicit vs. tacit collusion-the impact of communication in oligopoly experiments,” *European Economic Review* **56** (2012) 1759–1772.
- [66] M. P. Wellman, K. Tuyls, and A. Greenwald, “Empirical game-theoretic analysis: A survey,”. <https://arxiv.org/abs/2403.04018v1>.

A Additional Figures

Figure 3: Average (across repeats) joint action frequencies over training episodes. Rows (up to down): UCB, ϵ -greedy, TS, D3QN, MAPPO. Columns (left to right): $(n = 2, \sigma(\beta) = 0.0)$, $(n = 2, \sigma(\beta) = 0.5)$, $(n = 5, \sigma(\beta) = 0.0)$, $(n = 5, \sigma(\beta) = 0.5)$.

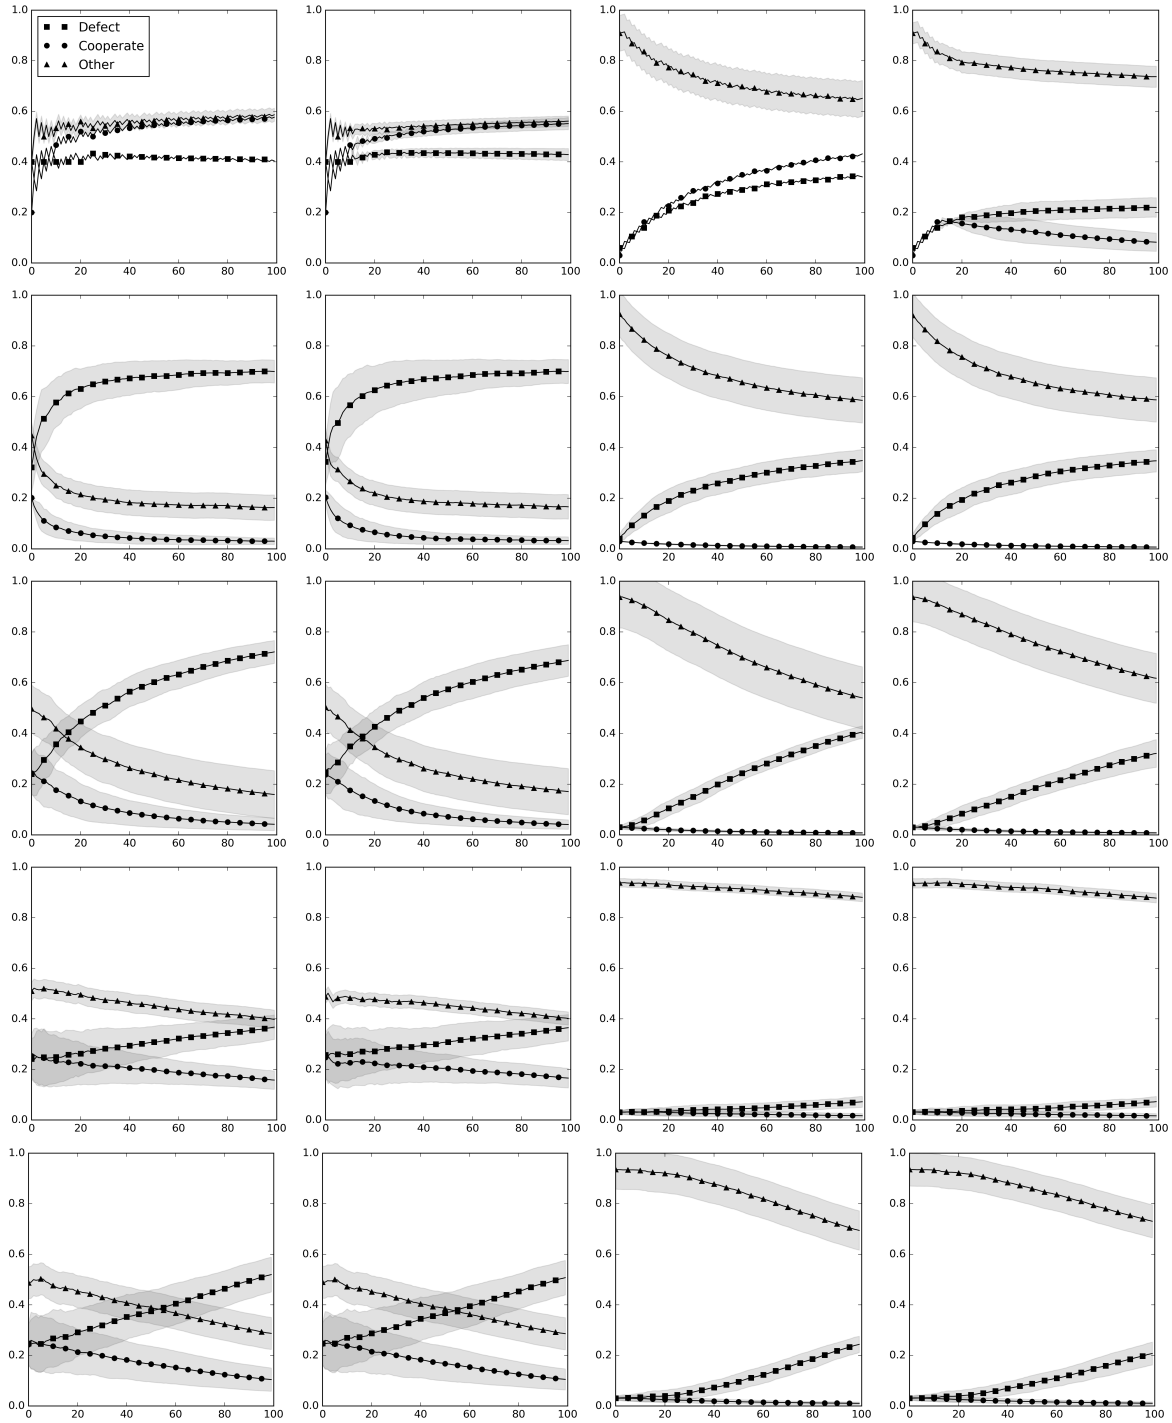


Figure 4: Average (across repeats) training losses. Rows (up to down): D3QN, MAPPO Actor, MAPPO Critic. Columns (left to right): $(n = 2, \sigma(\beta) = 0.0)$, $(n = 2, \sigma(\beta) = 0.5)$, $(n = 5, \sigma(\beta) = 0.0)$, $(n = 5, \sigma(\beta) = 0.5)$.




Figure 5: Average (across repeats) training policy values. Rows (up to down): D3QN Q-values, MAPPO policy values. Columns (left to right): $(n = 2, \sigma(\beta) = 0.0)$, $(n = 2, \sigma(\beta) = 0.5)$, $(n = 5, \sigma(\beta) = 0.0)$, $(n = 5, \sigma(\beta) = 0.5)$.





Figure 6: Average (across repeats) training losses. Rows (up to down): D3QN, MAPPO Actor, MAPPO Critic. Columns (left to right): $(n = 2, \sigma(\beta) = 0.0)$, $(n = 2, \sigma(\beta) = 0.5)$, $(n = 5, \sigma(\beta) = 0.0)$, $(n = 5, \sigma(\beta) = 0.5)$.



B Implementation Details

To replicate the study, follow GitHub.

E-Greedy Method. Our `EpsilonGreedyMAB` class utilizes an ϵ -greedy strategy with the following hyperparameters:

- `n_arms`: Number of possible actions (arms) for the bandit.
- `epsilon`: Exploration rate, set to 0.3 in the `initialize_agents` function.

The agent selects a random action with probability ϵ and the best known action with probability $1 - \epsilon$. The action values are updated using an incremental average formula

$$Q_{t+1}(a) = Q_t(a) + \frac{1}{N(a)}(R_t - Q_t(a))$$

where $Q_t(a)$ is the estimated value of action a at time t , $N(a)$ is the number of times action a has been selected, and R_t is the reward received at time t .

Upper Confidence Bound. Our `UCBMAB` class implements the UCB algorithm with the following hyperparameters:

- `n_arms`: Number of possible actions (arms) for the bandit.

The agent selects actions based on confidence bounds calculated as:

$$\text{UCB}_t(a) = Q_t(a) + \sqrt{\frac{2 \log t}{N(a)}}$$

where $Q_t(a)$ is the estimated value of action a at time t , $N(a)$ is the number of times action a has been selected, and t is the total number of pulls. Action values are updated using the incremental average formula mentioned above.

Thompson Sampling. Our `ThompsonSamplingMAB` class uses a Bayesian approach to action selection with the following hyperparameters:

- `n_arms`: Number of possible actions (arms) for the bandit.
- `alpha`: Beta distribution parameter
- `beta`: Beta distribution parameter

The agent maintains Beta distributions for each action, initialized with $\alpha = 1$ and $\beta = 1$ for each arm. Actions are selected based on samples drawn from these Beta distributions. The Beta distribution parameters are updated based on the received binary rewards as follows:

$$\begin{aligned}\alpha_{\text{new}} &= \alpha_{\text{old}} + \text{reward} \\ \beta_{\text{new}} &= \beta_{\text{old}} + 1 - \text{reward}\end{aligned}$$

This method effectively balances exploration and exploitation by sampling from the probability distributions of each action's expected reward.

Deep Dueling DQN (D3QN). The `DuelingDQN` class implements a Dueling Deep Q-Network with the following architecture:

- `feature_layer`: A fully connected layer with 128 units followed by a ReLU activation function.
- `value_stream`: A fully connected layer with 128 units followed by a ReLU activation function, and another fully connected layer with 1 unit representing the state value.
- `advantage_stream`: A fully connected layer with 128 units followed by a ReLU activation function, and another fully connected layer with `action_dim` units representing the advantage values for each action.

The Q-values are computed as:

$$Q(s, a) = V(s) + \left(A(s, a) - \frac{1}{|\mathcal{A}|} \sum_{a'} A(s, a') \right)$$

where $V(s)$ is the state value, $A(s, a)$ is the advantage of action a in state s , and \mathcal{A} is the set of all possible actions.

Multiagent Proximal Policy Optimization (MAPPO). The `MLPNetwork` class implements a multi-layer perceptron network with the following architecture:

- `state_dim`: Dimensionality of the state space.
- `action_dim`: Dimensionality of the action space.
- `layers`: A fully connected network with two hidden layers, each with 128 units followed by a ReLU activation function.

Our `MAPPOAgent` class implements a MAPPO agent with the following hyperparameters:

- `state_dim`: Dimensionality of the state space.
- `action_dim`: Dimensionality of the action space.
- `lr`: Learning rate, set to 1×10^{-3} .
- `gamma`: Discount factor, set to 0.99.
- `clip_param`: Clipping parameter for PPO, set to 0.2.
- `c1`: Coefficient for value function loss, set to 0.5.
- `c2`: Coefficient for entropy bonus, set to 0.01.
- `epsilon`: Initial exploration rate, set to 1.0.
- `epsilon_decay`: Decay rate of epsilon, set to 0.995.
- `epsilon_min`: Minimum exploration rate, set to 0.01.

The agent selects actions using an ϵ -greedy strategy:

- With probability ϵ , select a random action.
- With probability $1 - \epsilon$, select the action with the highest policy probability:

$$a_t = \arg \max_a \pi(a|s_t)$$

- The exploration rate ϵ is decayed over time:

$$\epsilon = \max(\epsilon \times \text{epsilon_decay}, \text{epsilon_min})$$

The agent stores experiences in a memory buffer for training:

- Store experiences as tuples of `(state, action, reward, next_state, done)`.

The Generalized Advantage Estimation (GAE) is computed using the following formula:

$$\delta_t = r_t + \gamma V(s_{t+1}) - V(s_t)$$

$$\hat{A}_t = \sum_{l=0}^{\infty} (\gamma \lambda)^l \delta_{t+l}$$

and the policy and value are updated as

- **Policy Update (Actor):**

$$\text{ratio} = \frac{\pi(a_t|s_t)}{\pi_{\text{old}}(a_t|s_t)}$$

$$\mathcal{L}^{\text{CLIP}}(\theta) = \mathbb{E}_t \left[\min(\text{ratio} \cdot \hat{A}_t, \text{clip}(\text{ratio}, 1 - \epsilon, 1 + \epsilon) \cdot \hat{A}_t) \right]$$

- **Value Update (Critic):**

$$\mathcal{L}^{\text{VF}} = \left(\hat{R}_t - V(s_t) \right)^2$$

- **Total Loss:**

$$\mathcal{L} = \mathcal{L}^{\text{CLIP}} + c_1 \mathcal{L}^{\text{VF}} - c_2 \mathcal{H}$$

where \mathcal{H} is the entropy bonus.