

Learning to mitigate externalities: the coase theorem with hindsight rationality

Authors	Scheid, Antoine;Capitaine, Aymeric;Boursier, Etienne;Moulines, Eric;Jordan, Michael I.;Durmus, Alain
Citation	G.-Y. Hao, J. Zhang, B. Huang, H. Wang, and K. Zhang, "Natural Counterfactuals With Necessary Backtracking," Adv Neural Inf Process Syst, vol. 37, pp. 14962–14995, Dec. 2024.
DOI	10.52202/079017-0484
Publisher	NEURIPS
Download date	2026-05-15 06:55:27
Link to Item	https://hdl.handle.net/20.500.14634/278

Learning to Mitigate Externalities: the Coase Theorem with Hindsight Rationality

Antoine Scheid¹

Aymeric Capitaine¹

Etienne Boursier²

Eric Moulines¹

Michael I. Jordan^{3,4}

Alain Durmus¹

¹ Centre de Mathématiques Appliquées – CNRS – École polytechnique – Palaiseau, 91120, France

² INRIA Saclay, Université Paris Saclay, LMO - Orsay, 91400, France

³ University of California, Berkeley

⁴ Inria, Ecole Normale Supérieure, PSL Research University - Paris, 75, France

Abstract

In economic theory, the concept of externality refers to any indirect effect resulting from an interaction between players that affects the social welfare. Most of the models within which externality has been studied assume that agents have perfect knowledge of their environment and preferences. This is a major hindrance to the practical implementation of many proposed solutions. To address this issue, we consider a two-player bandit setting where the actions of one of the players affect the other player and we extend the Coase theorem [Coase, 2013]. This result shows that the optimal approach for maximizing the social welfare in the presence of externality is to establish property rights, i.e., enable transfers and bargaining between the players. Our work removes the classical assumption that bargainers possess perfect knowledge of the underlying game. We first demonstrate that in the absence of property rights, the social welfare breaks down. We then design a policy for the players which allows them to learn a bargaining strategy which maximizes the total welfare, recovering the Coase theorem under uncertainty.

1 Introduction

The concept of *externality* is used in economics to capture phenomena that impact the global welfare stemming from economic interactions without any compensation [Buchanan and Stubblebine, 2006, Shah et al., 2018]. Externality is generally considered as market failure since they result in a loss of collective welfare. Given its practical importance [Dahlman, 1979, Greenfield et al., 2009], mechanisms that characterize and mitigate externalities are central to modern economic thinking.

A first approach to tackle the adverse effects of externalities in modern economics was based on quotas and taxation [see, e.g., Pigou, 2017, and the references therein]. However, Coase’s theorem [Coase, 2013] shows that in the presence of well-defined property rights and low transaction costs, parties affected by externalities can privately negotiate efficient solutions, and recover a welfare efficient allocation through transfers and bargaining.

Throughout the paper, we use the following simple example to illustrate our results.

Example 1. Consider two firms 1 (upstream) and 2 (downstream) respectively producing quantities $q_1 \geq 0$ and $q_2 \geq 0$ of a good sold at a fixed price $p > 0$. They incur strictly increasing and convex costs, captured by the differentiable cost functions $c_1 : q_1 \mapsto c_1(q_1)$ and $c_2 : q_2 \mapsto c_2(q_2)$, satisfying

$c_1(0) = 0$ and $c_2(0) = 0$. We assume that firm 1 exerts on firm 2 a constant externality $\alpha > 0$ per unit produced. In other words, their profit functions (or utilities) are given by

$$\pi_1(q_1) = pq_1 - c_1(q_1) \quad \text{and} \quad \pi_2(q_1, q_2) = pq_2 - c_2(q_2) - \alpha q_1 .$$

One simple concrete illustration consists in an upstream firm emitting pollutants that reduce the downstream firm's production. If the upstream firm owns the property rights, it may receive a payment from the downstream one to reduce its output and thereby pollution. On the other hand, if the downstream firm owns the property right, it may require compensation from the upstream firm to allow its operation.

The Coase theorem demonstrates that in both cases with appropriate property rights, the resulting levels of production would be welfare efficient. The theorem is typically explained in textbooks under the assumption that the players have perfect knowledge of their own utility or profit function, as well as that of others. However, this assumption is unlikely to hold in real-world scenarios, where players have to learn about their own preferences and those of their competitors.

This example is, of course, a simplification of real world scenarios. For example, [Abildtrup et al. \[2012\]](#) consider a more complex setting to model the interaction between farmers and waterworks in Denmark where the farmers have the property rights whereas the waterworks can pay to reduce pollution. In particular, they demonstrate the failure of the theorem, attributing it to the breakdown of the main assumptions: no transaction costs, maximizing behaviors and perfect information, with an important focus on strategic behaviors and the asymmetry of information. We restore the latter in our work and provide foundations for the theorem to hold in more realistic scenarios.

A key question is whether the Coase theorem holds when players learn their preferences over time.

We investigate this question within the framework of a multi-armed bandit learning. We build upon recent works that extend the classical bandit setting to economics in which there are two players interacting via principal-agent protocols [[Dogan et al., 2023b,a](#), [Scheid et al., 2024](#)]. This allows us to capture, for example, a version of the two-firm problem where firms are uncertain regarding both their profit functions and the degree of externality on other firms. We represent production decisions in this problem as arms which can be played by the firms at any round over time, with the goal of finding decisions that maximize their rewards. More precisely, we assume that the reward of the upstream firm only depends on its own action, while the reward of the downstream firm depends on both its action and the upstream firm's action. This dependency on both actions allows us to capture externalities.

The property rights in [Example 1](#), or more generally over a bandit instance, amount to giving a firm the possibility of engaging in monetary transfers that influence the arms that are played over time. The owner of the bandit instance will face a problem of bandit learning with transfers with an upstream player who is also learning his preferences.

To account for the efficiency of a policy in this setup, we extend the classical static notion of welfare efficiency to the online setting. We say that a policy is *Welfare efficient* if the social welfare regret is sub-linear. Proving the Coase theorem within our setup therefore boils down to show that if the bandit owner (who is without loss of generality the upstream player in this study¹) runs a no-regret bandit algorithm to learn and exploit his preferences, the downstream player can then choose an optimal transfer scheme leading to a sub-linear total regret.

Our contributions are as follows:

- We show that when an upstream agent exerts externality on a downstream agent, in the absence of property right, the social welfare breaks down. Put differently, no joint policy of the agents can be welfare efficient.
- We then introduce property rights and show how it affects the game. In this case, bargaining and transfers are available to the players. We propose a policy for the downstream player that leads to welfare efficiency when the upstream player follows any black-box no-regret policy under mild assumption. This solution addresses the breakdown issue at equilibrium. Put together, we show an online version of the *Coase theorem*.

¹The fact that efficiency is restored whoever is given the property rights is known as the *invariance* property of the Coase theorem.

2 Setup and Inefficiency of Externality

2.1 Bandit game

We consider a sequential bandit game in which two players (downstream and upstream) simultaneously play actions in a bandit instance for a horizon $T \in \mathbb{N}^*$. The action set for both players is $\mathcal{A} = \{1, \dots, K\}$, $K \in \mathbb{N}^*$.

The reward distributions of the agents differ. Given a family of distributions $\{\gamma_a : a \in \mathcal{A}\}$ indexed by \mathcal{A} , the upstream player's rewards are provided by an i.i.d. family of random variables

$$\{(Z_a(t))_{t \in [T]} : a \in \mathcal{A}\}, \quad \text{where } Z_a(t) \sim \gamma_a \quad \text{for any } t \in [T] \text{ and } a \in \mathcal{A}.$$

To model the externality exerted by the upstream on the downstream player, we assume that the latter has a reward that depends both on her action and on that of the upstream player. Formally, this is modeled through a family of distributions $\{\nu_{a,b} : a, b \in \mathcal{A}\}$ double-indexed by \mathcal{A} and an i.i.d. family of random variables

$$\{(X_{a,b}(t))_{t \in [T]} : a, b \in \mathcal{A}\}, \quad \text{where } X_{a,b}(t) \sim \nu_{a,b}$$

is the reward received by the downstream player at time t if she pulls the arm b and the upstream player pulls the arm a .

Players. We assume that players are risk-neutral expected-utility maximizers, and we define their expected utilities for any $(a, b) \in \mathcal{A} \times \mathcal{A}$ as

$$v^{\text{up}}(a) = \int z \gamma_a(dz) \in \mathbb{R} \quad \text{and} \quad v^{\text{down}}(a, b) = \int x \nu_{a,b}(dx) \in \mathbb{R}.$$

The distributions $(\gamma_a)_{a \in \mathcal{A}}$ and $(\nu_{a,b})_{(a,b) \in \mathcal{A}^2}$ are unknown to both the downstream and the upstream players and they aim to learn the distributions with best mean rewards by sequentially observing samples from $\{(Z_a(t))_{t \in [T]} : a \in \mathcal{A}\}$ and $\{(X_{a,b}(t))_{t \in [T]} : a, b \in \mathcal{A}\}$.

Moreover, we suppose that players are rational in hindsight; that is, they minimize their regret. Formally, the upstream player aims to minimize his regret defined as

$$\mathfrak{R}_n^{\text{up}}(T, \Pi_n^{\text{up}}) = T \mu^{*,\text{up}} - \mathbb{E} \left[\sum_{t=1}^T v^{\text{up}}(A_t) \right], \quad \text{where } \mu^{*,\text{up}} = \max_{a \in \mathcal{A}} v^{\text{up}}(a), \quad (1)$$

while the downstream player seeks to minimize her external regret defined as

$$\mathfrak{R}_n^{\text{down}}(T, \Pi_n^{\text{up}}, \Pi_n^{\text{down}}) = \mathbb{E} \left[\sum_{t=1}^T \max_{b \in \mathcal{A}} v^{\text{down}}(A_t, b) - v^{\text{down}}(A_t, B_t) \right], \quad (2)$$

where the players' actions $(A_t)_{t \in [T]}$, $(B_t)_{t \in [T]}$ as well as their policies Π_n^{up} , Π_n^{down} are defined below. Note that the utility of the downstream player also depends on the actions taken by the upstream player, which represents the externality exerted by the upstream player on the downstream player, hence the strategic dimension of our setting. We first consider a game where no property right is defined, so each player is free to pick his preferred arm irrespectively of the other player's choice. This will result in a breakdown of the total utility.

Policies without property rights. Consider first the upstream player. Based on a policy Π_n^{up} (for example a no-regret bandit algorithm such as the Upper Confidence Bounds algorithm (UCB) [Auer, 2002] or the ε -greedy algorithm [Robbins, 1952, Langford and Zhang, 2007]), we define his history $(\mathcal{H}_t^{\text{up},n})_{t \in [T]}$ by induction. We set $\mathcal{H}_0^{\text{up},n} = \emptyset$ and supposing that $\mathcal{H}_t^{\text{up},n}$ is defined for $t \in [T]$, then

$$\mathcal{H}_{t+1}^{\text{up},n} = \mathcal{H}_t^{\text{up},n} \cup \{A_{t+1}, V_{t+1}, Z_{A_{t+1}}(t+1)\},$$

where $(V_s)_{s \in \mathbb{N}^*}$ is a family of independent uniform random variables in $[0, 1]$, allowing for randomization in the policy, and A_{t+1} is provided by Π_n^{up} , following $\Pi_n^{\text{up}} : (V_{t+1}, \mathcal{H}_t^{\text{up},n}) \mapsto A_{t+1}$.

Second, consider the downstream player and an algorithm Π_n^{down} (specifically a no-regret bandit algorithm). We define her history $(\mathcal{H}_t^{\text{down},n})_{t \in [T]}$ by induction. We set $\mathcal{H}_0^{\text{down},n} = \emptyset$ and supposing that $\mathcal{H}_t^{\text{down},n}$ is defined for $t \in [T]$, then

$$\mathcal{H}_{t+1}^{\text{down},n} = \mathcal{H}_t^{\text{down},n} \cup \{A_{t+1}, B_{t+1}, U_{t+1} X_{A_{t+1}, B_{t+1}}(t+1)\},$$

where $(U_s)_{s \in \mathbb{N}^*}$ is a family of independent uniform random variables in $[0, 1]$ allowing for randomization in the policy and B_{t+1} is provided by Π_n^{down} , following $\Pi_n^{\text{down}} : (U_{t+1}, \mathcal{H}_t^{\text{down}, n}) \mapsto B_{t+1}$.

Welfare efficiency. We now introduce the notion of *Welfare efficiency* for our setup. The global utility, or *social welfare*, of the players at round t is defined as $v^{\text{up}}(A_t) + v^{\text{down}}(A_t, B_t)$. We define the *socially optimal action* $(a^{\text{sw}}, b^{\text{sw}}) \in \mathcal{A} \times \mathcal{A}$ of the game as

$$(a^{\text{sw}}, b^{\text{sw}}) \in \operatorname{argmax}_{a, b \in \mathcal{A}} v^{\text{up}}(a) + v^{\text{down}}(a, b), \quad (3)$$

as well as the *global regret* (or *social welfare regret*) associated with policies Π_n^{up} and Π_n^{down} as

$$\mathfrak{R}^{\text{sw}}(T, \Pi_n^{\text{up}}, \Pi_n^{\text{down}}) = T (v^{\text{up}}(a^{\text{sw}}) + v^{\text{down}}(a^{\text{sw}}, b^{\text{sw}})) - \sum_{t=1}^T \mathbb{E} [v^{\text{up}}(A_t) + v^{\text{down}}(A_t, B_t)]. \quad (4)$$

Then, the joint policies Π_n^{up} and Π_n^{down} for the players are said to be *Welfare efficient* if

$$\lim_{T \rightarrow +\infty} \mathfrak{R}^{\text{sw}}(T, \Pi_n^{\text{up}}, \Pi_n^{\text{down}})/T = 0.$$

Intuitively, this condition implies that the frequency of the socially optimal action $(a^{\text{sw}}, b^{\text{sw}})$ tends to 1 as T goes to infinity. In this sense, it mimics the usual, static Welfare efficiency criterion. As we will see, $(\Pi_n^{\text{up}}, \Pi_n^{\text{down}})$ is typically not *Welfare efficient* when there is a disalignment in the game between the players' individual interests based on their rationality and the social welfare.

2.2 Inefficiency without property rights

We first present a result that captures the adverse consequence of externality on social welfare. The upstream player does not take into account the indirect cost incurred by the downstream player when he chooses his action. This drives the social welfare away from its optimal level. We illustrate this fact within our simple bilateral externality example.

Example 1 (continuing from p. 1). *We show that the competitive outcome, where each firm maximizes its profit independently, is not welfare efficient in the presence of externality. Define the social welfare as the function*

$$W : (q_1, q_2) \mapsto \pi_1(q_1) + \pi_2(q_1, q_2) = p(q_1 + q_2) - (c_1(q_1) + c_2(q_2)) - \alpha q_1. \quad (5)$$

By definition, the welfare efficient outcome $(q_1^, q_2^*) \in \mathbb{R}_+^2$ satisfies $W(q_1^*, q_2^*) \geq W(q_1, q_2)$ for any $(q_1, q_2) \in \mathbb{R}_+^2$. Since W is differentiable and strictly concave, (q_1^*, q_2^*) is uniquely defined by the condition $\nabla W(q_1^*, q_2^*) = 0$, that is*

$$c_1'(q_1) - \alpha = p \quad \text{and} \quad c_2'(q_2) = p. \quad (6)$$

Note that at the welfare efficient optimum, firm 1 does not equalize marginal cost with marginal profit, but produces less to account for the negative effect of externality on firm 2. We now characterize the competitive outcome $(q_1', q_2') \in \mathbb{R}_+^2$. Since π_1 and π_2 are differentiable and strictly concave, (q_1', q_2') satisfies

$$c_1'(q_1') = p \quad \text{and} \quad c_2'(q_2') = p. \quad (7)$$

For the competitive outcome to be welfare efficient, we require, by Equation (6) and Equation (7),

$$c_1'(q_1') - \alpha = c_1'(q_1'), \quad \text{that is} \quad \alpha = 0.$$

This proves that whenever there are externalities, no competitive outcome is efficient.

We now show that in our model, when there is no property right and under mild assumptions, no achievable policy is welfare efficient whenever there is a misalignment between the players' interests and the social welfare. The upstream player's policy Π_n^{up} is said to be *no-regret* if $\lim_{T \rightarrow +\infty} \mathfrak{R}_n^{\text{up}}(T, \Pi_n^{\text{up}})/T = 0$, where $\mathfrak{R}_n^{\text{up}}$ is defined in (1).

Theorem 2. *Suppose that $\operatorname{argmax}_{a \in \mathcal{A}} v^{\text{up}}(a)$ is the singleton $\{a_*^{\text{u}}\}$ and that*

$$v^{\text{up}}(a^{\text{sw}}) + v^{\text{down}}(a^{\text{sw}}, b^{\text{sw}}) - v^{\text{up}}(a_*^{\text{u}}) + v^{\text{down}}(a_*^{\text{u}}, b) > 0, \quad (8)$$

for any $b \in \mathcal{A}$. In the absence of property rights and when the upstream player runs any no-regret policy Π_n^{up} , we have $\mathfrak{R}^{\text{sw}}(T, \Pi_n^{\text{up}}, \Pi_n^{\text{down}}) = \Omega(T)$. Therefore, $\mathfrak{R}^{\text{sw}}(T, \Pi_n^{\text{up}}, \Pi_n^{\text{down}}) = \Omega(T)$ and $(\Pi_n^{\text{up}}, \Pi_n^{\text{down}})$ is not welfare efficient.

Condition (8) in Theorem 2 represents the unalignment between the upstream player's preference and the optimal choice from a social welfare point of view. Note that the upstream and downstream players can both have an $o(T)$ external regret, while the social welfare regret still grows linearly with T because of the unfavorable interactions between their policies.

3 Online Property Game with Bargaining Players

3.1 Online Property Game

We now consider the same repeated game in the form of a *property game* where one of the players possesses the bandit instance (*upstream player*). As in the original setup of Coase [2013], the other player (*downstream player*) will provide the bandit owner with transfers to incentivize him to choose some specific action and influence the outcome of the game in her favor.

We show in Appendix B that our method applies similarly when property rights are given to the upstream player rather than the downstream player. Hence, there is no loss of generality in considering the aforementioned framework. In this sense, we recover the *invariance property* of the Coasean bargaining [Mas-Colell et al., 1995].

Example 1 (continuing from p. 1). *We now illustrate how Coasean bargaining re-instates efficiency. Suppose without loss of generality that property rights are such that firm 2 can pay $\tau \in \mathbb{R}_+$ to firm 1 for it to operate at a level \tilde{q}_1 . Profits become*

$$\bar{\pi}_2 : (q_1, q_2, \tau, \tilde{q}_1) \mapsto \pi_2(q_1, q_2) - \mathbb{1}_{\{q_1 = \tilde{q}_1\}} \tau \quad \text{and} \quad \bar{\pi}_1 : (q_1, \tau, \tilde{q}_1) \mapsto \pi(q_1) + \mathbb{1}_{\{q_1 = \tilde{q}_1\}} \tau.$$

Consider the competitive outcome $(q_1, q_2, \tau, \tilde{q}_1) \in \mathbb{R}_+^4$ which satisfies

$$q_1 = q_1(\tau, \tilde{q}_1) \in \arg \max_{q'_1 \geq 0} \bar{\pi}_1(q'_1, \tau, \tilde{q}_1) \quad \text{and} \\ \bar{\pi}_2(q_1(\tau, \tilde{q}_1), q_2, \tau, \tilde{q}_1) = \max_{q'_2, \tau', \tilde{q}'_1} \bar{\pi}_2(q_1(\tau', \tilde{q}'_1), q'_2, \tau', \tilde{q}'_1).$$

The condition on q_1 accounts for the rationality of the firm 1 and the fact that its choice depends on the payment (τ, \tilde{q}_1) . Obviously, the optimal solution is reached for $\tilde{q}_1 = q_1$ and $\tau = \max_{q'} \pi_1(q') - \pi_1(\tilde{q}_1)$. Plugging this back in the expression of $\bar{\pi}_2$ then yields

$$(q_1, q_2) = \operatorname{argmax}_{q'_1, q'_2} \pi_1(q'_1) + \pi_2(q'_2) = \operatorname{argmax}_{q'_1, q'_2} W(q'_1, q'_2),$$

so the competitive outcome (q_1, q_2) is welfare efficient.

The transfers at each step can be interpreted as a contract between two players [see, e.g., Bolton and Dewatripont, 2004, Salanié, 2005, for general contract theory] and providing the right amount of incentives relates to adjusting a contract in an online setting [see Dütting et al., 2019, Guruganesh et al., 2021, Zhu et al., 2022, Fallah and Jordan, 2023, Guruganesh et al., 2024, Ananthkrishnan et al., 2024, for learning-based perspectives about contracts].

Similarly to Example 1, we modify the players' policies to now account for the transfer $\tau(t)$ that the downstream player offers at round t to the upstream player if he picks action \tilde{a}_t . The downstream player's policy at round t does not only output an arm B_t but now a triple $(\tilde{a}_t, \tau(t), B_t)$, where B_t is the arm that she should play and \tilde{a}_t is the arm on which a transfer $\tau(t)$ is offered to the upstream player. On the upstream player's side, the policy still outputs an arm A_t to play but also takes as an input the incentive $(\tilde{a}_t, \tau(t))$. In addition, the instantaneous utility of the upstream player becomes $Z_{A_t}(t) + \mathbb{1}_{\tilde{a}_t}(A_t)\tau(t)$, whereas the downstream player receives $X_{A_t, B_t}(t) - \mathbb{1}_{\tilde{a}_t}(A_t)\tau(t)$.

Policies with property rights. Based on policies Π_p^{up} for the upstream player and Π_p^{down} for the downstream player, we define their histories $(\mathcal{H}_t^{\text{up}, \text{p}})_{t \in [T]}$ and $(\mathcal{H}_t^{\text{down}, \text{p}})_{t \in [T]}$ by induction. We set $\mathcal{H}_0^{\text{up}, \text{p}} = \emptyset$, $\mathcal{H}_0^{\text{down}, \text{p}} = \emptyset$ and supposing that $\mathcal{H}_t^{\text{up}, \text{p}}$, $\mathcal{H}_t^{\text{down}, \text{p}}$ are defined for $t \in [T]$, then

$$\mathcal{H}_{t+1}^{\text{up}, \text{p}} = \mathcal{H}_t^{\text{up}, \text{p}} \cup \{\tilde{a}_{t+1}, \tau(t+1), A_{t+1}, V_{t+1}, Z_{A_{t+1}}(t+1)\}$$

and

$$\mathcal{H}_{t+1}^{\text{down}, \text{p}} = \mathcal{H}_t^{\text{down}, \text{p}} \cup \{\tilde{a}_{t+1}, \tau(t+1), A_{t+1}, B_{t+1}, U_{t+1}, X_{A_{t+1}, B_{t+1}}(t+1)\},$$

where $(V_s)_{s \in \mathbb{N}^*}$, $(U_s)_{s \in \mathbb{N}^*}$ are two families of independent uniform random variables in $[0, 1]$ allowing for randomization in the policies, and the remaining quantities are given by $\Pi_p^{\text{down}}: (U_{t+1}, \mathcal{H}_t^{\text{down,p}}) \mapsto (\tilde{a}_{t+1}, \tau(t+1), B_{t+1})$ and $\Pi_p^{\text{up}}: (\tilde{a}_{t+1}, \tau(t+1), V_{t+1}, \mathcal{H}_t^{\text{up,p}}) \mapsto A_{t+1}$.

Players' goal. Given a transfer τ from the downstream to the upstream player on arm \tilde{a} , actions a and b respectively are chosen by the upstream and the downstream player, the upstream player's expected utility reads $v^{\text{up}}(a) + \mathbb{1}_{\tilde{a}}(a)\tau$ while the downstream player's expected utility is $v^{\text{down}}(a, b) - \mathbb{1}_{\tilde{a}}(a)\tau$. This defines the upstream player's expected regret for a horizon T as

$$\mathfrak{R}_p^{\text{up}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) = \mathbb{E} \left[\sum_{t=1}^T \max_{a \in \mathcal{A}} \{v^{\text{up}}(a) + \mathbb{1}_{\tilde{a}_t}(a)\tau(t)\} - (v^{\text{up}}(A_t) + \mathbb{1}_{\tilde{a}_t}(A_t)\tau(t)) \right]. \quad (9)$$

Based on the upstream player's utility, the downstream player aims on a single round at proposing an optimal transfer τ^{opt} on an arm $a^{\text{opt}} \in \mathcal{A}$ as well as picking an arm $b^{\text{opt}} \in \mathcal{A}$ which solves

$$\begin{aligned} & \text{maximize } (a, b, \tau) \mapsto v^{\text{down}}(a, b) - \tau \\ & \text{such that } \tau \in \mathbb{R}_+, b \in \mathcal{A}, a \in \operatorname{argmax}_{a' \in \mathcal{A}} \{v^{\text{up}}(a') + \mathbb{1}_a(a')\tau\}. \end{aligned} \quad (10)$$

Her regret for any horizon T is defined as

$$\mathfrak{R}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) = T\mu^{*,\text{down}} - \mathbb{E} \left[\sum_{t=1}^T v^{\text{down}}(A_t, B_t) - \mathbb{1}_{\tilde{a}_t}(A_t)\tau(t) \right], \quad (11)$$

where we define $\mu^{*,\text{down}} = v^{\text{down}}(a^{\text{opt}}, b^{\text{opt}}) - \tau^{\text{opt}}$ as the optimal utility she can aim for. We can see that the downstream player's influence is exerted through her action choice B_t as well as through transfers which enable her to influence the upstream player's actions. Hence, the notion of external regret is obsolete here. The game has now the form of a repeated *Stackelberg game* [Von Stackelberg, 2010].

Lemma 1. *Recall that $\mu^{*,\text{down}}$ is the downstream player's optimal reward as defined as a solution of (10). We have $\mu^{*,\text{down}} = \max_{a,b \in \mathcal{A}} \{v^{\text{down}}(a,b) + v^{\text{up}}(a)\} - \max_{a' \in \mathcal{A}} \{v^{\text{up}}(a')\}$, as well as $(a^{\text{opt}}, b^{\text{opt}}) = (a^{\text{sw}}, b^{\text{sw}})$ and $\mu^{*,\text{up}} + \mu^{*,\text{down}} = v^{\text{up}}(a^{\text{sw}}) + v^{\text{down}}(a^{\text{sw}}, b^{\text{sw}}) = \max_{a,b \in \mathcal{A}} \{v^{\text{up}}(a) + v^{\text{down}}(a,b)\}$, where $\mu^{*,\text{up}}$ is defined in Equation (1). Moreover, for any integer $T \in \mathbb{N}^*$, and policies $\Pi_p^{\text{up}}, \Pi_p^{\text{down}}$, we have that*

$$\mathfrak{R}^{\text{sw}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq \mathfrak{R}_p^{\text{up}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) + \mathfrak{R}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}).$$

This lemma has an interesting economic interpretation: if both players individually seek for their own interest within this online property game, they will together converge towards the optimal global utility. Individual rationality moves the outcome of the game towards the optimal social welfare. The transfers allow the players to align their goals and share the global reward, in line with the Coase theorem. Consequently, if both players run no-regret policies Π_p^{up} and Π_p^{down} , the social welfare regret will also be in $o(T)$. The rest of the paper shows that such no-regret policies exist. To this end, we introduce the following assumptions.

Without loss of generality, we assume that the upstream player's utility is rescaled and shifted, which corresponds to the following assumption on the reward distribution $(\gamma_a)_{a \in \mathcal{A}}$ in \mathbb{R} .

H1. *For any $a \in \mathcal{A}$, we have $v^{\text{up}}(a) \in [0, 1]$.*

We now make a high probability bound assumption on the upstream player's regret.²

H2. *There exist $C, \zeta > 0, \kappa \in [0, 1]$ such that for any $s, t \in [T]$ with $s+t \leq T$, any $\{\tau_a\}_{a \in [K]} \in \mathbb{R}_+^K$ and any policy Π_p^{down} that offers almost surely a transfer $(\tilde{a}_l, \tau(l)) = (\tilde{a}_l, \tau_{\tilde{a}_l})$ for any $l \in \{s+1, \dots, s+t\}$, the batched regret of the upstream player following Π_p^{up} satisfies, with probability at least $1 - t^{-\zeta}$,*

$$\sum_{l=s+1}^{s+t} \max_{a \in \mathcal{A}} \{v^{\text{up}}(a) + \mathbb{1}_{\tilde{a}_l}(a)\tau_{\tilde{a}_l}\} - (v^{\text{up}}(A_l) + \mathbb{1}_{\tilde{a}_l}(A_l)\tau_{\tilde{a}_l}) \leq Ct^\kappa.$$

²A similar assumption is made in the work of Donahue et al. [2024] but with a stronger instantaneous regret bound which does not encompass the UCB's regret bound.

The constraint on the downstream player's algorithm Π_p^{down} enforces constant incentives associated with any arm $a \in \mathcal{A}$ within the batch, while the incentivized actions $(\tilde{a}_l)_{l \in \{s+1, \dots, s+t\}}$ may change. Proposition 2 in Appendix C shows that an adaptation of UCB taking account the incentives satisfies **H2** with $C = 8\sqrt{K \log(KT^3)}$, $\kappa = 1/2$ and $\zeta = 2$. Note that usual bandit algorithms such as AAE, ETC or EXP-IX also satisfy the assumption [see, e.g., Donahue et al., 2024, Lattimore and Szepesvári, 2020].

3.2 Downstream player's procedure

We fix the policy Π_p^{up} which can be any algorithm satisfying **H2** for the upstream player and introduce the algorithm BELGIC (Bandits and Externalities for a Learning Game with Incentivized Coase) which provides a policy achieving sub-linear regret for the downstream player. It can be seen as an online bargaining strategy to mitigate externalities. Simply put, BELGIC unfolds in two steps. First note that for any action $a \in \mathcal{A}$, the optimal (lowest) transfer to offer to the upstream player to make him choose a is

$$\tau_a^* = \max_{a' \in \mathcal{A}} v^{\text{up}}(a') - v^{\text{up}}(a), \quad (12)$$

as detailed in Appendix A. Therefore, a batched binary search procedure (Algorithm 2) first allows the downstream player to estimate the optimal transfers $\tau_1^*, \dots, \tau_K^*$ with a good precision level of $1/T^\beta$, where $\beta > 0$. More precisely, the downstream player offers a constant incentive $(\tilde{a}, \tau_{\tilde{a}})$ for a batch of time steps of length $\tilde{T} = \lceil T^\alpha \rceil$. The observation of $T_{\tilde{a}}^\neq$, the number of steps from the batch for which the upstream player does not pick \tilde{a} allows her to estimate whether $\tau_{\tilde{a}}$ is above or below $\tau_{\tilde{a}}^*$ and adjust it, following Lemma 2 in Appendix A under the condition that α, β satisfy

$$\beta/\alpha < (1 - \kappa). \quad (13)$$

The procedure needs to be run for $K \lceil T^\alpha \rceil \lceil \log_2 T^\beta \rceil$ rounds since we have to make $\lceil \log_2 T^\beta \rceil$ batches of binary search of length $\lceil T^\alpha \rceil$ on each of the K arms [see Scheid et al., 2024]. This corresponds to the first phase of BELGIC as described in Algorithm 2. At the end of this stage, the estimated transfers $(\hat{\tau}_a)_{a \in \mathcal{A}}$ satisfy the bound in Proposition 1. These are then used to feed the subroutine Bandit-Alg.

Proposition 1. *Under **H1** and **H2**, after the first phase of BELGIC which consists in $K \lceil T^\alpha \rceil \lceil \log T^\beta \rceil$ steps of binary search grouped in $\lceil \log_2 T^\beta \rceil$ batches per arm $a \in \mathcal{A}$, we have that*

$$\mathbb{P}\left(\text{for any } a \in \mathcal{A}, \hat{\tau}_a - 4/T^\beta - CT^{(\kappa-1)/2} \leq \tau_a^* \leq \hat{\tau}_a\right) \geq 1 - K \lceil \log_2 T^\beta \rceil / T^{\alpha\zeta}.$$

The additional term $1/T^\beta + CT^{\kappa-1}$ in $\hat{\tau}_a$ ensures that if **H2** holds, the upstream player necessarily plays the incentivized action \tilde{a}_t at round t with high probability. Lemmas 2 and 6 from Appendix C show how the binary search batches in BELGIC allow us to estimate τ_a^* depending on $\tilde{T} - T_{\tilde{a}}^\neq$, the number of times that arm a has been pulled by the upstream player during the batch.

Then, any bandit subroutine Bandit-Alg, such as UCB or ε -greedy, for instance, can be run in a black-box fashion on the shifted bandit instance, where the rewards are shifted by the upper estimated transfers $(\hat{\tau}_a)_{a \in \mathcal{A}}$. The downstream player computes a shifted history $\tilde{\mathcal{H}}_t^{\text{down},p}$ such that for any $t \leq K \lceil T^\alpha \rceil \lceil \log_2 T^\beta \rceil$, $\tilde{\mathcal{H}}_t^{\text{down},p} = \emptyset$ and for any $t > K \lceil T^\alpha \rceil \lceil \log_2 T^\beta \rceil$

$$\tilde{\mathcal{H}}_t^{\text{down},p} = \begin{cases} \{\tilde{a}_t, B_t, \tau(t), U_t, X_{\tilde{a}_t, B_t}(t) - \hat{\tau}_{\tilde{a}_t}\} \cup \tilde{\mathcal{H}}_{t-1}^{\text{down},p} & \text{if } \tilde{a}_t = A_t \\ \tilde{\mathcal{H}}_{t-1}^{\text{down},p} & \text{otherwise,} \end{cases} \quad (14)$$

which serves to feed Bandit-Alg, following

$$\text{Bandit-Alg: } (U_t, \tilde{\mathcal{H}}_{t-1}^{\text{down},p}) \mapsto (\tilde{a}_t, B_t) \in \mathcal{A} \times \mathcal{A}. \quad (15)$$

For any family of constant incentives $\{\tau_a\}_{a \in \mathcal{A}} \in \mathbb{R}_+^K$, we define $\mathfrak{R}_{\text{Bandit-Alg}}(T, \nu, \{\tau_a\}_{a \in \mathcal{A}})$ as the regret for the downstream player's subroutine Bandit-Alg on the bandit instance with shifted means over T rounds, following

$$\mathfrak{R}_{\text{Bandit-Alg}}(T, \nu, \{\tau_a\}_{a \in \mathcal{A}}) = T \max_{a, b \in \mathcal{A}^2} \mathbb{E}[v_{a,b}^{\text{down}}(1) - \tau_a] - \mathbb{E}\left[\sum_{t=1}^T v^{\text{down}}(\tilde{a}_t, B_t) - \tau_{\tilde{a}_t}\right].$$

Note that here, Bandit-Alg aims to maximize the shifted reward $(v^{\text{down}}(a, b) - \tau_a)_{(a,b) \in \mathcal{A}^2}$.

Algorithm 1 BELGIC

1: **Input:** Set of actions $\mathcal{A} = [K]$, time horizon T , subroutine Π_p^{up} , upstream player's regret constants C, κ , parameters α and β .
2: Compute $\tilde{\mathcal{H}}_s^{\text{down,p}} = \emptyset$ for any $s \leq K \lceil \log_2 T^\beta \rceil \lceil T^\alpha \rceil$.
3: **for** $a \in \mathcal{A}$ **do**
4: # See Algorithm 2
5: $\underline{\tau}_a, \bar{\tau}_a = \text{Binary Search}(a, \lceil \log_2 T^\beta \rceil, \lceil T^\alpha \rceil, 0, 1)$
6: **end for**
7: For any action $a \in \mathcal{A}$, $\hat{\tau}_a = \bar{\tau}_a + 1/T^\beta + CT^{(\kappa-1)/2}$.
8: **for** $t = K \lceil T^\alpha \rceil \lceil \log_2 T^\beta \rceil + 1, \dots, T$ **do**
9: Get recommended actions by Bandit-Alg on the $\mathcal{A} \times \mathcal{A}$ bandit instance, $(\tilde{a}_t, B_t) = \text{Bandit-Alg}(U_t, \tilde{\mathcal{H}}_{t-1}^{\text{down,p}})$.
10: Offer a transfer $\hat{\tau}_{\tilde{a}_t}$ on action \tilde{a}_t , nothing for any other action $a' \in \mathcal{A}$ and play action B_t .
11: Observe $A_t = \Pi_p^{\text{up}}(\tilde{a}_{t+1}, \tau(t+1), V_t, \mathcal{H}_{t-1}^{\text{up,p}}), X_{\tilde{a}_t, B_t}(t)$
12: **if** $A_t = \tilde{a}_t$ **then** update history $\tilde{\mathcal{H}}_t^{\text{down,p}}$.
13: **end if**
14: Update upstream player's history $\mathcal{H}_t^{\text{up,p}}$.
15: **end for**

Algorithm 2 Binary Search Subroutine

1: **Input:** action $a, N_T, \tilde{T}, \underline{\tau}_a, \bar{\tau}_a$.
2: **for** $d = 0, \dots, N_T - 1$ **do**
3: Compute $\tau_a^{\text{mid}} = (\bar{\tau}_a(d) + \underline{\tau}_a(d))/2, T_a^\neq = 0$.
4: **for** $t = d\tilde{T} + 1, \dots, d\tilde{T} + \tilde{T}$ **do**
5: Propose transfer $\tau_a^{\text{mid}}(d)$ on arm a and nothing for any other action $a' \in \mathcal{A}$.
6: $A_t = \Pi_p^{\text{up}}(t, \tau_a^{\text{mid}}(a), a, V_t, \mathcal{H}_{t-1}^{\text{up,p}})$
7: **if** $A_t \neq a$ **then:** $T_a^\neq + = 1$
8: **end if**
9: Update upstream player's history $\mathcal{H}_t^{\text{up,p}}$.
10: **end for**
11: **if** $CT^{\kappa+\beta/\alpha} < T_a^\neq < \tilde{T} - CT^{\kappa+\beta/\alpha}$ **then** Return $\underline{\tau}_a(d), \bar{\tau}_a(d)$.
12: **else if** $T_a^\neq \leq \tilde{T} - CT^{\kappa+\beta/\alpha}$ **then** $\bar{\tau}_a(d) = \tau_a^{\text{mid}}(d) + 1/T^\beta$ and update history $\tilde{\mathcal{H}}_t^{\text{down,p}}$.
13: **else** $\underline{\tau}_a(d) = \tau_a^{\text{mid}}(d) - 1/T^\beta$ and update history $\tilde{\mathcal{H}}_t^{\text{down,p}}$.
14: **end if**
15: **end for**

Theorem 3. Assume that **H1** and **H2** hold. Then BELGIC, run with α, β satisfying (13) and any bandit subroutine *Bandit-Alg*, has an overall regret $\mathfrak{R}_p^{\text{down}}$ such that

$$\mathfrak{R}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \text{BELGIC}) \leq 2(3 + 2C + \bar{v} - \underline{v}) \log_2(T) (2T^{1-\alpha\zeta} + T^{(\kappa+1)/2} + \lceil T^\alpha \rceil) + 4T^{1-\beta} + \mathfrak{R}_{\text{Bandit-Alg}}(T, \nu, \{\hat{\tau}_a\}_{a \in \mathcal{A}})$$

where, for ease of notation

$$\bar{v} = \max_{a, b \in \mathcal{A} \times \mathcal{A}} \{v^{\text{down}}(a, b)\} \quad \text{and} \quad \underline{v} = \min_{a, b \in \mathcal{A} \times \mathcal{A}} \{v^{\text{down}}(a, b)\}.$$

Knowledge of C and κ . An upper bound on C and κ is sufficient to compute the hyperparameters in BELGIC. Theorem 3 shows that the bigger C and κ are, the worse is the downstream player's regret, hence the interest of knowing them more precisely.

Corollary 1. Assume that the upstream player's distribution $(\gamma_a)_{a \in \mathcal{A}}$ is such that **H1** holds. In addition, suppose that the distributions $(\gamma_a)_{a \in \mathcal{A}}$ and $(\nu_{a,b})_{a, b \in \mathcal{A} \times \mathcal{A}}$ are 1-sub-Gaussian and that the upstream player plays $\Pi_p^{\text{up}} = \text{Algorithm 3}$ (a slight modification of UCB to take into account the incentives). Then the downstream player's regret when she runs BELGIC with parameters $\alpha = 3/4$ and $\beta = 1/4$ (which satisfy (13)) and subroutine *Bandit-Alg* = UCB satisfies the following upper

bound³

$$\mathfrak{R}_p^{\text{down}}(T, \text{UCB}, \text{BELGIC}) \leq (10 + 4K + 32\sqrt{K \log_2(KT^3)} + \bar{v} - \underline{v}) \log_2(T)(3 + 2T^{3/4}) + 3K^2(\bar{v} - \underline{v}).$$

The upper bound on the social welfare regret in Lemma 1 together with Corollary 1 shows that when the upstream player runs $\Pi_p^{\text{up}} = \text{UCB}$ and the downstream player runs BELGIC, the social welfare regret then satisfies $\mathfrak{R}^{\text{sw}}(T, \Pi_p^{\text{up}}, \text{BELGIC}) = \mathcal{O}(K \log(T)T^{(\kappa+1)/2})$.

In other words, if the downstream player runs BELGIC which produces a policy $\Pi_{\text{down}}^{\text{n}}$, for any upstream policy Π_p^{up} , $(\Pi_{\text{down}}^{\text{n}}, \Pi_p^{\text{up}})$ is welfare efficient.

Influence of the upstream performance. It is interesting to note that in the downstream player’s regret bound, the upstream player’s regret bound in $\mathcal{O}(T^\kappa)$ plays a significant role: the downstream player never learns faster than the upstream player. The latter’s performance determines the social welfare convergence rate towards the social optimum. We can observe that the players’ bounded rationality [Selten, 1990, Jones, 1999] and personal interest make the game converge towards the optimal social welfare equilibrium—even though they are both learning here.




Figure 1: Empirical frequencies of the upstream player’s actions when property rights are not defined (left) and when they are defined (right).

Experiments. We conclude this section with experiments showing the empirical convergence of our algorithm to a social optimum. In the simulation, we consider two firms, with firm 1 being upstream and firm 2 being downstream. Their profit functions are respectively given by

$$\pi_1 := \max \left\{ q_1 - 2 \left(\frac{q_1}{10} \right)^2 + 2 \frac{q_1}{10}, 0 \right\},$$

and

$$\pi_2(q_1, q_2) := \max \left\{ -16 \left(\frac{q_1}{10} - \frac{6}{10} \right)^2 + 8 \left(q_1 - \frac{6}{10} \right)^2 + \frac{1}{50} q_2, 0 \right\}.$$

Thus, firm 1’s and firm 2’s profit functions depends quadratically on q_1 with an firm 1 optimum at $q_1^* = 5$ and a social optimum at $q_1^* = 8$. Note that in the expression of π_2 , q_2 has very little influence as compared to q_1 - which allows to plot profits for only one value of q_2 .

We discretize the setup, consider a bandit instance (horizon $T = 5 \cdot 10^6$, 10 arms, average over 10 rounds) and we assume that UCB is used as a subroutine. In the first setting, there are no property rights and each firm runs UCB on their side. Second, property rights are defined and firm 2 runs BELGIC as its policy. The plots in Figure 1 display the empirical frequencies and show empirically the effectiveness of BELGIC to mitigate externalities.

4 Related work

Our work addresses the impact of externalities and is therefore related to taxation theory [see, e.g., Mirrlees et al., 2011, Salanie, 2011, and the references therein], a prominent solution for this issue, as

³Note that it is K and not \sqrt{K} here, since the action space is of cardinality K^2 for the downstream player.

exemplified by the *Pigouvian tax* [see [Pigou, 2017](#)]. Taxation is a fundamental aspect of all developed economies, with 30% to 50% of national income derived from taxes. The topic has been fruitful for various scenarios, including the carbon tax [[Carattini et al., 2018](#), [Metcalf and Weisbach, 2009](#)], alcohol markets [[Griffith et al., 2019](#)], or business taxation [[Boadway and Bruce, 1984](#)]. Taxation can also be studied through an operations research lens, where it is used to enhance system efficiency or manage specific games [[Roughgarden, 2010](#), [Caragiannis et al., 2010](#), [Bild and Vinci, 2019](#)]. Recent work by [Cui et al. \[2024\]](#) explores online mechanisms to maximize efficiency in congestion games.

Mechanism designs [[Myerson, 1989](#), [Nisan and Ronen, 1999](#), [Laffont and Martimort, 2009](#)] allow to design games that have specific desired outcomes. Deploying these mechanisms in their classical economical form assume that players’ utility functions are known a priori, which is often unrealistic. There is a major need to blend mechanism design with machine learning.

However, our approach differs, since, drawing inspiration from Coase’s theory, we implement an online version of his theorem [[Coase, 2013](#), [Cooter, 1982](#)], incorporating uncertainty to tackle the breakdown of social welfare in an online setting. We use the bandit setup [see [Lattimore and Szepesvári, 2020](#), [Slivkins et al., 2019](#)] as a general and convenient way to model the game introduced by Coase. However, our work differs from considering a single agent playing a bandit game. Instead, we focus on the more general problem of multi-players bandits, a field receiving a growing attention from the community [see e.g., [Boursier and Perchet, 2019, 2022](#), [Sankararaman et al., 2019](#)].

Our approach is inspired by the principal-agent model introduced by [Dogan et al. \[2023b\]](#), which was further extended by [Dogan et al. \[2023a\]](#), [Scheid et al. \[2024\]](#). However, unlike the models proposed in the work of [Dogan et al. \[2023b,a\]](#), we do not specify a particular bandit algorithm for the upstream player and instead, we allow him to use any no-regret algorithm satisfying $\mathbf{H2}$ in a black-box fashion. Conversely, the model of [Scheid et al. \[2024\]](#) assumes that the upstream player is always fully informed and best-responding, whereas we assume that he is also learning. [Chen et al. \[2023b\]](#) leverages a similar model to study information acquisition by a principal through an agent’s actions. However, in their model, the agent is also almighty and knows exactly the costs associated with each action. Designing incentives in an unknown environment is related to auction theory incorporating uncertainty, as it is explored in the work of [Feng et al. \[see 2018\]](#), [Li et al. \[see 2023\]](#). Similar issues have been explored in the *Reinforcement Learning* framework within a leader-follower game [see [Chen et al., 2023a](#), with quantal responses by the follower] or in a principal-agent game with incentive design as done by [Ben-Porat et al. \[2023\]](#). [Donahue et al. \[2024\]](#) also study a two-players repeated Stackelberg game on a bandit instance but instead of allowing for transfers, their main focus concerns the achievability of a *Stackelberg equilibrium* through iterations of bandit policies: the same kind of goal also appears in [Collina et al. \[2023\]](#). Such principal-agent setups are of some interest to model various real-world situations such as the design of fundings for hospitals [[Wang et al., 2024](#)] or have been studied with multiple agents through the lens of auction design in dynamic setups [[Bergemann and Said, 2010](#), [Chen et al., 2023c](#)], or to account for fairness [[Fallah et al., 2024](#)].

In our game, the downstream player needs to learn the optimal transfers/incentives to offer to the upstream player. This is related to the *Incentivized Exploration* literature [[Mansour et al., 2016](#), [Simchowitz and Slivkins, 2023](#), [Esmaeili et al., 2023](#)], which is often cast in terms of a benevolent planner who aims to optimize the global welfare of agents via plausible recommendations. A related model is *Bayesian Persuasion* [[Kamenica and Gentzkow, 2011](#)], where a sender influences a receiver’s action through sending a signal. This model has begun to be studied in learning settings [see, e.g., [Castiglioni et al., 2020](#), [Bernasconi et al., 2022](#), [Wu et al., 2022b,a](#)].

5 Conclusion

This paper studies a model of externalities in a two-players sequential game where both players learn their optimal actions. We first show that when the players act independently, then a misalignment between the players’ interests and the social welfare leads to a breakdown of the global utility. We then introduce interactions through transfers, which restores a social welfare optimum, representing the online version of the *Coase theorem*. To that purpose, we propose a policy for the downstream player which allows her to estimate the optimal transfers as well as choosing the best actions. The mathematical difficulty comes from the learning aspect on both sides. Since our work is coined in a learning framework for mechanism design, several directions for research are open, as for instance extensions to the multi-agent setting, which raises many questions.

Acknowledgements

Funded by the European Union (ERC, Ocean, 101071601). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- Jens Abildtrup, Frank Jensen, and Alex Dubgaard. Does the coase theorem hold in real markets? an application to the negotiations between waterworks and farmers in denmark. *Journal of environmental management*, 93(1):169–176, 2012.
- Nivasini Ananthakrishnan, Stephen Bates, Michael Jordan, and Nika Haghtalab. Delegating data collection in decentralized machine learning. In *International Conference on Artificial Intelligence and Statistics*, pages 478–486. PMLR, 2024.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- Omer Ben-Porat, Yishay Mansour, Michal Moshkovitz, and Boaz Taitler. Principal-agent reward shaping in mdps. *arXiv preprint arXiv:2401.00298*, 2023.
- Dirk Bergemann and Maher Said. Dynamic auctions: A survey. *Wiley Encyclopedia of Operations Research and Management Science*, 2010.
- Martino Bernasconi, Matteo Castiglioni, Alberto Marchesi, Nicola Gatti, and Francesco Trovò. Sequential information design: Learning to persuade in the dark. *Advances in Neural Information Processing Systems*, 35:15917–15928, 2022.
- Vittorio Bilò and Cosimo Vinci. Dynamic taxes for polynomial congestion games. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–36, 2019.
- Robin Boadway and Neil Bruce. A general proposition on the design of a neutral business tax. *Journal of Public Economics*, 24(2):231–239, 1984.
- Patrick Bolton and Mathias Dewatripont. *Contract theory*. MIT press, 2004.
- Etienne Boursier and Vianney Perchet. Sic-mmab: Synchronisation involves communication in multiplayer multi-armed bandits. *Advances in Neural Information Processing Systems*, 32, 2019.
- Etienne Boursier and Vianney Perchet. A survey on multi-player bandits. *arXiv preprint arXiv:2211.16275*, 2022.
- Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.
- James M Buchanan and Wm Craig Stubblebine. Externality. In *Inframarginal Contributions to Development Economics*, pages 55–73. World Scientific, 2006.
- Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. Taxes for linear atomic congestion games. *ACM Transactions on Algorithms (TALG)*, 7(1):1–31, 2010.
- Stefano Carattini, Maria Carvalho, and Sam Fankhauser. Overcoming public resistance to carbon taxes. *Wiley Interdisciplinary Reviews: Climate Change*, 9(5):e531, 2018.
- Matteo Castiglioni, Andrea Celli, Alberto Marchesi, and Nicola Gatti. Online bayesian persuasion. *Advances in Neural Information Processing Systems*, 33:16188–16198, 2020.
- Siyu Chen, Mengdi Wang, and Zhuoran Yang. Actions speak what you want: Provably sample-efficient reinforcement learning of the quantal stackelberg equilibrium from strategic feedbacks. *arXiv preprint arXiv:2307.14085*, 2023a.

- Siyu Chen, Jibang Wu, Yifan Wu, and Zhuoran Yang. Learning to incentivize information acquisition: Proper scoring rules meet principal-agent model. In *International Conference on Machine Learning*, pages 5194–5218. PMLR, 2023b.
- Yurong Chen, Qian Wang, Zhijian Duan, Haoran Sun, Zhaohua Chen, Xiang Yan, and Xiaotie Deng. Coordinated dynamic bidding in repeated second-price auctions with budgets. In *International Conference on Machine Learning*, pages 5052–5086. PMLR, 2023c.
- Ronald Coase. The problem of social cost. *Journal of Law and Economics*, 56(4):837 – 877, 2013. URL <https://EconPapers.repec.org/RePEc:ucp:jlawec:doi:10.1086/674872>.
- Natalie Collina, Eshwar Ram Arunachaleswaran, and Michael Kearns. Efficient stackelberg strategies for finitely repeated games. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 643–651, 2023.
- Robert Cooter. The cost of coase. *The Journal of Legal Studies*, 11(1):1–33, 1982.
- Qiwen Cui, Maryam Fazel, and Simon S Du. Learning optimal tax design in nonatomic congestion games. *arXiv preprint arXiv:2402.07437*, 2024.
- Carl J Dahlman. The problem of externality. *The journal of law and economics*, 22(1):141–162, 1979.
- Ilgın Dogan, Zuo-Jun Max Shen, and Anil Aswani. Estimating and incentivizing imperfect-knowledge agents with hidden rewards. *arXiv preprint arXiv:2308.06717*, 2023a.
- Ilgın Dogan, Zuo-Jun Max Shen, and Anil Aswani. Repeated principal-agent games with unobserved agent rewards and perfect-knowledge agents. *arXiv preprint arXiv:2304.07407*, 2023b.
- Kate Donahue, Nicole Immorlica, Meena Jagadeesan, Brendan Lucier, and Aleksandrs Slivkins. Impact of decentralized learning on player utilities in stackelberg games. *arXiv preprint arXiv:2403.00188*, 2024.
- Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. Simple versus optimal contracts. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 369–387, 2019.
- Seyed A Esmaili, Suho Shin, and Aleksandrs Slivkins. Robust and performance incentivizing algorithms for multi-armed bandits with strategic agents. *arXiv preprint arXiv:2312.07929*, 2023.
- Alireza Fallah and Michael I Jordan. Contract design with safety inspections. *arXiv preprint arXiv:2311.02537*, 2023.
- Alireza Fallah, Michael I Jordan, and Annie Ulichney. Fair allocation in dynamic mechanism design. *arXiv preprint arXiv:2406.00147*, 2024.
- Zhe Feng, Chara Podimata, and Vasilis Syrgkanis. Learning to bid without knowing your value. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 505–522, 2018.
- Thomas K Greenfield, Yu Ye, William Kerr, Jason Bond, Jürgen Rehm, and Norman Giesbrecht. Externalities from alcohol consumption in the 2005 us national alcohol survey: implications for policy. *International journal of environmental research and public health*, 6(12):3205–3224, 2009.
- Rachel Griffith, Martin O’Connell, and Kate Smith. Tax design in the alcohol market. *Journal of public economics*, 172:20–35, 2019.
- Guru Guruganesh, Jon Schneider, and Joshua R Wang. Contracts under moral hazard and adverse selection. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 563–582, 2021.
- Guru Guruganesh, Yoav Kolumbus, Jon Schneider, Inbal Talgam-Cohen, Emmanouil-Vasileios Vlatakis-Gkaragkounis, Joshua R Wang, and S Matthew Weinberg. Contracting with a learning agent. *arXiv preprint arXiv:2401.16198*, 2024.
- Bryan D Jones. Bounded rationality. *Annual review of political science*, 2(1):297–321, 1999.

- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6): 2590–2615, 2011.
- Jean-Jacques Laffont and David Martimort. The theory of incentives: the principal-agent model. In *The theory of incentives*. Princeton university press, 2009.
- John Langford and Tong Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. *Advances in neural information processing systems*, 20, 2007.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Ningyuan Li, Yunxuan Ma, Yang Zhao, Zhijian Duan, Yurong Chen, Zhilin Zhang, Jian Xu, Bo Zheng, and Xiaotie Deng. Learning-based ad auction design with externalities: the framework and a matching-based approach. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1291–1302, 2023.
- Yishay Mansour, Aleksandrs Slivkins, Vasilis Syrgkanis, and Zhiwei Steven Wu. Bayesian exploration: Incentivizing exploration in bayesian games. *arXiv preprint arXiv:1602.07570*, 2016.
- Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- Gillbert E Metcalf and David Weisbach. The design of a carbon tax. *Harv. Envtl. L. Rev.*, 33:499, 2009.
- James Mirrlees et al. *Tax by design: The Mirrlees review*. OUP Oxford, 2011.
- Roger B Myerson. *Mechanism design*. Springer, 1989.
- Noam Nisan and Amir Ronen. Algorithmic mechanism design. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 129–140, 1999.
- Vianney Perchet, Philippe Rigollet, Sylvain Chassang, and Erik Snowberg. Batched bandit problems. *The Annals of Statistics*, 44:660–681, 2016.
- Arthur Pigou. *The economics of welfare*. Routledge, 2017.
- Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, page 527–535, 1952.
- Tim Roughgarden. Algorithmic game theory. *Communications of the ACM*, 53(7):78–86, 2010.
- Bernard Salanié. *The economics of contracts: a primer*. MIT press, 2005.
- Bernard Salanie. *The economics of taxation*. MIT press, 2011.
- Abishek Sankararaman, Ayalvadi Ganesh, and Sanjay Shakkottai. Social learning in multi agent multi armed bandits. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 3(3):1–35, 2019.
- Antoine Scheid, Daniil Tiapkin, Etienne Boursier, Aymeric Capitaine, El Mahdi El Mhamdi, Éric Moulines, Michael I Jordan, and Alain Durmus. Incentivized learning in principal-agent bandit games. *ICML*, 2024.
- Reinhard Selten. Bounded rationality. *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft*, 146(4):649–658, 1990.
- Virag Shah, Jose Blanchet, and Ramesh Johari. Bandit learning with positive externalities. *Advances in Neural Information Processing Systems*, 31, 2018.
- Max Simchowitz and Aleksandrs Slivkins. Exploration and incentives in reinforcement learning. *Operations Research*, 2023.
- Aleksandrs Slivkins et al. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2):1–286, 2019.

- Heinrich Von Stackelberg. *Market structure and equilibrium*. Springer Science & Business Media, 2010.
- Serena Wang, Stephen Bates, P Aronow, and Michael Jordan. On counterfactual metrics for social welfare: Incentives, ranking, and information asymmetry. In *International Conference on Artificial Intelligence and Statistics*, pages 1522–1530. PMLR, 2024.
- Jibang Wu, Zixuan Zhang, Zhe Feng, Zhaoran Wang, Zhuoran Yang, Michael I Jordan, and Haifeng Xu. Markov persuasion processes and reinforcement learning. In *ACM Conference on Economics and Computation*, 2022a.
- Jibang Wu, Zixuan Zhang, Zhe Feng, Zhaoran Wang, Zhuoran Yang, Michael I Jordan, and Haifeng Xu. Sequential information design: Markov persuasion process and its efficient reinforcement learning. *arXiv preprint arXiv:2202.10678*, 2022b.
- Banghua Zhu, Stephen Bates, Zhuoran Yang, Yixin Wang, Jiantao Jiao, and Michael I Jordan. The sample complexity of online contract design. *arXiv preprint arXiv:2211.05732*, 2022.

A Algorithmic Subroutine for the Binary Search

Optimal Transfer. For any given round $t \geq 1$, action $a \in \mathcal{A}$ and $\varepsilon > 0$, the downstream player can incentivize any best-responding upstream player to choose a by offering a transfer, $\tau_a^{*,\varepsilon} \in \mathbb{R}_+$, defined as:

$$\tau_a^{*,\varepsilon} = \max_{a' \in \mathcal{A}} v^{\text{up}}(a') - v^{\text{up}}(a) + \varepsilon .$$

With this transfer, it holds that for any $a' \in \mathcal{A}$, $a' \neq a$, we have $v^{\text{up}}(a') < v^{\text{up}}(a) + \tau_a^{*,\varepsilon}$, ensuring the upstream player's action $A_t = a$, since action a yields a superior reward. Consequently,

$$\tau_a^* = \lim_{\varepsilon \rightarrow 0} \tau_a^{*,\varepsilon} = \max_{a' \in \mathcal{A}} v^{\text{up}}(a') - v^{\text{up}}(a)$$

represents the infimal transfer necessary to make arm a the best upstream player's choice.

First step of BELGIC: estimation of the optimal transfers. Suppose that we consider an arm $a \in \mathcal{A}$ and that the downstream player offers an incentive τ_a to the upstream player if he picks this arm. We consider this procedure with a constant incentive τ_a for a batch of time steps of length $\tilde{T} = \lceil T^\alpha \rceil$ due to the fact that the upstream player is learning [see [Perchet et al., 2016](#), for batched bandits in the usual multi-armed setting]. Lemma 2 shows that for BELGIC to accurately estimate τ_a^* with high probability, it must hold

$$C\tilde{T}^{\kappa+\beta/\alpha} < \tilde{T}/2 , \text{ which is equivalent to } CT^{\kappa\alpha+\beta-\alpha} < 1/2 , \quad (16)$$

This is why we impose the condition (13) on α and β , namely $\beta/\alpha < 1 - \kappa$, which ensures that $\kappa(\alpha - 1) + \beta < 0$ and therefore $\lim_{T \rightarrow +\infty} T^{\kappa(\alpha-1)+\beta} = 0$. More precisely, we define for $a \in [K]$

$$\Lambda_a = (a - 1) \lceil \log_2 T^\beta \rceil \tilde{T} ,$$

which is the step after which starts the binary search procedure on arm a . For any $d \in \{1, \dots, \lceil \log_2 T^\beta \rceil\}$ on arm a , we define

$$k_{a,d} = \Lambda_a + (d - 1)\tilde{T} \quad (17)$$

as the step after which starts the d -th batch iteration on arm a , and we consider

$$T_{a,d}^\neq = \text{Card} \{t \in \{k_{d,a} + 1, \dots, k_{d,a} + \tilde{T}\} \text{ such that } A_t \neq a\} , \quad (18)$$

where $(A_t)_{t \in \{1, \dots, K \lceil \log_2 T^\beta \rceil \tilde{T}\}}$ is given by Algorithm 2. Lemma 2 shows that for any $a \in \mathcal{A}$, $d \in \{1, \dots, \lceil \log_2 T^\beta \rceil\}$, with high probability

$$\text{if } T_{a,d}^\neq < \tilde{T} - C\tilde{T}^{\kappa+\beta/\alpha} \text{ and } T_{a,d}^\neq > C\tilde{T}^{\kappa+\beta/\alpha}, \text{ then } |\tau_a^* - \hat{\tau}_{d,a}| \leq 1/T^\beta , \quad (19)$$

where $\hat{\tau}_{d,a}$ is the current estimate of τ_a^* offered for iteration $k_{a,d}$. In case (18) does not hold, it means that the upstream player has misplayed and has chosen most of the steps a suboptimal action, leading to an instantaneous regret for him larger than the bound given in **H2**. This is why (19) holds with high probability. To sum up, the first phase of BELGIC consists in $\lceil \log_2 T^\beta \rceil$ batches of binary search on each arm $a \in \mathcal{A}$ to obtain a precision level $1/T^\beta$ on the optimal transfer τ_a^* .

During this phase in Algorithm 2, we define $\bar{\tau}_a(d) \in \mathbb{R}_+$ as the upper estimate and $\underline{\tau}_a(d) \in \mathbb{R}_+$ as the lower estimate of τ_a^* after $d \in \{1, \dots, \lceil \log_2 T^\beta \rceil\}$ rounds of binary search on arm a . For any $t \in [T]$ and $a \in \mathcal{A}$, we define $\tau_a^{\text{mid}}(t) = (\bar{\tau}_a(t) + \underline{\tau}_a(t))/2$. $\bar{\tau}_a(d)$, $\underline{\tau}_a(d)$, $\tau_a^{\text{mid}}(d)$ are updated at the end of the d -th binary search batch of length \tilde{T} on arm a . We define $N_T = \lceil \log_2 T^\beta \rceil$ as the number of binary search steps per arm.

After this first binary search phase, the downstream player computes estimates of the optimal transfers τ_a^*

$$(\hat{\tau}_a)_{a \in \mathcal{A}} = (\bar{\tau}_a(\lceil \log_2 T^\beta \rceil) + 1/T^\beta + CT^{(\kappa-1)/2})_{a \in \mathcal{A}} ,$$

and offers these transfers $(\tau_a^*)_{a \in \mathcal{A}}$ to make the upstream player play any action $\tilde{a} \in \mathcal{A}$ she wants.

Second Step. After the first phase during which the optimal incentives are estimated by the downstream player through $(\hat{\tau}_a)_{a \in \mathcal{A}}$, she runs in the second phase the subroutine **BANDIT-ALG** on the $\mathcal{A} \times \mathcal{A}$ bandit instance driven by her action B_t and the upstream player's one A_t . More precisely, any bandit subroutine **BANDIT-ALG**, such as UCB or ε -greedy, for instance, can be run in a black-box fashion

on the shifted bandit instance, where rewards are shifted by the upper estimated transfers $(\hat{\tau}_a)_{a \in \mathcal{A}}$. The downstream player computes a shifted history $\tilde{\mathcal{H}}_t^{\text{down,p}} = \emptyset$ for any $t \leq K \lceil T^\alpha \rceil \lceil \log_2 T^\beta \rceil$ and for any $t > K \lceil T^\alpha \rceil \lceil \log_2 T^\beta \rceil$

$$\tilde{\mathcal{H}}_t^{\text{down,p}} = \begin{cases} \{\tilde{a}_t, B_t, \tau(t), U_t, X_{\tilde{a}_t, B_t}(t) - \hat{\tau}_{\tilde{a}_t}\} \cup \tilde{\mathcal{H}}_{t-1}^{\text{down,p}} & \text{if } \tilde{a}_t = A_t \\ \tilde{\mathcal{H}}_{t-1}^{\text{down,p}} & \text{otherwise.} \end{cases}$$

which serves to feed `Bandit-Alg`, following `Bandit-Alg`: $(U_t, \tilde{\mathcal{H}}_{t-1}^{\text{down,p}}) \mapsto (\tilde{a}_t, B_t) \in \mathcal{A} \times \mathcal{A}$. Note that here, `Bandit-Alg` aims to maximize the shifted reward $(\nu^{\text{down}}(a, b) - \hat{\tau}_a)_{(a,b) \in \mathcal{A}^2}$.

Based on this decision by `Bandit-Alg`, Π_p^{down} offers the incentive $\hat{\tau}_{\tilde{a}_t}$ associated with action \tilde{a}_t and plays action B_t . Lemma 5 ensures that \tilde{a}_t is the upstream player's best choice. Therefore, **H2** ensures that the upstream player will not deviate from the downstream player's recommendation with high probability.

B Invariance when the property rights are given to the downstream player

Our focus in the paper was the case where the upstream player possesses the bandit instance and receives monetary payments. We argue here that the symmetric situation, i.e. when the property rights are given to the downstream player, can be analysed in the exact same way.

Assume that the the downstream player owns the bandit instance. This implies that (i) they can prescribe what arm the upstream player has to play at each round, and (ii) the upstream player may perform a monetary transfer to influence the arm they are allowed to pull.

Consider the same bandit setup as before. Formally, the downstream player's action is $(A_t, B_t) \in \mathcal{A} \times \mathcal{A}$ where A_t is the arm that the upstream player is prescribed to pull (he cannot deviate since the downstream player has the property rights), while B_t is the arm played by the downstream player. On the other hand, the upstream player's policy outputs at each round the action $(\tilde{a}_t, \tau(t)) \in \mathcal{A} \times \mathbb{R}_+$, where \tilde{a}_t is the arm they choose to incentivize and $\tau(t)$ is the amount of transfer. It means that the downstream player receives a transfer $\tau(t)$ if she prescribes action $A_t = \tilde{a}_t$. As a consequence, the instantaneous utility of the upstream player is $Z_{A_t} - \mathbb{1}_{\tilde{a}_t}(A_t)\tau(t)$, while the downstream player receives $X_{A_t, B_t} + \mathbb{1}_{\tilde{a}_t}(A_t)\tau(t)$. In that case, the upstream player may perform a binary search on each arm $\bar{a} \in \mathcal{A}$ to identify the optimal incentive $\tau_{\bar{a}}^*$, by considering `Card` $\{t \in [T] \text{ such that } \tilde{a}_t = \bar{a}\}$ during batches designed for the binary search and then play on the shifted bandit instance as we explained. This situation and the upstream player's strategy are now equivalent to the one presented before.

C Proofs and Technical Results

Recall that we defined the shifted history $\tilde{\mathcal{H}}_t^{\text{down,p}}$ that will serve to feed Π_p^{down} at time t as $\tilde{\mathcal{H}}_t^{\text{down,p}} = (\tilde{a}_s, \tau(s), A_s, V_s, Z_{A_s}(s))_{s \leq t}$.

Theorem 4. *Suppose that $\text{argmax}_{a \in \mathcal{A}} v^{\text{up}}(a)$ is the singleton $\{a_\star^{\text{u}}\}$ and that*

$$v^{\text{up}}(a^{\text{sw}}) + v^{\text{down}}(a^{\text{sw}}, b^{\text{sw}}) - v^{\text{up}}(a_\star^{\text{u}}) + v^{\text{down}}(a_\star^{\text{u}}, b) > 0,$$

for any $b \in \mathcal{A}$. In the absence of property rights and when the upstream player runs any no-regret policy Π_n^{up} , we have $\mathfrak{R}^{\text{sw}}(T, \Pi_n^{\text{up}}, \Pi_n^{\text{down}}) \geq T \Delta^{\text{sw}} - \mathfrak{R}_n^{\text{up}}(T, \Pi_n^{\text{up}}) \Delta^{\text{sw}} / \Delta^{\text{up}}$, where $\Delta^{\text{up}} = \min_{a' \in \mathcal{A} \setminus \{a_\star^{\text{u}}\}} v^{\text{up}}(a_\star^{\text{u}}) - v^{\text{up}}(a')$ and $\Delta^{\text{sw}} = v^{\text{up}}(a^{\text{sw}}) + v^{\text{down}}(a^{\text{sw}}, b^{\text{sw}}) - \max_{b \in \mathcal{A}} (v^{\text{up}}(a_\star^{\text{u}}) + v^{\text{down}}(a_\star^{\text{u}}, b))$. Therefore, $\mathfrak{R}^{\text{sw}}(T, \Pi_n^{\text{up}}, \Pi_n^{\text{down}}) = \Omega(T)$ and $(\Pi_n^{\text{up}}, \Pi_n^{\text{down}})$ is not welfare efficient.

Proof of Theorem 4. Since $\text{argmax}_{a \in \mathcal{A}} v^{\text{up}}(a)$ is the singleton $\{a_\star^{\text{u}}\}$, we define $\Delta^{\text{up}} = \min_{a' \in \mathcal{A} \setminus \{a_\star^{\text{u}}\}} v^{\text{up}}(a_\star^{\text{u}}) - v^{\text{up}}(a')$ as the upstream player reward gap and $\Delta^{\text{sw}} = v^{\text{up}}(a^{\text{sw}}) + v^{\text{down}}(a^{\text{sw}}, b^{\text{sw}}) - \max_{b \in \mathcal{A}} (v^{\text{up}}(a_\star^{\text{u}}) + v^{\text{down}}(a_\star^{\text{u}}, b))$ as the social welfare reward gap if the upstream player plays his most preferred action.

Denote $N_\star^{\text{up}}(T)$ the number of pulls of the upstream player up to time T on the arm a_\star^{u} . By definition of Δ^{up} , we have that for any step $t \in [T]$ such that $A_t \neq a_\star^{\text{u}}$, $\max_{a \in \mathcal{A}} \{v^{\text{up}}(a)\} - v^{\text{up}}(A_t) =$

$v^{\text{up}}(a_\star^{\text{up}}) - v^{\text{up}}(A_t) \geq \min_{a' \in \mathcal{A} \setminus \{a_\star^{\text{u}}\}} v^{\text{up}}(a_\star^{\text{u}}) - v^{\text{up}}(a') = \Delta^{\text{up}}$. There are $T - N_\star^{\text{up}}(T)$ such steps, which leads to

$$\begin{aligned} \mathfrak{R}_n^{\text{up}}(T, \Pi_n^{\text{up}}, \Pi_p^{\text{down}}) &\geq \mathbb{E} \left[\sum_{t=1}^T \max_{a \in \mathcal{A}} \{v^{\text{up}}(a)\} - v^{\text{up}}(A_t) \right] \\ &\geq (T - \mathbb{E}[N_\star^{\text{u}}(T)]) \Delta^{\text{up}}, \end{aligned}$$

and we obtain

$$\mathbb{E}[N_\star^{\text{u}}(T)] \geq T - \mathfrak{R}_n^{\text{up}}(T, \Pi_n^{\text{up}}) / \Delta^{\text{up}}.$$

Moreover, for any $t \in [T]$ such that $A_t = a_\star^{\text{up}}$ and any $B_t \in \mathcal{A}$, $v^{\text{up}}(a^{\text{sw}}) + v^{\text{down}}(a^{\text{sw}}, b^{\text{sw}}) - (v^{\text{up}}(A_t) + v^{\text{down}}(A_t, B_t)) = v^{\text{up}}(a^{\text{sw}}) + v^{\text{down}}(a^{\text{sw}}, b^{\text{sw}}) - (v^{\text{up}}(a_\star^{\text{up}}) + v^{\text{down}}(a_\star^{\text{up}}, B_t)) \geq \Delta^{\text{sw}}$ by definition, which leads to

$$\mathfrak{R}^{\text{sw}}(T, \Pi_n^{\text{up}}, \Pi_n^{\text{down}}) \geq \mathbb{E}[N_\star^{\text{u}}(T)] \Delta^{\text{sw}},$$

and we obtain

$$\mathfrak{R}^{\text{sw}}(T, \Pi_n^{\text{up}}, \Pi_n^{\text{down}}) \geq T \Delta^{\text{sw}} - \mathfrak{R}_n^{\text{up}}(T, \Pi_n^{\text{up}}) \Delta^{\text{sw}} / \Delta^{\text{up}}.$$

Since $\lim_{T \rightarrow +\infty} \mathfrak{R}_n^{\text{up}}(T, \Pi_n^{\text{up}}) / T = 0$, we have $\lim_{T \rightarrow +\infty} \mathfrak{R}^{\text{sw}}(T, \Pi_n^{\text{up}}, \Pi_n^{\text{down}}) / T = \Delta^{\text{sw}} > 0$, hence the result. \square

The proof of Theorem 2 is an immediate consequence of Theorem 4.

Lemma 1. Recall that $\mu^{\star, \text{down}}$ is the downstream player's optimal reward as defined as a solution of (10). We have $\mu^{\star, \text{down}} = \max_{a, b \in \mathcal{A}} \{v^{\text{down}}(a, b) + v^{\text{up}}(a)\} - \max_{a' \in \mathcal{A}} \{v^{\text{up}}(a')\}$, as well as $(a^{\text{opt}}, b^{\text{opt}}) = (a^{\text{sw}}, b^{\text{sw}})$ and $\mu^{\star, \text{up}} + \mu^{\star, \text{down}} = v^{\text{up}}(a^{\text{sw}}) + v^{\text{down}}(a^{\text{sw}}, b^{\text{sw}}) = \max_{a, b \in \mathcal{A}} \{v^{\text{up}}(a) + v^{\text{down}}(a, b)\}$, where $\mu^{\star, \text{up}}$ is defined in Equation (1). Moreover, for any integer $T \in \mathbb{N}^*$, and policies $\Pi_p^{\text{up}}, \Pi_p^{\text{down}}$, we have that

$$\mathfrak{R}^{\text{sw}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq \mathfrak{R}_p^{\text{up}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) + \mathfrak{R}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}).$$

Proof of Lemma 1. Recall that $\mu^{\star, \text{down}}$ is defined as $\mu^{\star, \text{down}} = \sup_{a, b \in \mathcal{A}^2, \tau \in \mathbb{R}_+} \{v^{\text{down}}(a, b) - \tau\}$, such that $a \in \arg\max_{a' \in \mathcal{A}} \{v^{\text{up}}(a') + \mathbb{1}_a(a')\tau\}$ and $a_\star^{\text{up}} = \arg\max_{a' \in \mathcal{A}} v^{\text{up}}(a')$. Note that we can write

$$\mu^{\star, \text{down}} = \max \{ \sup_{a, b \in \mathcal{A}^2, \tau \in \mathbb{R}_+} \mathbb{1}_{\tilde{\mathcal{A}}}(a, \tau) (v^{\text{down}}(a, b) - \tau), \max_{b \in \mathcal{A}} v^{\text{down}}(a_\star^{\text{up}}, b) \},$$

where $\tilde{\mathcal{A}} = \{(a, \tau) : v^{\text{up}}(a) + \tau \geq \max_{a'} v^{\text{up}}(a') + \mathbb{1}_a(a')\tau\}$ which is the set of pairs $(a, \tau) \in \mathcal{A} \times \mathbb{R}_+$ such that the constraint binds. However, we also have by definition that $v^{\text{up}}(a_\star^{\text{up}}) + 0 \geq \max_{a' \in \mathcal{A}} v^{\text{up}}(a') + \mathbb{1}_{a_\star^{\text{up}}}(a') \cdot 0$ and hence, $(a_\star^{\text{up}}, 0) \in \tilde{\mathcal{A}}$. Therefore, since $v^{\text{down}}(a_\star^{\text{up}}, b) \geq 0$ for any $b \in \mathcal{A}$, we can write

$$\mu^{\star, \text{down}} = \sup_{(a, \tau) \in \tilde{\mathcal{A}}, b \in \mathcal{A}} \{v^{\text{down}}(a, b) - \tau\}.$$

First note that if $(a, \tau) \in \tilde{\mathcal{A}}$, then for any $a' \in \mathcal{A}$, $\tau \in \mathbb{R}_+$, $v^{\text{up}}(a) + \tau \geq v^{\text{up}}(a') + \mathbb{1}_a(a')\tau$, which gives

$$v^{\text{up}}(a) - v^{\text{up}}(a') \geq (\mathbb{1}_a(a') - 1)\tau.$$

However, either $a = a'$ and hence $v^{\text{up}}(a) - v^{\text{up}}(a') = 0$, either $a \neq a'$ and hence $\mathbb{1}_a(a') = 0$. Therefore, we have that

$$v^{\text{up}}(a) - v^{\text{up}}(a') \geq -\tau,$$

which implies by definition of the optimal incentives that $\tau \geq v^{\text{up}}(a') - v^{\text{up}}(a)$ for any $a' \in \mathcal{A}$, and hence $\tau_a^* \leq \tau$ for any $(a, \tau) \in \tilde{\mathcal{A}}$.

In addition, $(a, \tau_a^*) \in \tilde{\mathcal{A}}$ by definition. Consequently,

$$\mu^{\star, \text{down}} = \max_{a, b \in \mathcal{A}^2} \{v^{\text{down}}(a, b) - \tau_a^*\} = \max_{a, b \in \mathcal{A}} \{v^{\text{down}}(a, b) - \max_{a' \in \mathcal{A}} \{v^{\text{up}}(a')\} + v^{\text{up}}(a)\},$$

hence the first part of the result. Since $\mu^{*,\text{up}}$ is defined as $\mu^{*,\text{up}} = \max_{a \in \mathcal{A}} v^{\text{up}}(a)$, we have that

$$\begin{aligned} \mu^{*,\text{down}} + \mu^{*,\text{up}} &= \max_{a,b \in \mathcal{A}} \{v^{\text{down}}(a,b) - \max_{a' \in \mathcal{A}} \{v^{\text{up}}(a')\} + v^{\text{up}}(a)\} + \max_{a \in \mathcal{A}} v^{\text{up}}(a) \\ &= \max_{a,b \in \mathcal{A}} \{v^{\text{down}}(a,b) + v^{\text{up}}(a)\} \\ &= v^{\text{up}}(a^{\text{sw}}) + v^{\text{down}}(a^{\text{sw}}, b^{\text{sw}}). \end{aligned}$$

Now summing $\mathfrak{R}_p^{\text{down}}$ and $\mathfrak{R}_p^{\text{up}}$ as defined in (9) and (11), we obtain

$$\begin{aligned} &\mathfrak{R}_p^{\text{up}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) + \mathfrak{R}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \\ &= \mathbb{E} \left[\sum_{t=1}^T \max_{a \in \mathcal{A}} \{v^{\text{up}}(a) + \mathbb{1}_{\tilde{a}_t}(a)\tau(t)\} - (v^{\text{up}}(A_t) + \mathbb{1}_{\tilde{a}_t}(A_t)\tau(t)) \right] \\ &\quad + T \max_{a,b \in \mathcal{A}^2} \{v^{\text{down}}(a,b) - \max_{a' \in \mathcal{A}} \{v^{\text{up}}(a')\} + v^{\text{up}}(a)\} - \mathbb{E} \left[\sum_{t=1}^T v^{\text{down}}(A_t, B_t) - \mathbb{1}_{\tilde{a}_t}(A_t)\tau(t) \right] \\ &\geq T \max_{a \in \mathcal{A}} \{v^{\text{up}}(a)\} - \mathbb{E} \left[\sum_{t=1}^T v^{\text{up}}(A_t) + \mathbb{1}_{\tilde{a}_t}(A_t)\tau(t) \right] - \mathbb{E} \left[\sum_{t=1}^T v^{\text{down}}(A_t, B_t) - \mathbb{1}_{\tilde{a}_t}(A_t)\tau(t) \right] \\ &\quad + T \max_{a,b \in \mathcal{A}^2} \{v^{\text{down}}(a,b) + v^{\text{up}}(a)\} - T \max_{a' \in \mathcal{A}} \{v^{\text{up}}(a')\} \\ &= T \max_{a,b \in \mathcal{A}^2} \{v^{\text{up}}(a) + v^{\text{down}}(a,b)\} - \mathbb{E} \left[\sum_{t=1}^T v^{\text{up}}(A_t) + v^{\text{down}}(A_t, B_t) \right] \\ &= \mathfrak{R}^{\text{sw}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}), \end{aligned}$$

hence the result. \square

For a downstream player's policy Π_p^{down} , we define $\tilde{\mathfrak{R}}_p^{\text{up}}$ as the upstream player's regret without expectation, following

$$\tilde{\mathfrak{R}}_p^{\text{up}}(\{s+1, \dots, s+t\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) = \sum_{l=s+1}^{s+t} \max_{a \in \mathcal{A}} \{v^{\text{up}}(a) + \mathbb{1}_{\tilde{a}_l}(a)\tau(l)\} - (v^{\text{up}}(A_l) + \mathbb{1}_{\tilde{a}_l}(A_l)\tau(l)),$$

where $(\tilde{a}_l, \tau(l))_{l \in [T]}$ are the incentives output by Π_p^{down} and $(A_l)_{l \in [T]}$ are the output of Π_p^{up} . Recall the assumption that we use on the upstream player's regret for a policy Π_p^{up} . We show here that it is satisfied by typical no-regret bandit algorithms.

H2. *There exist $C, \zeta > 0, \kappa \in [0, 1)$ such that for any $s, t \in [T]$ with $s+t \leq T$, any $\{\tau_a\}_{a \in [K]} \in \mathbb{R}_+^K$ and any policy Π_p^{down} that offers almost surely a transfer $(\tilde{a}_l, \tau(l)) = (\tilde{a}_l, \tau_{\tilde{a}_l})$ for any $l \in \{s+1, \dots, s+t\}$, the batched regret of the upstream player following Π_p^{up} satisfies, with probability at least $1 - t^{-\zeta}$,*

$$\sum_{l=s+1}^{s+t} \max_{a \in \mathcal{A}} \{v^{\text{up}}(a) + \mathbb{1}_{\tilde{a}_l}(a)\tau_{\tilde{a}_l}\} - (v^{\text{up}}(A_l) + \mathbb{1}_{\tilde{a}_l}(A_l)\tau_{\tilde{a}_l}) \leq Ct^\kappa.$$

We present the upstream player's UCB subroutine.

Proposition 2. *Let $s, t \in [T]$ such that $s+t \leq T$. Suppose that there exists a family $(\tau_a)_{a \in \mathcal{A}}$ of constant incentives associated with each arm $a \in \mathcal{A}$ such that we have in Algorithm 3: $(\tilde{a}_l, \tau(l))_{s+l \in \{s+1, \dots, s+t\}} = (\tilde{a}_l, \tau_{\tilde{a}_l})_{l \in \{s+1, \dots, s+t\}}$ as an output of Π_p^{down} . Suppose that the distributions γ_a are 1-sub-Gaussian. Then with probability at least $1 - T^{-2}$, the regret of the version of UCB given in Algorithm 3 run by the upstream player satisfies*

$$\tilde{\mathfrak{R}}_p^{\text{up}}(\{s+1, \dots, s+t\}, \text{UCB}, \Pi_p^{\text{down}}) \leq 8\sqrt{\log(KT^3)}\sqrt{tK}.$$

Note that the major difference between this assumption and the regret bounds that we generally consider in multi-armed bandit problems is that we consider the regret without expectation here.

Algorithm 3 Upstream player's UCB

- 1: **Input:** Set of arms K , horizon T .
 - 2: **Initialize:** For any arm $a \in [K]$, set $\hat{\mu}_a = 0$, $T_a = 0$.
 - 3: **for** $1 \leq t \leq K$: **do**
 - 4: Pull arm $A_t = t$
 - 5: Update $\hat{\mu}_{A_t} = X_{A_t}(t)$, $T_{A_t}(t) = 1$
 - 6: **end for**
 - 7: **for** $t \geq K + 1$ **do**
 - 8: Observe the incentive $(\tilde{a}_t, \tau(t))$.
 - 9: Pull arm $A_t \in \operatorname{argmax}_{a \in [K]} \left\{ \hat{\mu}_a(t-1) + 2\sqrt{\frac{\log(KT^3)}{T_a(t-1)}} + \mathbb{1}_{\tilde{a}_t}(a)\tau(t) \right\}$
 - 10: Update $T_{A_t}(t) = T_{A_t}(t-1) + 1$, $\hat{\mu}_{A_t}(t) = \frac{1}{T_{A_t}(t)}(T_{A_t}(t-1)\hat{\mu}_{A_t}(t-1) + X_{A_t}(t))$
 - 11: **end for**
-

Proof of Proposition 2. The proof is adapted from the proof of Bubeck et al. [2012, Theorem 2.1]. Let $s, t \in [T]$ such that $s + t \leq T$. For any integer $l \in \{s + 1, \dots, s + t\}$, we write $n_l(a) = \operatorname{Card} \{l' \in [l] \text{ such that } A_{l'} = a\}$ for the number of pulls of arm a and $\hat{\mu}_a(l)$ for the empirical mean utility of the arm $a \in \mathcal{A}$ estimated on the batch $\{1, \dots, l\}$: $\hat{\mu}_a(l) = n_l(a)^{-1} \sum_{k=1}^l \mathbb{1}_a(A_k) X_a(k)$. Since the rewards on the incentivized bandit instance are 1-sub-Gaussian, a Hoeffding bound gives that for any $\delta \in (0, 1)$, $a \in \mathcal{A}$, $l \in \{s + 1, \dots, s + t\}$, and any family of arms $(a_{l'})_{l' \in \{1, \dots, l\}}$ such that $\operatorname{Card} \{j: a_j = a\} = k$, we have that

$$\mathbb{P} \left(|\hat{\mu}_a(l) - v^{\text{up}}(a)| \geq 2\sqrt{\log(2/\delta)/k} \mid (A_{l'})_{l' \in \{1, \dots, s+l\}} = (a_{l'})_{l' \in \{1, \dots, s+l\}} \right) \leq \delta.$$

Therefore, for any $\delta \in (0, 1)$, $a \in \mathcal{A}$, $l \in \{s + 1, \dots, s + t\}$, we have the following bound

$$\begin{aligned} & \mathbb{P} \left(|\hat{\mu}_a(l) - v^{\text{up}}(a)| \geq 2\sqrt{\log(2/\delta)/n_l(a)} \right) \\ &= \mathbb{P} \left(\bigcup_{k=1}^l \bigcup_{\substack{(a_{l'})_{l' \text{ s.t.}} \\ \operatorname{Card} \{l' \in \{1, \dots, l\}: a_{l'} = a\} = k}} \{ |\hat{\mu}_a(l) - v^{\text{up}}(a)| \geq 2\sqrt{\log(2/\delta)/k} \} \right) \\ &\leq \sum_{k=1}^l \mathbb{P} \left(\bigcup_{\substack{(a_{l'})_{l' \text{ s.t.}} \\ \operatorname{Card} \{l' \in \{1, \dots, l\}: a_{l'} = a\} = k}} \{ |\hat{\mu}_a(s+l) - v^{\text{up}}(a)| \geq 2\sqrt{\log(2/\delta)/k} \} \right) \\ &\leq \sum_{k=1}^l \sum_{\substack{(a_{l'})_{l' \text{ s.t.}} \\ \operatorname{Card} \{l' \in \{1, \dots, l\}: a_{l'} = a\} = k}} \mathbb{P} \left(|\hat{\mu}_a(s+l) - v^{\text{up}}(a)| \geq 2\sqrt{\log(2/\delta)/k} \mid A_l = a_l \right) \mathbb{P}(A_l = a_l) \\ &\leq T\delta, \end{aligned}$$

and with an union bound, we obtain that

$$\mathbb{P} \left(\exists l \in [t], a \in \mathcal{A} \text{ such that } |\hat{\mu}_a(l) - v^{\text{up}}(a)| \geq 2\sqrt{\log(2/\delta)/n_l(a)} \right) \leq T^2 K \delta,$$

Considering the probability of the opposite event, we have that

$$\mathbb{P} \left(\text{for any } l \in [t], a \in \mathcal{A}, |\hat{\mu}_a(l) - v^{\text{up}}(a)| \leq 2\sqrt{\log(2T^2K/\delta)/n_l(a)} \right) \geq 1 - \delta, \quad (20)$$

where we rescaled δ as δ/T^2K .

For the remaining of the proof, we take $\delta = T^{-2}$ and define $a_l^* = \operatorname{argmax}_{a \in \mathcal{A}} \{v^{\text{up}}(a) + \mathbb{1}_{\tilde{a}_l}(a)\tau\}$. We now assume that the event $\{\text{for any } l \in \{s + 1, \dots, s + t\}, a \in \mathcal{A}, |\hat{\mu}_a(l) - v^{\text{up}}(a)| \leq 2\sqrt{\log(2T^2K/\delta)/n_l(a)}\}$ holds. If at some step $l \in \{s + 1, \dots, s + t\}$, action A_l is chosen in Algorithm 3, it means that

$$\hat{\mu}_{A_l}(l) + 2\sqrt{\log(tK/\delta)/n_l(A_l)} + \mathbb{1}_{\tilde{a}_l}(A_l)\tau_{\tilde{a}_l} \geq \hat{\mu}_{a_l^*}(l) + 2\sqrt{\log(tK/\delta)/n_l(a_l^*)} + \mathbb{1}_{\tilde{a}_l}(a_l^*)\tau_{\tilde{a}_l},$$

with regards to the choice of actions in UCB based on the upper confidence bound. We now decompose the whole regret on the batch $\{s+1, \dots, s+t\}$ defined as $\tilde{\mathfrak{R}}_p^{\text{up}}(\{s+1, \dots, s+t\}, \text{UCB}, \Pi_p^{\text{down}})$. (20) ensures that with probability at least $1 - 1/T^2$

$$\begin{aligned}
\tilde{\mathfrak{R}}_p^{\text{up}}(\{s+1, \dots, s+t\}, \text{UCB}, \Pi_p^{\text{down}}) &= \sum_{l=s+1}^{s+t} v^{\text{up}}(a_l^*) + \mathbb{1}_{\tilde{a}_l}(a_l^*)\tau_{\tilde{a}_l} - (v^{\text{up}}(A_l) + \mathbb{1}_{\tilde{a}_l}(A_l)\tau_{\tilde{a}_l}) \\
&\leq \sum_{l=s+1}^{s+t} \hat{\mu}_{a_l^*}(l) + 2\sqrt{\log(Kt/\delta)/n_l(a_l^*)} - v^{\text{up}}(A_l) + \mathbb{1}_{\tilde{a}_l}(a_l^*)\tau_{\tilde{a}_l} - \mathbb{1}_{\tilde{a}_l}(A_l)\tau_{\tilde{a}_l} \\
&\leq \sum_{l=s+1}^{s+t} \hat{\mu}_{A_l}(l) + 2\sqrt{\log(Kt/\delta)/n_l(A_l)} - v^{\text{up}}(A_l) \\
&\leq \sum_{l=s+1}^{s+t} \hat{\mu}_{A_l}(l) + 2\sqrt{\log(Kt/\delta)/n_l(A_l)} - (\hat{\mu}_{A_l}(l) - 2\sqrt{\log(Kt/\delta)/n_l(A_l)}) \\
&\leq 4\sqrt{\log(Kt/\delta)} \sum_{l=s+1}^{s+t} \sqrt{1/n_l(A_l)},
\end{aligned}$$

and we have

$$\begin{aligned}
\sum_{l=s+1}^{s+t} \sqrt{1/n_l(A_l)} &\leq \sum_{i=1}^K \sum_{l=s+1}^{s+t} \sqrt{\mathbb{1}_{A_l}(i)/n_l(i)} \\
&\leq \sum_{i=1}^K \sum_{j=n_{s+1}(i)}^{n_{s+t}(i)} 1/\sqrt{j} \leq 2 \sum_{i=1}^K \sqrt{n_{s+t}(i) - n_s(i)}, \tag{21}
\end{aligned}$$

where the last step holds because for any integers $s, t \in \mathbb{N}^*$, we have that

$$\sum_{l=s+1}^{s+t} \frac{1}{\sqrt{l}} = \sum_{l=s+1}^{s+t} \frac{1}{\sqrt{l}} \int_{x=l-1}^l dx \leq \sum_{l=s+1}^{s+t} \int_{x=l-1}^l \frac{dx}{\sqrt{x}} = \int_{x=s}^{s+t} \frac{dx}{\sqrt{x}} = 2(\sqrt{s+t} - \sqrt{s}).$$

Using Cauchy-Schwarz inequality we obtain from (21)

$$1/K \sum_{l=s+1}^{s+t} \sqrt{1/n_l(A_l)} \leq 2\sqrt{1/K \sum_{i=1}^K n_{s+t}(i) - n_s(i)} = 2\sqrt{t/K},$$

which gives $\sum_{l=s+1}^{s+t} \sqrt{1/n_l(A_l)} \leq 2\sqrt{tK}$.

Finally plugging all the terms together, since $\delta = 1/T^2$, we obtain that with probability at least $1 - 1/T^2$

$$\tilde{\mathfrak{R}}_p^{\text{up}}(\{s+1, \dots, s+t\}, \text{UCB}) \leq 8\sqrt{\log(KT^3)}\sqrt{tK}. \quad \square$$

Lemma 2. Assume **H2** holds and consider some arm $a \in \mathcal{A}$ such that we run the d -th batch of binary search on a with $d \in \{1, \dots, \lceil \log_2 T^\beta \rceil\}$: for any $t \in \{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}$, we have $(\tilde{a}_t, \tau(t)) = (a, \tau_a)$ with $\tau_a = \tau_a^{\text{mid}}(d)$ and $\tilde{T} = \lceil T^\alpha \rceil$. Recall that we defined in (18): $T_{a,d}^\neq = \text{Card} \{t \in \{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\} \text{ such that } A_t \neq a\}$. Let $\beta \in (0, 1)$ be such that $\beta < \alpha(1 - \kappa)$. Given that the event $\{\tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq C\tilde{T}^\kappa\}$ holds, we have that

- If $T_{a,d}^\neq < \tilde{T} - C\tilde{T}^{\kappa+\beta/\alpha}$, then $\tau_a^* < \tau_a + 1/T^\beta$.
- If $T_{a,d}^\neq > C\tilde{T}^{\kappa+\beta/\alpha}$, then $\tau_a^* > \tau_a - 1/T^\beta$.

Consequently, with probability at least $1 - 2T^{-\alpha\zeta}$, if $C\tilde{T}^{\kappa+\beta/\alpha} < T_{a,d}^\neq < \tilde{T} - C\tilde{T}^{\kappa+\beta/\alpha}$, then $|\tau_a^* - \tau_a| \leq 1/T^\beta$.

Proof of Lemma 2. The whole proof is done conditionally on the event

$$\{\tilde{\mathfrak{A}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq C\tilde{T}^{\kappa}\}.$$

Note that it holds with probability at least $1 - \tilde{T}^{-\zeta} \geq 1 - T^{-\alpha\zeta}$ since we suppose that Π_p^{up} satisfies **H2**.

Suppose that we have $\tau_a \geq \tau_a^* + 1/T^\beta$. By definition of the optimal incentives, we obtain, using by assumption $\tau_a \geq \tau_a^* + 1/T^\beta$

$$\begin{aligned} \mathbb{1}_a(a)\tau_a + v^{\text{up}}(a) &\geq \tau_a^* + v^{\text{up}}(a) + 1/T^\beta \\ &= \max_{a' \in \mathcal{A}} v^{\text{up}}(a') - v^{\text{up}}(a) + v^{\text{up}}(a) + 1/T^\beta \\ &= \max_{a'} v^{\text{up}}(a') + 1/T^\beta, \end{aligned}$$

which ensures that a is the optimal arm for the upstream player during the batch $\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}$ that we consider. In that case, since a is the best arm, by definition of the upstream player's utility, the reward gap $\mathbb{E}[v^{\text{up}}(a) + \tau_a - (v^{\text{up}}(A_t) + \mathbb{1}_a(A_t)\tau_a)]$ for the upstream player at any step $t \in \{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}$ is at least $v^{\text{up}}(a) + \tau_a - \max_{a' \in \mathcal{A}} v^{\text{up}}(a')$, and we obtain

$$C\tilde{T}^{\kappa} \geq \tilde{\mathfrak{A}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \geq T_{a,d}^{\neq}(\tau_a + v^{\text{up}}(a) - \max_{a' \in \mathcal{A}} \{v^{\text{up}}(a')\}),$$

by **H2**, $T_{a,d}^{\neq}$ being the number of steps for which a suboptimal arm has been chosen. Therefore, by definition of the optimal incentives, we obtain that

$$C\tilde{T}^{\kappa} \geq T_{a,d}^{\neq}(\tau_a - \tau_a^*) \geq T_{a,d}^{\neq}/T^\beta \geq T_{a,d}^{\neq} \tilde{T}^{-\beta/\alpha},$$

which gives: $T_{a,d}^{\neq} \leq C\tilde{T}^{\kappa+\beta/\alpha}$. Therefore, if we take the contrapositive, we obtain that during the sequence $\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}$, if $T_{a,d}^{\neq} > C\tilde{T}^{\kappa+\beta/\alpha}$, then with probability at least $1 - T^{-\alpha\zeta}$, $\tau_a < \tau_a^* + 1/T^\beta$, or equivalently $\tau_a^* > \tau_a - 1/T^\beta$.

Now suppose that $\tau_a \leq \tau_a^* - 1/T^\beta$. By definition of the optimal incentives, we obtain

$$\begin{aligned} \mathbb{1}_a(a)\tau_a + v^{\text{up}}(a) &\leq \tau_a^* + v^{\text{up}}(a) \\ &\leq \max_{a' \in \mathcal{A}} v^{\text{up}}(a') - v^{\text{up}}(a) + v^{\text{up}}(a) - 1/T^\beta \\ &= \max_{a'} v^{\text{up}}(a') - 1/T^\beta, \end{aligned}$$

which ensures that a is a suboptimal arm for the upstream player during this batch of time steps. Therefore, arm a which has a reward gap bigger than $1/T^\beta$ since $\max_{a' \in \mathcal{A}} \{v^{\text{up}}(a) + \mathbb{1}_a(a')\} - (v^{\text{up}}(a) + \tau_a) \geq 1/T^\beta$ and arm a has been picked $\tilde{T} - T_{a,d}^{\neq}$ times. Consequently, we have that

$$\begin{aligned} C\tilde{T}^{\kappa} &\geq \tilde{\mathfrak{A}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \\ &\geq (\tilde{T} - T_{a,d}^{\neq}) \underbrace{(\max_{a' \in \mathcal{A}} v^{\text{up}}(a') - (\tau_a + v^{\text{up}}(a)))}_{\geq 1/T^\beta} \\ &\geq (\tilde{T} - T_{a,d}^{\neq}) \tilde{T}^{-\beta/\alpha}, \end{aligned}$$

which gives $T_{a,d}^{\neq} \geq \tilde{T} - C\tilde{T}^{\kappa+\beta/\alpha}$. Therefore, if we take the contrapositive, we obtain that if $T_{a,d}^{\neq} < \tilde{T} - C\tilde{T}^{\kappa+\beta/\alpha}$, then with probability at least $1 - T^{-\alpha\zeta}$, $\tau_a > \tau_a^* - 1/T^\beta$, or equivalently $\tau_a^* < \tau_a + 1/T^\beta$.

For the second part of the proof, suppose that: $C\tilde{T}^{\kappa+\beta/\alpha} < T_{a,d}^{\neq} < \tilde{T}^{\kappa+\beta/\alpha} - C\tilde{T}^{\kappa+\beta/\alpha}$.

From the above result, we have that $\tau_a^* < \tau_a + 1/T^\beta$ and $\tau_a^* > \tau_a - 1/T^\beta$ with probability at least $1 - 2T^{-\alpha\zeta}$. Plugging these inequalities in the absolute value $|\tau_a^* - \tau_a|$ concludes the proof. \square

Lemma 3. Assume that we run Algorithm 2 and consider some binary search batch iteration $d \in \lceil \log T^\beta \rceil$ run on arm $a \in \mathcal{A}$. Then $0 \leq \underline{\tau}_a(d) \leq \tau_a^{\text{mid}}(d) \leq \bar{\tau}_a(d) \leq 1$.

Proof of Lemma 3. The proof proceeds by induction. Considering some action $a \in \mathcal{A}$, we have for the initialisation before any binary search is run: $\underline{\tau}_a(0) = 0, \bar{\tau}_a(0) = 1$ and therefore $\tau_a^{\text{mid}}(0) \in [\underline{\tau}_a(0), \bar{\tau}_a(0)]$. We now consider that a number d of binary search batches has been run on a . Suppose that we run an additional binary search batch on action a . We have

$$\tau_a^{\text{mid}}(d+1) = \frac{\bar{\tau}_a(d) + \underline{\tau}_a(d)}{2} \text{ which gives } \tau_a^{\text{mid}}(d+1) \in [\underline{\tau}_a(d), \bar{\tau}_a(d)]. \quad (22)$$

After this iteration of binary search, we either have $\bar{\tau}_a(d+1) = \tau_a^{\text{mid}}(d+1) + 1/T^\beta$ and $\underline{\tau}_a(d+1) = \underline{\tau}_a(d)$ or $\bar{\tau}_a(d+1) = \tau_a^{\text{mid}}(d)$ and $\underline{\tau}_a(d+1) = \tau_a^{\text{mid}}(d+1) - 1/T^\beta$. Therefore, we still have $0 \leq \underline{\tau}_a(d+1) \leq \tau_a^{\text{mid}}(d+1) \leq \bar{\tau}_a(d+1) \leq 1$, hence the result for any $d \in \lceil \log_2 T^\beta \rceil$ by induction. \square

Lemma 4. Consider some arm $a \in \mathcal{A}$ and suppose that we have run $D \in \mathbb{N}^*$ batches of binary search of length \tilde{T} on a . Then, we have that

$$\mathbb{P} \left(\bigcap_{d \in [D]} \{ \tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq C\tilde{T}^\kappa \} \right) \geq 1 - D/T^{\alpha\zeta}.$$

Proof of Lemma 4. First observe that for any batch $d \in \{1, \dots, D\}$ of binary search run on arm a during steps $\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}$, we have by **H2** that

$$\mathbb{P} \left(\tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \geq C\tilde{T}^\kappa \right) \leq 1/\tilde{T}^\zeta,$$

and applying a union bound over the D batches, we have that

$$\begin{aligned} & \mathbb{P} \left(\bigcup_{d \in [D]} \{ \tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \geq C\tilde{T}^\kappa \} \right) \\ & \leq \sum_{j=1}^D \mathbb{P} \left(\{ \tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \geq C\tilde{T}^\kappa \} \right) \\ & \leq D/\tilde{T}^\zeta, \end{aligned}$$

which gives that

$$\mathbb{P} \left(\bigcap_{d \in [D]} \{ \tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq C\tilde{T}^\kappa \} \right) \geq 1 - D/T^{\alpha\zeta},$$

hence the result. \square

Lemma 5. Suppose that the upstream player runs a subroutine Π_p^{up} satisfying **H2**. Consider some action $a \in \mathcal{A}$ and the D -th binary search batch of length \tilde{T} run on arm a with $D \in \{1, \dots, \lceil \log_2 T^\beta \rceil\}$. Then

$$\bigcap_{d \in [D]} \{ \tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq C\tilde{T}^\kappa \} \subseteq \{ \tau_a^* \in [\underline{\tau}_a(D), \bar{\tau}_a(D)] \},$$

and the probability of these events is at least $1 - \lceil \log_2 T^\beta \rceil / T^{\alpha\zeta}$. We also have that

$$|\bar{\tau}_a(D) - \underline{\tau}_a(D)| \leq 1/2^D + 2/T^\beta \text{ holds almost surely.}$$

Proof of Lemma 5. Suppose that the conditions of the lemma hold and consider some arm a .

We show by induction on the number of binary search batches D that have been run on a that $\bigcap_{d \in [D]} \{ \tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq C\tilde{T}^\kappa \} \subseteq \{ \tau_a^* \in [\underline{\tau}_a(D), \bar{\tau}_a(D)] \}$. If it is true, Lemma 4 completes this first part of the proof.

The initialisation holds since $\underline{\tau}_a(0) = 0, \bar{\tau}_a(0) = 1$ and $\tau_a^* = \max_{a' \in \mathcal{A}} v^{\text{up}}(a') - v^{\text{up}}(a) \in [0, 1]$ with probability 1 - since $\max_{a' \in \mathcal{A}} v^{\text{up}}(a') \in [0, 1]$ and $v^{\text{up}}(a) \in [0, 1]$.

We suppose that the property is true for some integer $D < \lceil \log_2 T^\beta \rceil$ and that we have run one more binary search on arm a . We have

$$\tau_a^{\text{mid}}(D+1) = \frac{\bar{\tau}_a(D) + \underline{\tau}_a(D)}{2},$$

$\tau_a^{\text{mid}}(D+1)$ being the incentive offered to the upstream player if he chooses action a during the D -th batch $\{k_{a,D} + 1, \dots, k_{a,D} + \tilde{T}\}$. After this batch, if $T_{a,D}^\neq < \tilde{T} - C\tilde{T}^{\kappa+\beta/\alpha}$, BELGIC updates $\bar{\tau}_a(D+1) = \tau_a^{\text{mid}}(D+1) + 1/T^\beta$, $\underline{\tau}_a(D+1) = \underline{\tau}_a(D)$ and Lemma 2 ensures that $\underline{\tau}_a(D+1) < \tau_a^* < \bar{\tau}_a(D+1)$ given $\{\tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,D+1} + 1, \dots, k_{a,D+1} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq C\tilde{T}^\kappa\}$. Thus the induction holds.

Otherwise, if $T_{a,D}^\neq > C\tilde{T}^{\kappa+\beta/\alpha}$, BELGIC updates $\underline{\tau}_a(D+1) = \tau_a^{\text{mid}}(D+1) - 1/T^\beta$, $\bar{\tau}_a(D+1) = \bar{\tau}_a(D)$ and Lemma 2 ensures that $\underline{\tau}_a(D+1) < \tau_a^* < \bar{\tau}_a(D+1)$ given $\{\tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,D+1} + 1, \dots, k_{a,D+1} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq C\tilde{T}^\kappa\}$. The induction still holds.

Consequently, we have that for any number D of binary search batches run on arm a , $\tau_a^* \in [\underline{\tau}_a(D), \bar{\tau}_a(D)]$ with probability $1 - D/T^{\alpha\zeta}$

For the second part of the proof, we define $u(D) = \bar{\tau}_a(D) - \underline{\tau}_a(D) \geq 0$ as the length of the interval containing τ_a^* with probability at least $1 - D/T^{\alpha\zeta}$. We have $u(0) = 1$. Suppose that after D iterations of binary search batches, the next batch of binary search $\{k_{a,D+1} + 1, \dots, k_{a,D+1} + \tilde{T}\}$ outputs $T_{a,D+1}^\neq < \tilde{T} - C\tilde{T}^{\kappa+\beta/\alpha}$. Then, the update of Algorithm 2 gives

$$\begin{aligned} u_{D+1} &= \bar{\tau}_a(D+1) - \underline{\tau}_a(D+1) \\ &= \tau_a^{\text{mid}}(D+1) + 1/T^\beta - \underline{\tau}_a(D) \\ &= \frac{\bar{\tau}_a(D) + \underline{\tau}_a(D)}{2} - \underline{\tau}_a(D) + 1/T^\beta \\ &= \frac{\bar{\tau}_a(D) - \underline{\tau}_a(D)}{2} + 1/T^\beta \\ &= u_D/2 + 1/T^\beta. \end{aligned}$$

On the other hand, if $T_{a,D+1}^\neq > C\tilde{T}^{\kappa+\beta/\alpha}$, the update gives

$$\begin{aligned} u_{D+1} &= \bar{\tau}_a(D+1) - \underline{\tau}_a(D+1) \\ &= \bar{\tau}_a(D) - (\tau_a^{\text{mid}}(D+1) - 1/T^\beta) \\ &= \bar{\tau}_a(D) - \frac{\bar{\tau}_a(D) + \underline{\tau}_a(D)}{2} + 1/T^\beta \\ &= \frac{\bar{\tau}_a(D) - \underline{\tau}_a(D)}{2} + 1/T^\beta \\ &= u_D/2 + 1/T^\beta. \end{aligned}$$

We can see that $(u_D)_{D \geq 0}$ is an arithmetico-geometric sequence defined by $u_{D+1} = u_D/2 + 1/T^\beta$ with an initial term $u_0 = 1$. Writing $r = 1/T^\beta/(1 - 1/2) = 2/T^\beta$, we obtain that

$$|\bar{\tau}_a(D) - \underline{\tau}_a(D)| = u_D = 1/2^D(1 - r) + r = 1/2^D(1 - 2/T^\beta) + 2/T^\beta \leq 1/2^D + 2/T^\beta,$$

for any $D \in \{1, \dots, \lceil \log_2 T^\beta \rceil\}$, hence the result. \square

Lemma 6. *Suppose that the upstream player runs a policy Π_p^{up} satisfying **H2**. Considering some action $a \in \mathcal{A}$, we have that after the binary search batch $D = \lceil \log_2 T^\beta \rceil$: $\mathbb{P}(\underline{\tau}_a(D) \leq \tau_a^* \leq \bar{\tau}_a(D) \leq \underline{\tau}_a + 3/T^\beta) \geq 1 - D/T^{\alpha\zeta}$.*

Proof of Lemma 6. We suppose that the event $\bigcap_{d \in [\lceil \log_2 T^\beta \rceil]} \{\tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq C\tilde{T}^\kappa\}$ holds. Lemma 4 ensures that this event holds with probability at least $1 - \lceil \log_2 T^\beta \rceil / T^{\alpha\zeta}$.

After $D = \lceil \beta \log_2 T \rceil$ batches of binary search on arm a , we have by Lemma 5 that

$$|\bar{\tau}_a(D) - \underline{\tau}_a(D)| \leq 1/2^D + 2/T^\beta \leq 1/2^{\beta \log_2 T} + 2/T^\beta = 3/T^\beta.$$

Lemma 5 guarantees that $\tau_a^* \in [\underline{\tau}_a(D); \bar{\tau}_a(D)]$ with probability at least $1 - D/T^{\alpha\zeta}$, and we obtain $\underline{\tau}_a(D) \leq \tau_a^* \leq \bar{\tau}_a(D) \leq \underline{\tau}_a(D) + 3/T^\beta$ with the same probability. \square

Proposition 1. *Under H1 and H2, after the first phase of BELGIC which consists in $K \lceil T^\alpha \rceil \lceil \log T^\beta \rceil$ steps of binary search grouped in $\lceil \log_2 T^\beta \rceil$ batches per arm $a \in \mathcal{A}$, we have that*

$$\mathbb{P}\left(\text{for any } a \in \mathcal{A}, \hat{\tau}_a - 4/T^\beta - CT^{(\kappa-1)/2} \leq \tau_a^* \leq \hat{\tau}_a\right) \geq 1 - K \lceil \log_2 T^\beta \rceil / T^{\alpha\zeta}.$$

Proof of Proposition 1. We consider a number of binary search batches $D = \lceil \log T^\beta \rceil$ and we define the event \mathcal{G} as $\mathcal{G} = \{\text{for any } a \in \mathcal{A}, \underline{\tau}_a(D) \leq \tau_a^* \leq \bar{\tau}_a(D) \leq \underline{\tau}_a(D) + 3/T^\beta\}$. We have that

$$\begin{aligned} \mathbb{P}(\mathcal{G}) &= \mathbb{P}\left(\bigcap_{a \in \mathcal{A}} \{\underline{\tau}_a(D) \leq \tau_a^* \leq \bar{\tau}_a(D) \leq \underline{\tau}_a(D) + 3/T^\beta\}\right) \\ &= 1 - \mathbb{P}\left(\bigcup_{a \in \mathcal{A}} \{\underline{\tau}_a(D) \leq \tau_a^* \leq \bar{\tau}_a(D) \leq \underline{\tau}_a(D) + 3/T^\beta\}^c\right) \\ &\geq 1 - \sum_{a \in \mathcal{A}} \mathbb{P}\left(\{\underline{\tau}_a(D) \leq \tau_a^* \leq \bar{\tau}_a(D) \leq \underline{\tau}_a(D) + 3/T^\beta\}^c\right), \end{aligned}$$

where the last inequality holds with an union bound. Lemma 6 with $D = \lceil \log_2 T^\beta \rceil$ ensures that we have

$$\mathbb{P}\left(\{\underline{\tau}_a(D) \leq \tau_a^* \leq \bar{\tau}_a(D) \leq \underline{\tau}_a(D) + 3/T^\beta\}^c\right) \leq D/T^{\alpha\zeta},$$

and since $\text{Card}\{\mathcal{A}\} = K$, we obtain

$$\mathbb{P}(\mathcal{G}) \geq 1 - K \lceil \log_2 T^\beta \rceil / T^{\alpha\zeta}.$$

Since the estimated incentives are defined as $\hat{\tau}_a = \bar{\tau}_a(\lceil \log_2 T^\beta \rceil) + 1/T^\beta + CT^{(\kappa-1)/2}$, we can conclude

$$\mathbb{P}(\text{for any } a \in \mathcal{A}, \hat{\tau}_a - 4/T^\beta - CT^{(\kappa-1)/2} \leq \tau_a^* \leq \hat{\tau}_a) \geq 1 - K \lceil \log_2 T^\beta \rceil / T^{\alpha\zeta},$$

since whenever $\underline{\tau}_a(\lceil \log_2 T^\beta \rceil) \leq \tau_a^* \leq \bar{\tau}_a(\lceil \log_2 T^\beta \rceil) \leq \underline{\tau}_a(\lceil \log_2 T^\beta \rceil) + 1/T^\beta$, we also have by definition: $\hat{\tau}_a(\lceil \log_2 T^\beta \rceil) - 4/T^\beta - CT^{(\kappa-1)/2} \leq \underline{\tau}_a(\lceil \log_2 T^\beta \rceil) \leq \tau_a^* \leq \hat{\tau}_a(\lceil \log_2 T^\beta \rceil)$. \square

Theorem 3. *Assume that H1 and H2 hold. Then BELGIC, run with α, β satisfying (13) and any bandit subroutine *Bandit-Alg*, has an overall regret $\mathfrak{R}_p^{\text{down}}$ such that*

$$\begin{aligned} \mathfrak{R}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \text{BELGIC}) &\leq 2(3 + 2C + \bar{v} - \underline{v}) \log_2(T) (2T^{1-\alpha\zeta} + T^{(\kappa+1)/2} + \lceil T^\alpha \rceil) + 4T^{1-\beta} \\ &\quad + \mathfrak{R}_{\text{Bandit-Alg}}(T, \nu, \{\hat{\tau}_a\}_{a \in \mathcal{A}}) \end{aligned}$$

where, for ease of notation

$$\bar{v} = \max_{a, b \in \mathcal{A} \times \mathcal{A}} \{v^{\text{down}}(a, b)\} \quad \text{and} \quad \underline{v} = \min_{a, b \in \mathcal{A} \times \mathcal{A}} \{v^{\text{down}}(a, b)\}.$$

Proof of Theorem 3. Suppose that the conditions of Theorem 3 are satisfied. By definition, $\Lambda_{K+1} + 1 \in [T]$ is the step at which starts the run of the subroutine *Bandit-Alg*, since $\Lambda_{K+1} = K \lceil T^\alpha \rceil \lceil \beta \log T \rceil$. All the binary search batches have length $\tilde{T} = \lceil T^\alpha \rceil$. For any $a \in \mathcal{A}, d \in \{1, \dots, \lceil \log_2 T^\beta \rceil\}$, we define the event

$$\mathbf{B}_{a,d} = \left\{ \tilde{\mathfrak{R}}_p^{\text{up}}(\{k_{a,d} + 1, \dots, k_{a,d} + \tilde{T}\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq C\tilde{T}^\kappa \right\},$$

as well as

$$\mathcal{E} = \bigcap_{\substack{a \in [K] \\ d \in [\lceil \beta \log T \rceil]}} \mathbf{B}_{a,d} \cap \left\{ \tilde{\mathfrak{R}}_p^{\text{up}}(\{\Lambda_{K+1} + 1, \dots, T\}, \Pi_p^{\text{up}}) \leq C(T - \Lambda_{K+1})^\kappa \right\},$$

and by **H2** and Lemma 4, with an union bound, we have that

$$\begin{aligned}
\mathbb{P}(\mathcal{E}) &= 1 - \mathbb{P} \left(\bigcup_{\substack{a \in [K] \\ d \in \lceil \log T^\beta \rceil}} \mathbb{B}_{a,d}^c \bigcup \left\{ \tilde{\mathfrak{R}}_p^{\text{up}}(\{\Lambda_{K+1} + 1, \dots, T\}, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq C(T - \Lambda_{K+1})^\kappa \right\} \right) \\
&\geq 1 - \sum_{\substack{a \in [K] \\ d \in \lceil \log T^\beta \rceil}} \mathbb{P}(\mathbb{B}_{a,d}^c) + \mathbb{P}(\tilde{\mathfrak{R}}_p^{\text{up}}(\{\Lambda_{K+1} + 1, \dots, T\}, \Pi_p^{\text{up}}) \leq C(T - \Lambda_{K+1})^\kappa) \\
&\geq 1 - K \lceil \beta \log T \rceil \tilde{T}^{-\zeta} - (T - \Lambda_{K+1})^{-\zeta} \\
&\geq 1 - K \lceil \beta \log T \rceil T^{-\alpha\zeta} - T^{-\zeta},
\end{aligned}$$

and we now decompose

$$\mathfrak{R}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) = \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \tilde{\mathfrak{R}}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \right] + \left[\mathbb{1}(\mathcal{E}^c) \tilde{\mathfrak{R}}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \right], \quad (23)$$

where $\tilde{\mathfrak{R}}_p^{\text{down}}$ is defined as

$$\tilde{\mathfrak{R}}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) = T\mu^{*,\text{down}} - \sum_{t=1}^T (v^{\text{down}}(A_t, B_t) - \mathbb{1}_{\hat{a}_t}(A_t)\tau(t)).$$

By definition, $\mu^{*,\text{down}} \geq \min_{a,b \in \mathcal{A} \times \mathcal{A}} v^{\text{down}}(a, b) \geq \underline{v}$ and since Lemma 1 allows to write $\mu^{*,\text{down}} = \max_{a,b \in \mathcal{A} \times \mathcal{A}} \{v^{\text{down}}(a, b) + v^{\text{up}}(a)\} - \max_{a' \in \mathcal{A}} \{v^{\text{up}}(a')\}$, we have that

$$\underline{v} \leq \mu^{*,\text{down}} \leq 1 + \bar{v},$$

and note that since $\kappa < 1$, for any $a \in \mathcal{A}$, $\hat{\tau}_a \leq 2 + CT^{(\kappa-1)/2} \leq 2 + C$. Consequently, $T\underline{v} \leq \tilde{\mathfrak{R}}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \leq (3 + C + \bar{v} - \underline{v})T$ almost surely. Therefore

$$\mathbb{E} \left[\mathbb{1}(\mathcal{E}^c) \tilde{\mathfrak{R}}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \right] \leq (3 + C + \bar{v} - \underline{v})(K \lceil \log T^\beta \rceil T^{1-\alpha\zeta} + T^{1-\zeta}). \quad (24)$$

We consider the second term in (23). We decompose it between the steps of binary search $\{1, \dots, K \lceil T^\alpha \rceil \lceil \log_2 T^\beta \rceil\}$ during which we run the Binary Search Subroutine and the following ones when we run Bandit-Alg, which gives

$$\begin{aligned}
&\mathbb{E} \left[\mathbb{1}(\mathcal{E}) \tilde{\mathfrak{R}}_p^{\text{down}}(T, \Pi_p^{\text{up}}, \Pi_p^{\text{down}}) \right] \quad (25) \\
&\leq \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \underbrace{\sum_{t=1}^{K \lceil T^\alpha \rceil \lceil \log_2 T^\beta \rceil} \mu^{*,\text{down}} - (v^{\text{down}}(A_t, B_t) - \mathbb{1}_{\hat{a}_t}(A_t)\tau(t))}_{\text{(A)}} \right] \\
&+ \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \underbrace{\sum_{t=K \lceil T^\alpha \rceil \lceil \log_2 T^\beta \rceil + 1}^T \mu^{*,\text{down}} - (v^{\text{down}}(A_t, B_t) - \mathbb{1}_{\hat{a}_t}(A_t)\tau(t))}_{\text{(B)}} \right].
\end{aligned}$$

Similarly to (24), we use the bound on $\mu^{*,\text{down}}$ to bound (A), which gives

$$\mathbb{E}[\mathbb{1}(\mathcal{E})(\text{A})] \leq \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \sum_{t=1}^{\Lambda_{K+1}} 1 + \bar{v} - (\underline{v} - 2 - C) \right] \leq K \lceil T^\alpha \rceil (1 + \log_2 T) (3 + C + \bar{v} - \underline{v}). \quad (26)$$

After the binary search, at each step $t \in \{\Lambda_{K+1} + 1, \dots, T\}$, **Bandit-Alg** recommends $(\tilde{a}_t, B_t) \in \mathcal{A} \times \mathcal{A}$ following (15) and **BELGIC** offers an incentive $(\tilde{a}_t, \hat{\tau}_{\tilde{a}_t})$ with $\hat{\tau}_a = \bar{\tau}_a(\lceil \log_2 T^\beta \rceil) + 1/T^\beta + CT^{(\kappa-1)/2}$.

Lemma 5 ensures that $\mathcal{E} \subseteq \{\tau_a^* \in [\underline{\tau}_a(\lceil \log_2 T^\beta \rceil), \bar{\tau}_a(\lceil \log_2 T^\beta \rceil)] \text{ for any } a \in \mathcal{A}\} \cap \{\mathfrak{R}_p^{\text{up}}(\{\Lambda_{K+1} + 1, \dots, T\}, \Pi_p^{\text{up}}) \leq C(T - \Lambda_{K+1})^\kappa\}$. Therefore, if \mathcal{E} holds, for any $a \in \mathcal{A}$, we have that

$$\begin{aligned} v^{\text{up}}(a) + \hat{\tau}_a &= v^{\text{up}}(a) + \bar{\tau}_a + 1/T^\beta + CT^{(\kappa-1)/2} \\ &> v^{\text{up}}(a) + \tau_a^* + CT^{(\kappa-1)/2} \\ &= v^{\text{up}}(a) + \max_{a'' \in \mathcal{A}} v^{\text{up}}(a'') - v^{\text{up}}(a) + CT^{(\kappa-1)/2}, \end{aligned}$$

and therefore, for any $a' \in \mathcal{A}$ such that $a' \neq a$, we have that

$$v^{\text{up}}(a) + \hat{\tau}_a > v^{\text{up}}(a') + CT^{(\kappa-1)/2}. \quad (27)$$

This shows that at any steps $t \in \{\Lambda_{K+1} + 1, \dots, T\}$ after the binary search, we have on the event \mathcal{E} that

$$\tilde{a}_t = \operatorname{argmax}_{a' \in \mathcal{A}} \{v^{\text{up}}(a') + \mathbb{1}_{\tilde{a}_t}(a') \hat{\tau}_{\tilde{a}_t}\}, \quad (28)$$

and the reward gap at step t for any $a \neq \tilde{a}_t$ is defined as

$$\max_{a' \in \mathcal{A}} \{v^{\text{up}}(a') + \mathbb{1}_{\tilde{a}_t}(a') \hat{\tau}_{\tilde{a}_t}\} - (v^{\text{up}}(a) + \mathbb{1}_{\tilde{a}_t}(a) \hat{\tau}_{\tilde{a}_t}) = v^{\text{up}}(\tilde{a}_t) + \hat{\tau}_a - v^{\text{up}}(a). \quad (29)$$

Following (27), the reward gap from (29) satisfies

$$\max_{a' \in \mathcal{A}} \{v^{\text{up}}(a') + \mathbb{1}_{\tilde{a}_t}(a') \hat{\tau}_{\tilde{a}_t}\} - (v^{\text{up}}(a) + \mathbb{1}_{\tilde{a}_t}(a) \hat{\tau}_{\tilde{a}_t}) \geq CT^{(\kappa-1)/2}.$$

We now define two sets

$$\begin{aligned} I_T &= \{t \in \{K \lceil T^\alpha \rceil \lceil \log_2 T \rceil + 1, \dots, T\} \text{ such that } \tilde{a}_t = A_t\}, \\ J_T &= \{t \in \{K \lceil T^\alpha \rceil \lceil \log_2 T \rceil + 1, \dots, T\} \text{ such that } \tilde{a}_t \neq A_t\}, \end{aligned}$$

which satisfy $I_T \cup J_T = \{K \lceil T^\alpha \rceil \lceil \log_2 T \rceil + 1, \dots, T\}$ almost surely. As shown in (28), I_T corresponds to all the steps during which the upstream player picked the best arm and for any $t \in I_T$

$$v^{\text{up}}(A_t) + \mathbb{1}_{\tilde{a}_t}(A_t) \tau(t) \geq \max_{a \in \mathcal{A}} \{v^{\text{up}}(a) + \mathbb{1}_{\tilde{a}_t}(a) \tau(t)\},$$

while by (29), for any $t \in J_T$, we have that

$$\max_{a \in \mathcal{A}} \{v^{\text{up}}(a) + \mathbb{1}_{\tilde{a}_t}(a) \hat{\tau}_{\tilde{a}_t}\} - (v^{\text{up}}(A_t) + \mathbb{1}_{\tilde{a}_t}(A_t) \hat{\tau}_{\tilde{a}_t}) \geq CT^{(\kappa-1)/2}. \quad (30)$$

H2 ensures that if \mathcal{E} holds, then $\mathfrak{R}_p^{\text{up}}(\{\Lambda_{K+1} + 1, \dots, T\}, \Pi_p^{\text{up}}, \text{BELGIC}) \leq CT^\kappa$, and this condition together with (30) gives that

$$\text{Card}\{J_T\} CT^{(\kappa-1)/2} \leq \mathfrak{R}_p^{\text{up}}(\{\Lambda_{K+1} + 1, \dots, T\}, \Pi_p^{\text{up}}, \text{BELGIC}) \leq CT^\kappa,$$

and consequently $\text{Card}\{J_T\} \leq T^{(\kappa+1)/2}$. We now bound (\mathbf{B}) as follows

$$\begin{aligned}
\mathbb{E}[\mathbb{1}(\mathcal{E})(\mathbf{B})] &= \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \sum_{t=\Lambda_{\kappa+1}+1}^T \mu^{\star, \text{down}} - (v^{\text{down}}(A_t, B_t) - \hat{\tau}_{\tilde{a}_t}) \right] \\
&= \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \sum_{t \in I_T} \max_{a, b \in \mathcal{A} \times \mathcal{A}} \{v^{\text{down}}(a, b) - \tau_a^*\} - (v^{\text{down}}(\tilde{a}_t, B_t) - \hat{\tau}_{\tilde{a}_t}) \right] \\
&\quad + \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \sum_{t \in J_T} \underbrace{\mu^{\star, \text{down}} - (v^{\text{down}}(A_t, B_t) - \hat{\tau}_{\tilde{a}_t})}_{\leq 3+C+\bar{v}-\underline{v}} \right] \\
&\leq \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \sum_{t \in I_T} \max_{a, b \in \mathcal{A} \times \mathcal{A}} \{v^{\text{down}}(a, b) - \hat{\tau}_a - (v^{\text{down}}(\tilde{a}_t, B_t) - \hat{\tau}_{\tilde{a}_t})\} + \max_{a' \in \mathcal{A}} \{\hat{\tau}_{a'} - \tau_{a'}^*\} \right] \\
&\quad + (3 + C + \bar{v} - \underline{v}) \mathbb{E}[\mathbb{1}(\mathcal{E}) \text{Card}\{J_T\}] \\
&= \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \sum_{t \in I_T} \max_{a, b \in \mathcal{A} \times \mathcal{A}} \{v^{\text{down}}(a, b) - \hat{\tau}_a\} - (v^{\text{down}}(\tilde{a}_t, B_t) - \hat{\tau}_{\tilde{a}_t}) \right] \\
&\quad + \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \sum_{t \in I_T} \max_{a' \in \mathcal{A}} \{\hat{\tau}_{a'} - \tau_{a'}^*\} \right] + (3 + C + \bar{v} - \underline{v}) T^{(\kappa+1)/2} \\
&\leq \mathfrak{R}_{\text{Bandit-Alg}}(\text{Card}\{I_T\}, \nu, \{\hat{\tau}_a\}_{a \in \mathcal{A}}) + \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \text{Card}\{I_T\} \max_{a' \in \mathcal{A}} \{\hat{\tau}_{a'} - \tau_{a'}^*\} \right] \\
&\quad + (3 + C + \bar{v} - \underline{v}) T^{(\kappa+1)/2},
\end{aligned}$$

where the first step holds by Lemma 1. Using Lemma 5, as well as the definition of \mathcal{E} , we obtain

$$\begin{aligned}
\mathbb{E} \left[\mathbb{1}(\mathcal{E}) \text{Card}\{I_T\} \max_{a' \in \mathcal{A}} \{\hat{\tau}_{a'} - \tau_{a'}^*\} \right] &\leq \mathbb{E} \left[\mathbb{1}(\mathcal{E}) \text{Card}\{\mathcal{E}\} (4/T^\beta + CT^{(\kappa-1)/2}) \right] \\
&\leq \mathbb{E}[\mathbb{1}(\mathcal{E}) T] (4/T^\beta + CT^{(\kappa-1)/2}) \\
&\leq 4T^{1-\beta} + CT^{(\kappa+1)/2},
\end{aligned}$$

which finally gives

$$\mathbb{E}[\mathbb{1}(\mathcal{E})(\mathbf{B})] \leq \mathfrak{R}_{\text{Bandit-Alg}}(T, \nu, \{\hat{\tau}_a\}_{a \in \mathcal{A}}) + 4T^{1-\beta} + (3 + 2C + \bar{v} - \underline{v}) T^{(1+\kappa)/2}, \quad (31)$$

and plugging together (26) and (31) in the decomposition (25) gives the following bound

$$\begin{aligned}
\mathbb{E} \left[\mathbb{1}(\mathcal{E}) \tilde{\mathfrak{R}}_{\text{p}}^{\text{down}}(T, \Pi_{\text{p}}^{\text{up}}, \text{BELGIC}) \right] &\leq \mathfrak{R}_{\text{Bandit-Alg}}(T, \nu, \{\hat{\tau}_a\}_{a \in \mathcal{A}}) + 4T^{1-\beta} + (3 + 2C + \bar{v} - \underline{v}) T^{(1+\kappa)/2} \\
&\quad + (3 + C + \bar{v} - \underline{v}) K \lceil T^\alpha \rceil (1 + \log_2 T),
\end{aligned}$$

and summing the bounds on the events \mathcal{E} and \mathcal{E}^c finally gives

$$\begin{aligned}
\mathfrak{R}_{\text{p}}^{\text{down}}(T, \Pi_{\text{p}}^{\text{up}}, \text{BELGIC}) &\leq \mathfrak{R}_{\text{Bandit-Alg}}(T, \nu, \{\hat{\tau}_a\}_{a \in \mathcal{A}}) + 4T^{1-\beta} + (3 + 2C + \bar{v} - \underline{v}) T^{(1+\kappa)/2} \\
&\quad + (2 + C + \bar{v} - \underline{v}) K \lceil T^\alpha \rceil (1 + \log_2 T) \\
&\quad + (3 + C + \bar{v} - \underline{v}) (K(1 + \log_2 T) T^{1-\alpha\zeta} + T^{1-\zeta}) \\
&\leq \mathfrak{R}_{\text{Bandit-Alg}}(T, \nu, \{\hat{\tau}_a\}_{a \in \mathcal{A}}) + 4T^{1-\beta} + (3 + 2C + \bar{v} - \underline{v}) T^{(1+\kappa)/2} \\
&\quad + (3 + \bar{C} + \bar{v} - \underline{v}) ((1 + \log_2 T) (\lceil T^\alpha \rceil + T^{1-\alpha\zeta}) + T^{1-\zeta}) \\
&\leq (3 + 2C + \bar{v} - \underline{v}) (T^{1-\zeta} + T^{(\kappa+1)/2}) + (1 + \log_2 T) (\lceil T^\alpha \rceil + T^{1-\alpha\zeta}) \\
&\quad + \mathfrak{R}_{\text{Bandit-Alg}}(T, \nu, \{\hat{\tau}_a\}_{a \in \mathcal{A}}) + 4T^{1-\beta} \\
&\leq 2(3 + 2C + \bar{v} - \underline{v}) \log_2(T) (2T^{1-\alpha\zeta} + T^{(\kappa+1)/2}) + \lceil T^\alpha \rceil + 4T^{1-\beta} \\
&\quad + \mathfrak{R}_{\text{Bandit-Alg}}(T, \nu, \{\hat{\tau}_a\}_{a \in \mathcal{A}}).
\end{aligned}$$

□

Corollary 1. *Assume that the upstream player's distribution $(\gamma_a)_{a \in \mathcal{A}}$ is such that **H 1** holds. In addition, suppose that the distributions $(\gamma_a)_{a \in \mathcal{A}}$ and $(\nu_{a,b})_{a,b \in \mathcal{A} \times \mathcal{A}}$ are 1-sub-Gaussian and that the upstream player plays $\Pi_p^{\text{up}} = \text{Algorithm 3}$ (a slight modification of UCB to take into account the incentives). Then the downstream player's regret when she runs BELGIC with parameters $\alpha = 3/4$ and $\beta = 1/4$ (which satisfy (13)) and subroutine $\text{Bandit-Alg} = \text{UCB}$ satisfies the following upper bound⁴*

$$\begin{aligned} \mathfrak{R}_p^{\text{down}}(T, \text{UCB}, \text{BELGIC}) &\leq (10 + 4K + 32\sqrt{K \log_2(KT^3)} + \bar{v} - \underline{v}) \log_2(T)(3 + 2T^{3/4}) \\ &\quad + 3K^2(\bar{v} - \underline{v}). \end{aligned}$$

Proof of Corollary 1. First note that $\Pi_p^{\text{up}} = \text{Algorithm 3}$ satisfies **H2** with constants $\kappa = 1/2$, $\zeta = 2$, $C = 8\sqrt{K \log(KT^3)}$, following Proposition 2. Note that $\beta/\alpha = 1/3 < 1/2 = 1 - \kappa$, therefore Equation (13) is satisfied. Plugging these terms in the bound from Theorem 3 with $\alpha = 3/4, \beta = 1/4$ gives

$$\begin{aligned} \mathfrak{R}_p^{\text{down}}(T, \text{UCB}, \text{BELGIC}) &\leq 2(3 + 16\sqrt{K \log_2(KT^3)} + \bar{v} - \underline{v}) \log_2(T)(2T^{1-3/2} + T^{3/4} + \lceil T^{3/4} \rceil) \\ &\quad + 4T^{1/4} + 8\sqrt{K^2 \log_2(T)} T^{1/2} + 3K^2(\bar{v} - \underline{v}) \end{aligned}$$

where we use the bound for $\mathfrak{R}_{\text{Bandit-Alg}}(T, \nu, \{\hat{\tau}_a\}_{a \in \mathcal{A}})$ with $\text{Bandit-Alg} = \text{UCB}$ run on any bandit instance with K^2 arms, 1-subgaussian rewards and reward gaps of at most $\max_{a,b \in \mathcal{A} \times \mathcal{A}} v^{\text{down}}(a,b) - \min_{a,b \in \mathcal{A} \times \mathcal{A}} v^{\text{down}}(a,b) = \bar{v} - \underline{v}$, following [Lattimore and Szepesvári, 2020, Theorem 7.2]. Therefore, we have that

$$\begin{aligned} \mathfrak{R}_p^{\text{down}}(T, \text{UCB}, \text{BELGIC}) &\leq (6 + 32\sqrt{K \log_2(KT^3)} + \bar{v} - \underline{v}) \log_2(T)(2 + 1 + 2T^{3/4}) \\ &\quad + 4T^{1/4} + 3K^2(\bar{v} - \underline{v}) + 8K\sqrt{\log_2 T} T^{1/2} \\ &\leq (10 + 32\sqrt{K \log_2(KT^3)} + \bar{v} - \underline{v}) \log_2(T)(3 + 2T^{3/4}) \\ &\quad + 3K^2(\bar{v} - \underline{v}) + 8K\sqrt{\log_2(T)} T^{1/2}, \end{aligned}$$

which finally gives

$$\begin{aligned} \mathfrak{R}_p^{\text{down}}(T, \text{UCB}, \text{BELGIC}) &\leq (10 + 4K + 32\sqrt{K \log_2(KT^3)} + \bar{v} - \underline{v}) \log_2(T)(3 + 2T^{3/4}) \\ &\quad + 3K^2(\bar{v} - \underline{v}), \end{aligned}$$

hence the result. □

⁴Note that it is K and not \sqrt{K} here, since the action space is of cardinality K^2 for the downstream player.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract mostly presents the issue that we tackle in the paper, namely presenting an online version of the *Coase theorem*. This contribution is the most important part of our paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Our work presents a model to explore an online version of a theory from economics. Thereby, it implicitly has limitations due to the fact that a choice was made in the model.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Our work presents a lot of theorems supported by assumptions (see **H2**, Theorems **2** and **3**...). All of Appendix **C** is here to provide a theoretical support to these results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: Since our work is mostly a theoretic contribution, we do not present experiments. Therefore, our paper is not concerned by this issue.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: Since our work is mostly a theoretic contribution, we do not present experiments. Therefore, our paper is not concerned by this issue (same answer as for item 5).

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.

- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: Since our work is mostly a theoretic contribution, we do not present experiments. Therefore, our paper is not concerned by this issue (same answer as for item 5).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: Since our work is mostly a theoretic contribution, we do not present experiments. Therefore, our paper is not concerned by this issue (same answer as for item 5).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: Since our work is mostly a theoretic contribution, we do not present experiments. Therefore, our paper is not concerned by this issue (same answer as for item 5).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: As explained in the Related Works, an issue with externalities is the harm caused to all the parties involved in the considered setting. Here, we try to provide a setup so the players can negotiate and achieve a social welfare optimal equilibrium. Therefore, we are confident in the fact that we follow the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We consider an issue arising in several real-world situations and provide an algorithmic solution to it. We realize that this theoretical contribution still needs some work to be implemented but we hope that it will contribute to positive social impacts in the future.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.

- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [Yes]

Justification: Our work is a theoretical contribution. Thereby, it does not propose a release of data or any kind of trained model. We mix learning and theories from economics, which does not involve yet risks for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: We do not use any assets requiring such conditions here.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not present such kind of new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We do not incorporate any experiment involving human subjects. Therefore, we are not concerned by this item.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We do not incorporate any experiment involving human subjects, hence our answer.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.