

RESEARCH CENTRE

**Inria Saclay Centre at Institut
Polytechnique de Paris**

IN PARTNERSHIP WITH:

Institut Polytechnique de Paris, Criteo

2024

ACTIVITY REPORT

Project-Team

FAIRPLAY

Coopetitive AI: Fairness, Privacy, Incentives

IN COLLABORATION WITH: Centre de Recherche en Economie et
Statistique

DOMAIN

**Applied Mathematics, Computation and
Simulation**

THEME

**Optimization, machine learning and
statistical methods**

Inria

Contents

Project-Team FAIRPLAY	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
2.1 Broad context	3
2.1.1 Multi-agent systems	4
2.1.2 Societal aspects and ethics	4
2.2 A large variety of tools and techniques	5
2.3 General objectives	7
3 Research program	8
3.1 Theme 1: Developing fair mechanisms for auctions and matching problems	8
3.1.1 Fairness in auction-based systems	8
3.1.2 Fairness in matching and selection problems	9
3.2 Theme 2: Reconciling, and enforcing privacy with fairness	10
3.3 Theme 3: Exploiting structure in online algorithms and learning problems	11
3.3.1 Leveraging structure in online matching	11
3.3.2 Exploiting side-information in sequential learning	12
4 Application domains	13
4.1 Typical problems and use-cases	13
4.2 Application areas	14
4.2.1 Online advertisement	14
4.2.2 “Physical matching”	14
5 Highlights of the year	15
5.1 Awards	15
6 New results	15
6.1 Auctions and mechanism design	15
6.1.1 Auctions	15
6.1.2 Dynamic pricing	16
6.1.3 Mechanism design	16
6.2 Matching	17
6.3 Learning	17
6.3.1 Data valuation	17
6.3.2 Bandits and reinforcement learning	18
6.4 Online (and offline) algorithms	18
6.4.1 Prophet inequalities	18
6.4.2 Algorithms with prediction and scheduling	19
7 Partnerships and cooperations	20
7.1 International research visitors	20
7.1.1 Visits of international scientists	20
7.1.2 Visits to international teams	20
7.2 National initiatives	21
8 Dissemination	22
8.1 Promoting scientific activities	22
8.1.1 Scientific events: organisation	22
8.1.2 Scientific events: selection	22
8.1.3 Journal	23
8.1.4 Invited talks	23
8.1.5 Scientific expertise	24

8.2	Teaching - Supervision - Juries	24
8.2.1	Supervision	24
8.3	Teaching	24
8.4	Popularization	25
8.4.1	Specific official responsibilities in science outreach structures	25
9	Scientific production	25
9.1	Major publications	25
9.2	Publications of the year	25
9.3	Cited publications	28

Project-Team FAIRPLAY

Creation of the Project-Team: 2022 March 01

Keywords

Computer sciences and digital sciences

A4.8. – Privacy-enhancing technologies

A8.11. – Game Theory

A9.2. – Machine learning

A9.9. – Distributed AI, Multi-agent

Other research topics and application domains

B9.9. – Ethics

B9.10. – Privacy

1 Team members, visitors, external collaborators

Research Scientists

- Patrick Loiseau [Team leader, INRIA, Researcher, until Sep 2024, HDR]
- Patrick Loiseau [Team leader, INRIA, Senior Researcher, from Oct 2024, HDR]
- Marc Abeille [Criteo, Industrial member]
- Benjamin Heymann [Criteo, Industrial member]
- Simon Mauras [INRIA, Researcher, from Feb 2024]
- Hugo Richard [Criteo, Industrial member]
- Maxime Vono [Criteo, Industrial member]

Faculty Members

- Vianney Perchet [Team leader, Criteo & ENSAE, Professor, HDR]
- Cristina Butucea [ENSAE, Professor, HDR]
- Matthieu Lerasle [ENSAE, Professor, HDR]

Post-Doctoral Fellows

- Dorian Baudry [ENSAE, Post-Doctoral Fellow, until Jun 2024]
- Lorenzo Croissant [ENSAE]
- Simon Finster [INRIA, Post-Doctoral Fellow]
- Solenne Gaucher [ENSAE, Post-Doctoral Fellow]
- Nadav Merlis [ENSAE]
- Denis Sokolov [INRIA, Post-Doctoral Fellow]
- Bartholome Vieille [INRIA, Post-Doctoral Fellow, from Oct 2024]

PhD Students

- Ziyad Benomar [ENSAE]
- Maria Cherifa [Criteo, CIFRE]
- Hafedh El Ferichichi [ENSEA]
- Côme Fiegel [ENSAE]
- Reda Jalal [INRIA, from Feb 2024]
- Mike Liu [ENSAE]
- Mathieu Molina [INRIA]
- Giovanni Montanari [INRIA]
- Marius Potfer [EDF, CIFRE]
- Melissa Tamine [Criteo, CIFRE]

- Matilde Tullii [ENSAE]
- Rémi Verroneau [Footvision, CIFRE]
- Axel Xerri [INRIA, from Dec 2024]
- Minrui Xu [ENSAE, from Sep 2024]

Technical Staff

- Felipe Garrido Lucero [INRIA, Engineer, until Aug 2024]

Interns and Apprentices

- Yeniffer Munoz Fuentez [INRIA, Intern, until Mar 2024]

Administrative Assistant

- Melanie Da Silva [INRIA]

External Collaborators

- Clément Calauzènes [Criteo]
- Julien Combe [Ecole Polytechnique, HDR]
- Azadeh Khaleghi [ENSAE]
- Laurent Menard [UNIV PARIS NANTERRE, until Jun 2024, HDR]
- Jaouad Mourtada [ENSAE]

2 Overall objectives

2.1 Broad context

One of the principal objectives of Machine Learning (ML) is to automatically discover using past data some underlying structure behind a data generating process in order either to explain past observations or, perhaps more importantly, to make predictions and/or to optimize decisions made on future instances. The area of ML has exploded over the past decade and has had a tremendous impact in many application domains such as computer vision or bioinformatics.

Most of the current ML literature focuses on the case of a single agent (an algorithm) trying to complete some learning task based on gathered data that follows an exogenous distribution independent of the algorithm. One of the key assumptions is that this data has sufficient “regularity” for classical techniques to work. This classical paradigm of “a single agent learning on nice data”, however, is no longer adequate for many practical and crucial tasks that imply users (who own the gathered data) and/or other (learning) agents that are also trying to optimize their own objectives simultaneously, in a competitive or conflicting way. This is the case, for instance, in most learning tasks related to Internet applications (content recommendation/ranking, ad auctions, fraud detection, etc.). Moreover, as such learning tasks rely on users’ personal data and as their outcome affect users in return, it is no longer sufficient to focus on optimizing prediction performance metrics—it becomes crucial to consider societal and ethical aspects such as fairness or privacy.

The field of single agent ML builds on techniques from domains such as statistics, optimization, or functional analysis. When different agents are involved, a strategic aspect inherent in game theory enters the picture. Indeed, interactions—either positive or negative—between rational entities (firms, single user at home, algorithms, etc.) foster individual strategic behavior such as hiding information, misleading other agents, free-riding, etc. Unfortunately, this selfishness degrades the quality of the data or

of the predictions, prevents efficient learning and overall may diminish the social welfare. These strategic aspects, together with the decentralized nature of decision making in a multi-agent environment, also make it harder to build algorithms that meet fairness and privacy constraints.

The overarching objective of FAIRPLAY is to **create algorithms that learn for and with users—and techniques to analyze them—**, that is to create procedures able to perform classical learning tasks (prediction, decision, explanation) when the data is generated or provided by strategic agents, possibly in the presence of other competing learning agents, while respecting the fairness and privacy of the involved users. To that end, we will naturally rely on multi-agent models where the different agents may be either agents generating or providing data, or agents learning in a way that interacts with other agents; and we will put a special focus on societal and ethical aspects, in particular fairness and privacy. Note that in FAIRPLAY, we focus on the technical challenges inherent to formalizing mathematically and respecting ethical properties such as non-discrimination or privacy, often seen as constraints in the learning procedure. Nevertheless, throughout the team’s life, we will reflect on these mathematical definitions for the particular applications studied, in particular their philosophical roots and legal interpretation, through interactions with HSS researchers and with legal specialists (from Criteo).

2.1.1 Multi-agent systems

Any company developing and implementing ML algorithms is in fact one agent within a large network of users and other firms. Assuming that the data is i.i.d. and can be treated irrespectively of the environment response—as is done in the classical ML paradigm—might be a good first approximation, but should be overcome. Users, clients, suppliers, and competitors are adaptive and change their behavior depending on each other’s interactions. The future of many ML companies—such as Criteo—will consist in creating platforms matching the demand (created by their users) to the offer (proposed by their clients), under the system constraints (imposed by suppliers and competitors). Each of these agents have different, conflicting interests that should be taken into account in the model, which naturally becomes a multi-agent model.

Each agent in a multi-agent system may be modeled as having their own utility function u_i that can depend on the action of other agents. Then, there are two main types of objectives: individual or collective [105]. If each agent is making their own decision, then they can be modeled as each optimizing their own individual utility (which may include personal benefit as well as other considerations such as altruism where appropriate) unilaterally and in a decentralized way. This is why a mechanism providing correct incentives to agents is often necessary. At the other extreme, social welfare is the collective objective defined as the cumulative sum of utilities of all agents. To optimize it, it is almost always necessary to consider a centralized optimization or learning protocol. A key question in multi-agent systems is to apprehend the “social cost” of letting agents optimize their own utility by choosing unilaterally their decision compared to the one maximizing social welfare; this is often measured by the “price of anarchy”/“price of stability” [114]: the ratio of the maximum social welfare to the (worst/best) social welfare when agents optimize individually.

The natural language to model and study multi-agent systems is *game theory*—see below for a list of tools and techniques on which FAIRPLAY relies, game theory being the first of them. Multi-agent systems have been studied in the past; but not with a focus on learning systems where agents are either learning or providing data, which is our focus in FAIRPLAY and leads to a blend of game theory and learning techniques. We note here again that, wherever appropriate, we shall reflect (in part together with colleagues from HSS) on the soundness of the utility framework for the considered applications.

2.1.2 Societal aspects and ethics

There are several important ethical aspects that must be investigated in multi-agent systems involving users either as data providers or as individuals affected by the ML agent decision (or both).

Fairness and Discrimination When ML decisions directly affect humans, it is important to ensure that they do not violate fairness principles, be they based on ethical or legal grounds. As ML made its way in many areas of decision making, it was unfortunately repeatedly observed that it can lead to discrimination (regardless of whether or not it is intentional) based on gender, race, age, or other sensitive

attributes. This was observed in online targeted advertisement [99, 125, 39, 83, 99, 41], but also in many other applications such as hiring [70], data-driven healthcare [77], or justice [100]. Biases also have the unfortunate tendency to reinforce. An operating multi-agent learning system should be able in the long run to get rid by itself of inherent population biases, that is, be fair amongst users irrespective of the improperly constructed dataset.

The mathematical formulation of fairness has been debated in recent works. Although a few initial works proposed a notion of *individual fairness*, which mandates that “similar individuals” receive “similar outcomes” [72], this notion was quickly found unpractical because it relies on a metric to define closeness that makes the definition somewhat arbitrary. Most of the works then focused on notions of *group fairness*, which mandate equality of outcome “on average” across different groups defined by *sensitive attributes* (e.g., race, gender, religious belief, etc.). Most of the works on group fairness focus on the classification problem (e.g., classifying whether a job applicant is good or not for the job) where each data example (X_i, Y_i) contains a set of features X_i and a true label $Y_i \in \{0, 1\}$ and the goal is to make a prediction \hat{Y}_i based on the features X_i that has a high probability to be equal to the true label. Assuming that there is a single sensitive attribute s_i that can take two values a or b , this defines two groups: those for whom $s_i = a$ and those for whom $s_i = b$. There are several different concepts of group fairness that can be considered; we shall especially focus on *demographic parity* (DP), which prescribes $P(\hat{Y}_i = 1 | s_i = a) = P(\hat{Y}_i = 1 | s_i = b)$ and *equal opportunity* (EO) [84], which mandates that $P(\hat{Y}_i = 1 | s_i = a, Y_i = 1) = P(\hat{Y}_i = 1 | s_i = b, Y_i = 1)$.

The fair classification literature proposed, for each of these fairness notions, ways to train fair classifiers based on three main ideas: pre-processing [133], in-processing [131, 132, 128], and post-processing [84]. All of these works, however, focus on idealized situations where a single decision-maker has access to ground truth data with the sensitive features and labels in order to train classifiers that respect fairness constraints. We use similar group fairness definitions and extend them (in particular through causality), but our goal is to go further in terms of algorithms by modeling practical scenarios with multiple decision-makers and incomplete information (in particular lack of ground truth on the labels).

Privacy vs. Incentives ML algorithms, in particular in Internet applications, often rely on users’ personal information (whether it is directly their personal data or indirectly some hidden “type” – gender, ethnicity, behaviors, etc.). Nevertheless, users may be willing to provide their personal information if it increases their utility. This brings a number of key questions. First, how can we learn while protecting users’ privacy (and how should privacy even be defined)? Second, finding the right balance between those two a-priori incompatible concepts is challenging; how much (and even simply how) should an agent be compensated for providing useful and accurate data?

Differential privacy is the most widely used private learning framework [71, 73, 120] and ensures that the output of an algorithm does not significantly depend on a single element of the whole dataset. These privacy constraints are often too strong for economic applications (as illustrated before, it is sometimes optimal to disclose some private information). f -divergence privacy costs have thus been proposed in recent literature as a promising alternative [63]. These f -divergences, such as Kullback-Leibler, are also used by economists to measure the cost of information from a Bayesian perspective, as in the rational inattention literature [124, 107, 102]. It was only recently that this approach was considered to measure “privacy losses” in economic mechanisms [74]. In this model, the mechanism designer has some prior belief on the unobserved and private information. After observing the player’s action, this belief is updated and the cost of information corresponds to the KL between the prior and posterior distributions of this private information.

This privacy concept can be refined up to a single user level, into the so-called local differential privacy. Informally speaking, the algorithm output can also depend on a single user data that still must be kept private. Estimation are actually sometimes more challenging under this constraint, i.e., estimation rates degrade [121, 58, 59] but is sometimes more adapted to handle user-generated data [79].

Interestingly, we note that the notions of privacy and fairness are somewhat incompatible. This will motivate Theme 2 developed in our research program.

2.2 A large variety of tools and techniques

Analyzing multi-agent learning systems with ethical constraints will require us to use, develop, and merge several different theoretical tools and techniques. We describe the main ones here. Note that although

FAIRPLAY is motivated by practical use-cases and applications, part of the team's objectives is to improve those tools as necessary to tackle the problems studied.

Game theory and economics Game theory [78] is the natural mathematical tool to model multiple interacting decision-makers (called players). A game is defined by a set of players, a set of possible actions for each player, and a payoff function for each player that can depend on the actions of all the players (that is the distinguishing feature of a game compared to an optimization problem). The most standard solution concept is the so-called Nash equilibrium, which is defined as a strategy profile (i.e., a collection of possibly randomized action for each player) such that each player is at best response (i.e., has the maximum payoff given the others' strategies). It is a "static" (one-shot) solution concept, but there also exist dynamic solution concepts for repeated games [62, 109].

Online and reinforcement learning [55] In online learning (a.k.a. multi-armed bandit [56, 116]), data is gathered and treated on the fly. For instance, consider an online binary classification problem. Some unlabelled data $X_t \in \mathbb{R}^d$ is observed, and the agent predicts its label Y_t ; let us denote $\hat{Y}_t \in \pm 1$ the prediction. The agent potentially observes the loss $\mathbb{1}\{Y_t \neq \hat{Y}_t\}$ and then receives another new unlabeled data example X_{t+1} . In that specific problem, the typical learning objective is to perform asymptotically as good as the best classifier f^* in some given class \mathcal{F} , i.e., such that the loss $\sum_{t=1}^T \mathbb{1}\{Y_t \neq \hat{Y}_t\}$ is $o(T)$ -close to $\max_{f \in \mathcal{F}} \sum_{t=1}^T \mathbb{1}\{Y_t \neq f(X_t)\}$; the difference between those terms is called *regret*. The more general model with an underlying state of the world $S_t \in \mathcal{S}$ that evolves at each step following some Markov Decision Process (MDP, i.e., the transition matrix from S_t to S_{t+1} depend on the actions of the agent) and impacts the loss function is called reinforcement learning (RL). RL is an incredibly powerful learning technique, provided enough data are available since learning is usually quite slow. This is why the recent successes involve settings with heavy simulations (like games) or well-understood physical systems (like robots).

These techniques will be central to our approach as we aim to model problems where ground truth data is not available upfront and problems involving sequential decision making. There have been some successful first results in that direction. For instance, there are applications (e.g., cognitive radio) where several agents (users) aim at finding a matching with resources (the different bandwidth). They can do that by "probing" the resources, estimating their preferences and trying to find some stable matchings [53, 101].

Online algorithms [51] and **theoretical computer science** Online algorithms are closely related to online learning with a major twist. In online learning, the agent has "0-look ahead"; for instance, in the online binary classification example, the loss at stage t was $\mathbb{1}\{Y_t \neq \hat{Y}_t\}$ but Y_t was not known in advance. The comparison class, on the other hand, was the empirical performance of a given set of classifiers. In online algorithms, the agents have "1-look ahead"; in the classification example, this means that Y_t is known before choosing \hat{Y}_t . But the overall objective is obviously no longer the minimisation of the empirical error, but the minimisation of this error plus the total number of changes (say). The comparison class is then larger, namely a subset of admissible (or the whole set) sequences of prediction $\{\pm 1\}^T$. The typical and relevant example of online problem relevant for Criteo that will be investigated is the matching problem: agents and resources arrive sequentially and must be, if possible, paired together as fast as possible (and as successfully as possible). Variants of these problems include the optimal stopping time question (when/how make a final decision) such as prophet inequalities and related questions [68],

Optimal transport [126] Optimal transport is a quite old problem introduced by Monge where an agent aims at moving a pile of sand to fill a hole at the smallest possible price. Formally speaking, given two probability measures μ and ν on some space \mathcal{X} , the optimal transport problem consist in finding (if it exists, otherwise the problem can be relaxed) a transport map $T : \mathcal{X} \rightarrow \mathcal{X}$ that minimizes $\int_{\mathcal{X}} c(x, T(x)) d\mu(x)$ for some cost function $c : \mathcal{X}^2 \rightarrow \mathbb{R}$, under the constraint that $T\#\mu = \nu$, where $T\#\mu$ is the push-forward measure of μ by T . Interestingly, when μ and ν are empirical measures, i.e., $\mu = \frac{1}{N} \sum_{n=1}^N \delta_{x_n}$ and $\nu = \frac{1}{N} \sum_{n=1}^N \delta_{y_n}$, a transport map is nothing more than a matching between $\{x_n\}$ and $\{y_n\}$ that minimizes the cost $\sum_n c(x_n, T(x_n))$.

Recently, optimal transport gained a lot of interest in the ML community [115] thanks to its application to images and to new techniques to compute approximate matchings in a tractable way [118]. Even more

unexpected applications of optimal transport have been discovered: to protect privacy [54], fairness [47], etc. Those connections are promising, but only primitive for the moment. For instance, consider the problem of matching students to schools. The unfairness level of a school can be measured as the Wasserstein distance between the distribution of the students within that school compared to the overall distribution of students. Then the matching algorithms could have a constraint of minimizing the sum of (or its maximum) unfairness levels; alternatively, we could aim at designing mechanisms giving incentives to schools to be fair in their allocation (or at least in their list preferences), typically by paying a higher fee if the unfairness level is high.

2.3 General objectives

The overarching objective of FAIRPLAY is to **create algorithms to learn for and with users—and techniques to analyze them—**, through the study of **multi-agent learning systems where the agents can be cooperatively or competitively learning agents, or agents providing or generating data**, while guaranteeing that **fairness and privacy** constraints are satisfied for the involved users. We detail this global objective into a number of more specific ones.

Objective 1: Developing fair and private mechanisms

Our first objective is to incorporate ethical aspects of fairness and privacy in mechanisms used in typical problems occurring in Internet applications, in particular auctions, matching, and recommendation. We will focus on social welfare and consider realistic cases with multiple agents and sequential learning that occur in practice due to sequential decision making. Our objective is both to construct models to analyze the problem, to devise algorithms that respect the constraints at stake, and to evaluate the different trade-offs in standard notions of utility introduced by ethical constraints.

Objective 2: Developing multi-agent statistics and learning

Data is now acquired, treated and/or generated by a whole network of agents interacting with the environment. There are also often multiple agents learning either collaboratively or competitively. Our second objective is to build a new set of tools to perform statistics and learning tasks in such environments. To this end, we aim at modeling these situations as multi-agent systems and at studying the dynamics and equilibrium of these complex game-theoretic situations between multiple learning algorithms and data providers.

Objective 3: Improving the theoretical state of the art

Research must rely on theoretical, proven guarantees. We develop new results for the techniques introduced before, such as prophet inequalities, (online) matchings, bandits and RL, etc.

Objective 4: Proposing practical solutions and enhancing transfer from research to industry

Our last scientific objective is to apply and implement theoretical works and results to practical cases. This will be a crucial component of the project as we focus on transfer within Criteo.

Objective 5: Scientific Publications

We aim at publishing our results in top-tier machine learning conferences (NeurIPS, ICML, COLT, ICLR, etc.) and in top-tier game theory journals (Games and Economic Behavior, Mathematics of OR, etc.). We will also target conferences at the junction of those fields (EC, WINE, WebConf, etc.) as well as conferences specifically on security and privacy (IEEE S&P, NDSS, CSS, PETS, etc.) and on fairness (FAccT, AIES).

All the five objectives are interlaced. For instance, fairness and privacy constraints are important in Objective 2 whereas the multi-agent aspect is also important in Objective 1. Objectives 4 and 5 are transversal and present in all the first three objectives.

3 Research program

To reach the objectives laid out above, we organize the research in three themes. The first one focuses on developing fair mechanisms. The second one considers private mechanisms, and in particular considers the challenge of reconciling fairness and privacy—which are often conflicting notions. The last theme, somewhat transverse to the first two, consists in leveraging/incorporating structure in all those problems in order to speed up learning. Of course, all themes share common points on both the problems/applications considered and the methods and tools used to tackle them; hence there are cross-fertilization between the different themes.

3.1 Theme 1: Developing fair mechanisms for auctions and matching problems

3.1.1 Fairness in auction-based systems

Online ads platforms are nowadays used to advertise not just products, but also *opportunities* such as jobs, houses, or financial services. This makes it crucial for such platforms to respect fairness criteria (be it only for legal reasons), as an unfair ad system would deprive a part of the population of some potentially interesting opportunities. Despite this pressing need, there is currently no technical solution in place to provably prevent discriminations. One of the main challenge is that ad impression decisions are the outcome of an auction mechanism that involves bidding decisions of multiple self-interested agents controlling only a small part of the process, while group fairness notions are defined on the outcome of a large number of impressions. We propose to investigate two mechanisms to guarantee fairness in such a complex auction-based system (note that we focus on online ad auctions but the work has broader applicability).

Advertiser-centric (or bidder-centric) fairness We first focus on advertiser-centric fairness, i.e., the advertiser of a third-party needs to make sure that the reached audience is fair independently of the ad auction platform. A key difficulty is that the advertiser does not control the final decision for each ad impression, which depends on the bids of other advertisers competing in the same auction and on the platform’s mechanism. Hence, it is necessary that the advertiser keeps track of the auctions won for each of the groups and dynamically adjusts its bids in order to maintain the required balance.

A first difficulty is to model the behavior of other advertisers. We can first use a mean-field games approach similar to [87] that approximates the other bidders by an (unknown) distribution and checks equilibrium consistency; this makes sense if there are many bidders. We can also leverage refined mean-field approximations [81] to provide better approximations for smaller numbers of advertisers. Then a second difficulty is to find an optimal bidding policy that enforces the fairness constraint. We can investigate two approaches. One is based on an MDP (Markov Decision Process) that encodes the current fairness level and imposes a hard constraint. The second is based on modeling the problem as a contextual bandit problem. We note that in addition to fairness constraints, privacy constraints may complicate the optimal solution finding.

Platform-centric (or auction-centric) fairness We also consider the problem from the platform’s perspective, i.e., we assume that it is the platform’s responsibility to enforce the fairness constraint. We also focus here on demographic parity. To make the solution practical, we do not consider modification of the auction mechanism, instead we consider a given mechanism and let the platform adapt dynamically the bids of each advertiser to achieve the fairness guarantee. This approach would be similar to the pacing multipliers used by some platforms [67, 66], but using different multipliers for the different groups (i.e., different values of the sensitive attribute).

Following recent theoretical work on auction fairness [60, 86, 64] (which assumes that the targeted population of all ads is known in advance along with all their characteristics), we can formulate fairness as a constraint in an optimization problem for each advertiser. We study fairness in this static auction problem in which the auction mechanism is fixed (e.g., to second price). We then move to the online setting in which users (but also advertisers) are dynamic and in which decisions must be taken online, which we approach through dynamic adjustment of pacing multipliers.

3.1.2 Fairness in matching and selection problems

In this second part, we study fairness in selection and matching problems such as hiring or college admission. The selection problem corresponds to any situation in which one needs to select (i.e., assign a binary label to) data examples or individuals but with a constraint on the number of maximum number of positive labels. There are many applications of selection problems such as police checks, loan approvals, or medical screening. The matching problem can be seen as the more general variant with multiple selectors. Again, a particular focus is put here on cases involving repeated selection/matching problems and multiple decision makers.

Fair repeated multistage selection In our work [75], we identified that a key source of discrimination in (static) selection problems is *differential variance*, i.e., the fact that one has quality estimates that have different variances for different groups. In practice, however, the selection problem is often ran repeatedly (e.g., at each hiring campaign) and with partial (and increasing) information to exploit for making decisions.

Here, we consider the repeated multistage selection problem, where at each round a multistage selection problem is solved. A key aspect is that, at the end of a round, one learns extra information about the candidates that were selected—hence one can refine (i.e., decrease the variance of) the quality estimate for the groups in which more candidates were selected. We will first rethink fairness constraints in this type of repeated decision making problems. Then we will both study the discrimination that come out of natural (e.g., greedy) procedure as well as design (near) optimal ones for the constraints at stake. We also investigate how the constraints affect the selection utility.

Multiple decision-makers Next, we investigate cases with multiple decision-makers. We propose two cases in particular. The first one is the simple two-stage selection problem but where the decision-maker doing the first-stage selection is different from the decision-maker doing the second-stage selection. This is a typical case for instance for recruiting agencies that propose sublists of candidates to different firms that wish to hire. The second case is when multiple-decision makers are trying to make a selection simultaneously—a typical example of this being the college admission problem (or faculty recruitment). We intend to model it as a game between the different colleges and to study both static solutions as well as dynamic solutions with sequential learning, again modeling it as a bandit problem and looking for regret-minimizing algorithms under fairness constraints. A number of important questions arise here: if each college makes its selection independently and strategically (but based on quality estimates with variances that differ amongst groups), how does it affect the “global fairness” metrics (meaning in aggregate across the different colleges) and the “local fairness” metrics (meaning for an individual college)? What changes if there is a central admission system (such as Parcoursup)? And in this later case, how to handle fairness on the side of colleges (i.e., treat each college fairly in some sense)?

Fair matching with incentives in two-sided platforms We will study specifically the case of a platform matching demand on one side to offer on the other side, with fairness constraints on each side. This is the case for instance in online job markets (or crowdsourcing). This is similar to the previous case but, in addition, here there is an extra incentives problem: companies need to give the right incentives to job applicants to accept the jobs; while the platform doing the match needs to ensure fairness on both sides (job applicants and companies). This gives rise to a complicated interplay between learning and incentives that we will tackle again in the repeated setting.

We finally mention that, in many of these matching problems, there is an important time component: each agent needs to be matched “as soon as possible”, yielding a trade-off between the delay to be matched and the quality of the match. There is also a problem of participation incentives; that is how the matching algorithm used affect the behavior of the participants in the matching “market” (participation or not, information revelation, etc.). In the long-term, we will incorporate these aspects in the above models.

Throughout the work in this theme, we will also consider a question transverse and present in all the models above: how can we handle *multidimensional fairness*, that is, where there are multiple sensitive attributes and consequently an exponential number of sub-groups defined by all intersections; this combinatorial is challenging and, for the moment, still exploratory.

3.2 Theme 2: Reconciling, and enforcing privacy with fairness

In the previous theme, we implicitly assumed that we know the users' group, i.e., their sensitive attributes such as gender, age, or ethnicity. In practice, one of the key question when implementing fairness mechanisms is how to measure/control fairness metrics without having access to these protected attributes. This question relates to the link between privacy and fairness and the trade-off between them (as fairness requires data and privacy tends to protect it) [119, 69].

A first option to solve this problem would be (when it is possible) to build proxies [82, 122] for protected attributes using available information (e.g., websites visited or products bought) and to measure or control for fairness using those in place of the protected attributes. As the accuracy of these proxies cannot be assessed, however, they cannot be used for any type of "public certification"—that is, for a company to show reasonable fairness guarantees to clients (e.g., as a commercial argument), or (even less) to regulators. Moreover, in many cases, the entity responsible for fairness should not be accessing sensitive information, even through proxies, for privacy reasons.

In FAIRPLAY, we investigate a different means of **certifying fairness of decisions** without having access to sensitive attributes, by partnering with a **trusted third-party that collects protected attributes** (that could for instance be a regulator or a public entity, such as Nielsen, say). We distinguish two cases:

1. If the third-party and the company share a common identifier of users, then computing the fairness metric without leaking information to each other will boil down to a problem of secure multi-party computation (SMC). In such a case, there could be a need to be able to learn, which opens the topic of learning and privacy under SMC. This scenario, however, is likely not the most realistic one as having a common identifier requires a deep technical integration.
2. If the third-party and the company do not share a common identifier of users, but there are common features that they both observe [90], then it is possible only to partially identify the joint distribution. With additional structural assumptions on the distribution, however, it could be identified accurately enough to estimate biases and fairness metrics. This becomes a distribution identification problem and brings a number of questions such as: how to do the distribution identification? how to optimally query data from the third party to train fair algorithms with high utility? etc. An important point to keep in mind in such a study is that it is likely that the third party user-base is different from that of the company. It will therefore be key to handle the covariate shift from one distribution to the other while estimating biases.

This distribution identification problem will be important in the context of privacy, even independently of fairness issues. Indeed, in the near future, most learning will happen in a privacy-preserving setting (for instance, because of the Chrome privacy sandbox). This will require new learning schemes (different from e.g., Empirical Risk Minimization) as samples from the usual joint distribution (X, Y) of samples/labels will no longer be observed. Only aggregated data—e.g., (empirical) marginals of the form $E[Y|X_2 = 4, X_7 = \text{"lemonde.fr"}]$ —will be observed, with a limited budget of requests. This also brings questions such as how to mix it with ERM on some parts of the traffic, what is the (certainly adaptive or active) optimal strategy to query the marginals, etc. This problem will be further complicated by the fact that privacy (for instance through the variety of consents) will be heterogeneous: all features are not available all the time. This is therefore strongly related to learning with missing features and imputation [89].

In relation to the above problems, a key question is to determine what is the most appropriate definition of fairness to consider. Recall that it is well-known that usual fairness metrics are not compatible [94]. Moreover, in online advertising, fairness can be measured at multiple different levels: at the level of bids, at the level of audience reached, at the level of clicking users, etc. Fairness at the level of bids does not imply fairness of the audience reached (see Theme 1); yet external auditors would measure which ad is displayed—as was done for some ad platforms [123]—hence in terms of public image, that would be the appropriate level to define fairness. Intimately, the above problem relates to the question of measuring what is the relevant audience of an ad, which would define the label if one were to use the EO fairness notion. This label is typically not available. We can explore three ways to overcome this issue. The first is to find a sequential way to learn the label through users clicking on ads. The second and third options are to focus in a first step on individual fairness, or on counterfactual fairness [97], which has many possible different level of assumptions and was popularized in 2020 [98]. The notion of counterfactual is key

in causality [117]. A model is said counterfactually fair if its prediction does not change (too much) by intervening on the sensitive attribute. Several works already propose ways of designing models that are counterfactually fair [93, 130, 129]. This seems to be quite an interesting, but challenging direction to follow.

Finally, an alternative direction would be to pursue modeling the trade-off between privacy and fairness. For instance, in some game theoretic models, users can choose the quantity of data that they reveal [80, 54], so that the objective functions integrate different levels of fairness/privacy. Then those models should be studied both in terms of equilibrium and in the online setup, with the objective of identifying how the strategic privacy considerations affect the fairness-utility tradeoff.

3.3 Theme 3: Exploiting structure in online algorithms and learning problems

Our last research direction is somewhat transverse, with possible application to improving algorithms in the first three themes. We explore how the underlying structure can be exploited, in the online and learning problems considered before, to improve performance. Note that, in all these problems, we will incorporate the fairness and privacy aspects even if they are somewhat transverse to the structure considered.¹ The following sections are illustrating examples on how hidden structure can be leveraged in specific examples.

3.3.1 Leveraging structure in online matching

Finding large *matchings* in graphs is a longstanding problem with a rich history and many practical and theoretical applications [108, 85, 45, 44]. Recall that given a graph $G = (\mathcal{V}, \mathcal{E})$ —where \mathcal{V} is a set of vertices and \mathcal{E} is a set of edges—, a matching $\mathcal{M} \subset \mathcal{E}$ is a subset of edges such that each vertex belongs to at most one edge $e \in \mathcal{M}$. In that context, a perfect matching \mathcal{M} is a matching where each vertex $v \in \mathcal{V}$ is associated to an edge $e \in \mathcal{M}$, and a maximum matching is a matching of maximum size (one can also consider weights on edges). Here, we study an online setting, which is more adequate in applications such as Internet advertising where ad impressions must be assigned to available ad slots [108, 57]. Consider a bipartite graph, where $\mathcal{V} = U \cup V$ is the union of two disjoint sets. Nodes $u \in U$ are known beforehand, whereas nodes $v \in V$ are discovered one at a time, along with the edges they belong to, and must be either immediately matched to an available (i.e., unmatched yet) vertex $u \in U$ or discarded. Online bipartite matching is relevant in two-sided markets besides ad allocations such as assigning tasks to workers [85].

A natural measure for the quality of an online matching algorithm is the “competitive ratio” (CR): the ratio between the size of the created matching to the size of the optimal one [108]. The seminal work [92] introduced an optimal algorithm for the adversarial case [48], that guarantees a CR of $1 - \frac{1}{e}$; but focusing on a pessimistic worst-case. In practice, some relevant knowledge (either given *a priori* or learned during the process) on the underlying structure of the problem can be leveraged. The focus then shifted to models taking into account some type of stochasticity in the arrival model, mostly for the i.i.d. model where arriving vertices $v \in V$ are drawn from a fixed distribution \mathcal{D} [88, 46, 76, 91, 103, 104]. The classical approach consists in optimizing the CR over the distribution \mathcal{D} . Even in this seemingly optimistic framework, however, it is now known that there is no hope for a CR of more than 0.823 [104]. Moreover, this generally leads to very large linear programs (LP).

A more recent approach restricts the distribution \mathcal{D} over which the problem is optimized to classes of graphs with an underlying stochastic structure. The benefit of this approach is two-fold: it gives hope for higher competitive ratios, and for simpler algorithms. Experiments also proved that complex algorithms optimized on \mathcal{D} fared no better than simple greedy heuristics on “real-life” graphs [52]. A few results along these lines show that is a promising path. For instance, [57] studied the problem on graphs where each vertex has degree at least d and found a competitive ratio of $1 - (1 - 1/d)^d$. On d -regular graphs, [65] designed a $1 - O(\sqrt{\log d}/\sqrt{d})$ competitive algorithm. [106] showed that greedy algorithms were highly efficient on Erdős-Renyi random graphs, with a competitive ratio of 0.837 in the worst case. [43] showed that in a specific market with two types of matching agents, the behavior of the matching algorithm varies with the homogeneity of the market. Our goal here is to go beyond the independence assumption underlying all these works.

¹One may worry that the structure added adds biases. This is typically fine because we control for fairness on the output directly, but this may lead to a degradation of the benefit of adding structure.

Introducing correlation and inhomogeneity We will start by deriving and studying optimal online matching strategies on widely studied classes of graphs that present simple inhomogeneity or correlation structures (which are often present in applications). The stochastic block model [40] is often used to generate graphs with an underlying community structure. It presents a simple correlation structure: two vertices in the same community are more likely to have a common neighbors than two vertices in different communities. Another focus point will be a generalized version of the Erdős-Renyi model, where the in-place vertices $u \in \mathcal{U}$ are divided into sets s_i , where $u \in s_i$ generates an edge with probability $p_i = \frac{c_i}{n}$. These two settings should give us a better understanding of how heterogeneity and correlation affect the matching performance.

Deriving the competitive ratio implies to study the asymptotic size of maximum matchings in random graphs. Two methods are usually used. The first and constructive one is the study of the Karp-Sipser algorithm on the graph [49]. The second one involves the rich theory of graph weak local convergence [50]. A straightforward application of the methods, however, requires the graph to have independence properties; adapting them to graphs with a correlation structure will require new ideas.

Configuration models and random geometric graphs A configuration model is described as follows (in the bi-partite case). Each vertex $u \in \mathcal{U}$ has a number of half-edges drawn for the same distribution $\pi_{\mathcal{U}}$ and each vertex $v \in \mathcal{V}$ has a number of half-edges drawn from $\pi_{\mathcal{V}}$ (with the assumption that the expected total numbers of half edges from \mathcal{U} and \mathcal{V} are the same). Then a vertex $v \in \mathcal{V}$ that arrives in the sequential fashion has its half-edges “completed” by a (still free) half-edge of \mathcal{U} . This is a standard way of creating random graphs with (almost) fixed distribution of degrees. Here the question would simply be the competitive ratio of some greedy algorithm, whether the distributions $\pi_{\mathcal{U}}$ and $\pi_{\mathcal{V}}$ are known beforehand or learned on the fly. An interesting variant of this problem would be to assume the existence of a (hidden or not) geometric graph. Each $u \in \mathcal{U}$ is drawn i.i.d in \mathcal{R}^d (say a Gaussian centered at 0) and similarly for $v \in \mathcal{V}$. Then there is an edge between u and v with a probability depending on the distance between them. Here again, interesting variants can be explored depending on whether the distribution is known or not, and whether the locations of u and/or v are observed or not.

Learning while matching In practical applications, the full stochastic structure of the graphs may not be known beforehand. This begs the question: what will happen to the performance of the algorithms if the graph parameters are learned while matching? In the generalized Erdős-Renyi graph, this will correspond to learning the probability of generating edges. For the stochastic block model, the matching algorithm will have to perform online community detection.

3.3.2 Exploiting side-information in sequential learning

We end with an open direction that may be relevant to many of the problems considered above: how to use side-information to speed-up learning. In many sequential learning problems where one receives feedback for each action taken, it is actually possible to deduce, for free, extra information from the known structure of the problem. However, how to incorporate that information in the learning process is often unclear. We describe it through two examples.

One-sided feedback in auctions In online ad auctions, the advertisers’ strategy is to bid in a compact set of possible bids. After placing a bid, the advertiser learns whether they won the auction or not; but even if they do not observe the bids of other advertisers, they can deduce for free some extra information: if they win they learn that they would have won with any higher bid and if they lose they learn that they would have lost with any lower bid [127, 61]. We will investigate how to incorporate this extra information in RL procedures devised in Theme 1. One option is by leveraging the Kaplan-Meier estimator.

Side-information in dynamic resource allocation problems and matching Generalizing the idea above, one can observe side-information in many other problems [42]. Typically, in resource allocation problems (e.g., how to allocate a budget of ad impressions), one can leverage a monotony property: one would not have gained more by allocating less. Similarly, in matching with unknown weights, it is often possible upon doing a particular match to learn the weight of other potential pairs.

4 Application domains

4.1 Typical problems and use-cases

In FAIRPLAY, we focus mainly on problems involving learning that relate to Internet applications. We will tackle generic problems (in particular auctions and matching) with different applications in mind, in particular some applications in the context of Criteo’s business but also others. A crucial property of those applications is the aforementioned ethical concerns, namely fairness and/or privacy. The team was shaped and motivated by several such use-cases, from more practical (with possible short or middle term applications in particular in Criteo products) to more theoretical and exploratory ones. We describe first here the main types of generic problems and use-cases considered in this context.

Auctions [96] There are many different types of auctions that an agent can use to sell one or several items in her possession to n potential buyers. This is the typical way in which spots to place ads are sold to potential advertisers. In case of a single item, the seller asks buyers to bid $b_i \in [0, 1]$ on the item and the winner of the item is designating via an “allocation rule” that maps bids $b \in [0, 1]^n$ to a winner in $\{0, \dots, n\}$ (0 refers to the no winner case). Then the payment rule $p : [0, 1]^n \rightarrow [0, 1]^n$ indicates the amount of money that each bidder must pay to the seller. Auctions are specific cases of a broader family of “mechanisms”. Knowing the allocation and payment rules, bidders have incentives to bid strategically. Different auctions (or rules) end up with different revenue to the seller, who can choose the optimal rules. This is rather standard in economics, but these interactions become way more intricate when repeated over time (as in the online ad market [112]), when several items are sold at the same time (for instance in bundles), when the buyers have partial information about the actual value of the item [127] and/or reciprocally when the seller does not know the value distributions of the buyer. In that case, she might be tempted to try to learn them from the previous bids in order to design the optimal mechanism. Knowing this, the bidders have incentives to long term strategic behaviors, ending up in a quite complicated game between learning algorithms [113]. This setting of interacting algorithms is actually of interest by itself, irrespectively of ad auctions. It is noteworthy also that traditional auction mechanisms do not guarantee any fairness notion and that the literature on fixing that (for applications where it matters) is only nascent [60, 111, 64, 86].

Matching [95, 110] A matching is nothing more than a bi-partite graph between some agents (patients, doctors, students) and some resources (respectively, organs, hospital, schools). The objective is to figure out what is the “optimal” matching for a given criterion. Interestingly, there are two different—and mostly unrelated yet—concepts of “good matching”. The first one is called “stable” in the sense that each agent expresses preferences over resources (and vice-versa) and be such that no couple (agent-resource) that are un-paired would prefer to be paired together than with their current paired resource/agent. In the other concept of matching, agents and resources are described by some features (say vectors in \mathbb{R}^d , denoted by a_n for agents and r_m for resources) and pairing a_n to r_m incurs a cost of $c(a_n, r_m)$, for some a given function $c : (\mathbb{R}^d)^2 \rightarrow [0, 1]$. The objective is then to minimize the total cost of the matching $\sum_n c(a_n, r_{\sigma(n)})$, where $\sigma(n)$ is the resource allocated to agent n .

Matching is used in many different applications such as university admission (e.g., in Parcoursup). Notice that strategic interactions arise in matching if agents or resources can disclose their preferences/features to each other. Learning is also present as soon as not everything is known, e.g., the preferences or costs. Many applications of matching (again, such as college admission) are typical examples where fairness and privacy are of utmost importance. Finally, matching is also at the basis of several Internet applications and Criteo products, for instance to solve the problem of matching a given ad budget to fixed ad slots.

Ethical notions in those use-cases In both problems, individual users are involved and there is a clear need to consider fairness and privacy. However, the precise motivation and instantiation of these notions depends on the specific use-case. In fact, it is often part of the research question to decide which are the most relevant fairness and privacy notions, as mentioned in Section 2.1. We will throughout the team’s life put an important focus on this question, as well as on the question of the impact of the chosen notion on performance.

4.2 Application areas

In FAIRPLAY, we consider both applications to Criteo use-cases (online advertisement) and other applications (with other appropriate partners).

4.2.1 Online advertisement

Online advertising offers an application area for all of the research themes of FAIRPLAY; which we investigate primarily with Criteo.

First, online advertising is a typical application of online auctions and we consider applications of the work on auctions to Criteo use-cases, in particular the work on advertiser-centric fairness where the advertiser is Criteo. From a practical point of view, privacy will have to be enforced in such applications. For instance, when information is provided to advertisers to define audiences or to visualize the performance of their campaigns (insights) there is a possibility of leaking sensitive information on users. In particular, excellent proxies on protected attributes should probably not be leaked to advertisers, or transformed before (e.g., with the differential privacy techniques). This is therefore also an application of the fairness-vs-privacy research thread.

Note that, even before considering those questions, the first very important theoretical question is to determine what is the more appropriate definition of fairness (as there are, as mentioned above, many different variations) in those applications. We recall that it is well-known that usual fairness metrics are not compatible [94]. Moreover, in online advertising, fairness can be measured in term of bidding and recommendation or in term of what ads are actually displayed. Being fair on bidding does not lead to fairness in ads displaying [111], mainly because of the other advertising actors. While fairness in bidding and/or recommendation seem the most important because they only rely on our models, external auditors can easily get data on which ads we display.

We will also investigate applications of fair matching techniques to online advertising and to Criteo matching products—namely retargeting (personalized ads displayed on a website) and retail media (sponsored products on a merchant website). Indeed, one of Criteo major products, retail media can be cast as an online matching problem. On a given e-commerce website (say, target), several advertisers—currently brands—are running campaigns so that their products are “sponsored” or “boosted”, i.e., they appear higher on the list of results of a given query. The budgets (from a Criteo perspective) must be cleared (daily, monthly or annually). This constraint is easy thanks to the high traffic, but the main issue is that, without control/pacing/matching in times, the budget is depleted after only a few hours on a relatively low quality traffic (i.e., users that generate few conversions hence a small ROI for the advertisers). The question is therefore whether an individual user should be matched or not to boosted/sponsored products at a given time so that the ROI of the advertisers is maximized, the campaign budget is depleted and the retailer does not suffer too much from this active corruption of its organic results. Those are three different and concurrent objectives (for respectively the advertisers, Criteo and the retailers) that must be somehow conciliated. This problem (and more generally this type of problems) offers a rich application area to the FAIRPLAY research program. Indeed, it is crucial to ensure that fairness and privacy are respected. On the other hand, users, clicks, conversion arrival are not “worst case”. They rather follow some complicated—but certainly learnable—process; which allows applying our results on exploiting structure.

4.2.2 “Physical matching”

We investigate a number of other applications of matching: assignment of daycare slots to kids, mutation of professors to different academies, assignment of kidneys to patients, assignment of job applicants to jobs. In all these applications, there are crucial constraints of fairness that complicate the matching. We leverage existing partnership with the CNAF, the French Ministry of Education and the Agence de la biomédecine in Paris for the first three applications; for the last we will consolidate a nascent partnership with Pole Emploi and LinkedIn.

5 Highlights of the year

Simon Mauras joined the team in Feb 2024 as a CR.

The team's kickoff meeting took place on Oct 3rd in the Turing building.

The team had a record number of 11 papers accepted at NeurIPS'24.

5.1 Awards

Dr. Solenne Gaucher, a postdoc in the team, received the "L'Oréal-UNESCO Young Talents Award".

6 New results

6.1 Auctions and mechanism design

Participants: Benjamin Heymann, Patrick Loiseau, Simon Mauras, Vianney Perchet, Hugo Richard.

6.1.1 Auctions

Most of the display advertising inventory is sold through real-time auctions. The participants of these auctions are typically bidders (Google, Criteo, RTB House, and Trade Desk for instance) that participate on behalf of advertisers. In order to estimate the value of each display opportunity, they usually train advanced machine learning algorithms using historical data. In the labeled training set, the inputs are vectors of features representing each display opportunity, and the labels are the generated rewards. In practice, the rewards are given by the advertiser and are tied to whether a particular user converts. Consequently, the rewards are aggregated at the user level and never observed at the display level. A fundamental task that has, to the best of our knowledge, been overlooked is to account for this mismatch and split, or attribute, the rewards at the right granularity level before training a learning algorithm. We call this the label attribution problem. In [4], we develop an approach to the label attribution problem, which is both theoretically justified and practical. In particular, we develop a fixed point algorithm that allows for large-scale implementation and showcase our solution using a large-scale publicly available data set from Criteo, a large demand-side platform. We dub our approach the fixed point label attribution algorithm. Managerial implications: There is often a hidden leap of faith when transforming the advertiser's signal into display labeling. Demand Side Platforms providers should be careful when building their machine learning pipeline and carefully solve the label attribution step.

In [8], we use a fixed point gradient flow algorithm to compute the equilibria of first-price procurement auctions in the presence of losses and Bayesian priors. We use this efficient algorithm to compare optimal, first-price and VCG auctions. This allows us to numerically estimate the social cost of sub-optimality of the nodal pricing mechanism in wholesale electricity markets. We also derive a closed form expression of the optimal mechanism procurement cost when the types are uniformly distributed. In [9], we discuss a procurement problem with transportation losses and piecewise linear production costs. We first provide an algorithm based on Knaster-Tarski's fixed point theorem to solve the allocation problem in the quadratic losses case. We then identify a monotony condition on the types distribution under which the Bayesian cost minimizing mechanism takes a simple form.

In [11], motivated by online display advertising, we consider repeated second-price auctions, where agents sample their value from an unknown distribution with cumulative distribution function F . In each auction t , a decision-maker bound by limited observations selects n_t agents from a coalition of N to compete for a prize with p other agents, aiming to maximize the cumulative reward of the coalition across all auctions. The problem is framed as an N -armed structured bandit, each number of player sent being an arm n , with expected reward $r(n)$ fully characterized by F and $p + n$. We present two algorithms, Local-Greedy (LG) and Greedy-Grid (GG), both achieving constant problem-dependent regret. This relies on three key ingredients: 1. an estimator of $r(n)$ from feedback collected from any arm k , 2. concentration bounds of these estimates for k within an estimation neighborhood of n and 3. the unimodality property

of r under standard assumptions on F . Additionally, GG exhibits problem-independent guarantees on top of best problem-dependent guarantees. However, by avoiding to rely on confidence intervals, LG practically outperforms GG, as well as standard unimodal bandit algorithms such as OSUB or multi-armed bandit algorithms.

In [25], motivated by the strategic participation of electricity producers in electricity day-ahead market, we study the problem of online learning in repeated multi-unit uniform price auctions focusing on the adversarial opposing bid setting. The main contribution of this paper is the introduction of a new modeling of the bid space. Indeed, we prove that a learning algorithm leveraging the structure of this problem achieves a regret of $\tilde{O}(K^{4/3}T^{2/3})$ under bandit feedback, improving over the bound of $\tilde{O}(K^{7/4}T^{3/4})$ previously obtained in the literature. This improved regret rate is tight up to logarithmic terms. Inspired by electricity reserve markets, we further introduce a different feedback model under which all winning bids are revealed. This feedback interpolates between the full-information and bandit scenarios depending on the auctions' results. We prove that, under this feedback, the algorithm that we propose achieves regret $\tilde{O}(K^{5/2}\sqrt{T})$.

6.1.2 Dynamic pricing

In [26], we proposed improved algorithms for contextual dynamic pricing. In contextual dynamic pricing, a seller sequentially prices goods based on contextual information. Buyers will purchase products only if the prices are below their valuations. The goal of the seller is to design a pricing strategy that collects as much revenue as possible. We focus on two different valuation models. The first assumes that valuations linearly depend on the context and are further distorted by noise. Under minor regularity assumptions, our algorithm achieves an optimal regret bound of $\tilde{O}(T^{2/3})$, improving the existing results. The second model removes the linearity assumption, requiring only that the expected buyer valuation is β -Hölder in the context. For this model, our algorithm obtains a regret $\tilde{O}(T^{d+2\beta/d+3\beta})$, where d is the dimension of the context space.

6.1.3 Mechanism design

In [17], we study auction design within the widely acclaimed model of interdependent values, introduced by Milgrom and Weber [1982]. In this model, every bidder i has a private signal s_i for the item for sale, and a public valuation function $v_i(s_1, \dots, s_n)$ which maps every vector of private signals (of all bidders) into a real value. A recent line of work established the existence of approximately-optimal mechanisms within this framework, even in the more challenging scenario where each bidder's valuation function v_i is also private. This body of work has primarily focused on single-item auctions with two natural classes of valuations: those exhibiting submodularity over signals (SOS) and d -critical valuations. In this work we advance the state of the art on interdependent values with private valuation functions, with respect to both SOS and d -critical valuations. For SOS valuations, we devise a new mechanism that gives an improved approximation bound of 5 for single-item auctions. This mechanism employs a novel variant of an "eating mechanism", leveraging LP-duality to achieve feasibility with reduced welfare loss. For d -critical valuations, we broaden the scope of existing results beyond single-item auctions, introducing a mechanism that gives a $(d+1)$ -approximation for any environment with matroid feasibility constraints on the set of agents that can be simultaneously served. Notably, this approximation bound is tight, even with respect to single-item auctions.

In [19], we study best-of-both-worlds guarantees for the fair division of indivisible items among agents with subadditive valuations. Our main result establishes the existence of a random allocation that is simultaneously ex-ante $\frac{1}{2}$ -envy-free, ex-post $\frac{1}{2}$ -EFX and ex-post EF1, for every instance with subadditive valuations. We achieve this result by a novel polynomial-time algorithm that randomizes the well-established envy cycles procedure in a way that provides ex-ante fairness. Notably, this is the first best-of-both-worlds fairness guarantee for subadditive valuations, even when considering only EF1 without EFX.

In [6], we propose and theoretically analyze a measure to encourage greater voluntary contributions to public goods. Our measure is a simple intervention that restricts individuals' strategy sets by imposing a minimum individual contribution level while still allowing for full free riding for those who do not want to contribute. We show that for a well-chosen value of the minimum individual contribution level, this

measure does not incentivize any additional free riding while strictly increasing the total contributions relative to the situation without the minimum contribution level. Our measure is appealing because it is nonintrusive and in line with the principle of “freedom of choice.” It is easily implementable for many different public goods settings where more intrusive measures are less accepted. This approach has been implemented in practice in some applications, such as charities.

6.2 Matching

Participants: Matthieu Lerasle, Patrick Loiseau, Vianney Perchet.

In [7], we tackle the pair-matching problem. The pair-matching problem appears in many applications where one wants to discover good matches between pairs of entities or individuals. Formally, the set of individuals is represented by the nodes of a graph where the edges, unobserved at first, represent the good matches. The algorithm queries pairs of nodes and observes the presence/absence of edges. Its goal is to discover as many edges as possible with a fixed budget of queries. Pair-matching is a particular instance of multi-armed bandit problem in which the arms are pairs of individuals and the rewards are edges linking these pairs. This bandit problem is non-standard though, as each arm can only be played once. Given this last constraint, sub-linear regret can be expected only if the graph presents some underlying structure. This paper shows that sub-linear regret is achievable in the case where the graph is generated according to a Stochastic Block Model (SBM) with two communities. Optimal regret bounds are computed for this pair-matching problem. They exhibit a phase transition related to the Kesten-Stigum threshold for community detection in SBM. The pair-matching problem is considered in the case where each node is constrained to be sampled less than a given amount of times. We show how optimal regret rates depend on this constraint. The paper is concluded by a conjecture regarding the optimal regret when the number of communities is larger than 2. Contrary to the two communities case, we argue that a statistical-computational gap would appear in this problem.

In [18], we address the challenge of actively ranking a set of items/players with varying values/strengths. The comparison outcomes are random, with a greater noise the closer the values. A crucial requirement is that, at each iteration of the algorithm, all items must be compared once, i.e., an iteration is a perfect matching. Furthermore, we presume that comparing two players with closely matched strengths incurs no cost and, in contrast, a unit cost is associated with comparing players whose strength difference is more substantial. Our secondary objective is to determine an optimal matching between players based on this cost function: we propose and analyze an algorithm that draws on concepts from both AKS sorting networks and bandit theory. Our algorithm achieves both objectives with high probability, and the total cost is optimal (up to logarithmic terms).

6.3 Learning

Participants: Cristina Butucea, Matthieu Lerasle, Patrick Loiseau, Vianney Perchet, Maxime Vono, Hugo Richard.

6.3.1 Data valuation

In [20], we consider the dataset valuation problem, that is the problem of quantifying the incremental gain, to some relevant pre-defined utility of a machine learning task, of aggregating an individual dataset to others. The Shapley value is a natural tool to perform dataset valuation due to its formal axiomatic justification, which can be combined with Monte Carlo integration to overcome the computational tractability challenges. Such generic approximation methods, however, remain expensive in some cases. In this paper, we exploit the knowledge about the structure of the dataset valuation problem to devise more efficient Shapley value estimators. We propose a novel approximation, referred to as discrete uniform Shapley, which is expressed as an expectation under a discrete uniform distribution with support

of reasonable size. We justify the relevancy of the proposed framework via asymptotic and non-asymptotic theoretical guarantees and illustrate its benefits via an extensive set of numerical experiments.

6.3.2 Bandits and reinforcement learning

In [5], we provide a survey of recent results in multi-player bandits.

In [12], we consider the classical multi-armed bandit problem, but with strategic arms. In this context, each arm is characterized by a bounded support reward distribution and strategically aims to maximize its own utility by potentially retaining a portion of its reward, and disclosing only a fraction of it to the learning agent. This scenario unfolds as a game over T rounds, leading to a competition of objectives between the learning agent, aiming to minimize their regret, and the arms, motivated by the desire to maximize their individual utilities. To address these dynamics, we introduce a new mechanism that establishes an equilibrium wherein each arm behaves truthfully and discloses as much of its rewards as possible. With this mechanism, the agent can attain the second-highest average (true) reward among arms, with a cumulative regret bounded by $O(\log(T)/\Delta)$ (problem-dependent) or $O(\sqrt{T \log(T)})$ (worst-case).

In [24], we look at The Value of Reward Lookahead in Reinforcement Learning. In reinforcement learning (RL), agents sequentially interact with changing environments while aiming to maximize the obtained rewards. Usually, rewards are observed only after acting, and so the goal is to maximize the expected cumulative reward. Yet, in many practical settings, reward information is observed in advance – prices are observed before performing transactions; nearby traffic information is partially known; and goals are oftentimes given to agents prior to the interaction. In this work, we aim to quantifiably analyze the value of such future reward information through the lens of competitive analysis. In particular, we measure the ratio between the value of standard RL agents and that of agents with partial future-reward lookahead. We characterize the worst-case reward distribution and derive exact ratios for the worst-case reward expectations. Surprisingly, the resulting ratios relate to known quantities in offline RL and reward-free exploration. We further provide tight bounds for the ratio given the worst-case dynamics. Our results cover the full spectrum between observing the immediate rewards before acting to observing all the rewards before the interaction starts. In [23], we study reinforcement learning (RL) problems in which agents observe the reward or transition realizations at their current state before deciding which action to take. Such observations are available in many applications, including transactions, navigation and more. When the environment is known, previous work shows that this lookahead information can drastically increase the collected reward. However, outside of specific applications, existing approaches for interacting with unknown environments are not well-adapted to these observations. In this work, we close this gap and design provably-efficient learning algorithms able to incorporate lookahead information. To achieve this, we perform planning using the empirical distribution of the reward and transition observations, in contrast to vanilla approaches that only rely on estimated expectations. We prove that our algorithms achieve tight regret versus a baseline that also has access to lookahead information - linearly increasing the amount of collected reward compared to agents that cannot handle lookahead information.

6.4 Online (and offline) algorithms

Participants: Patrick Loiseau, Simon Mauras, Vianney Perchet.

6.4.1 Prophet inequalities

In [13], we study Lookback Prophet Inequalities. Prophet inequalities are fundamental optimal stopping problems, where a decision-maker observes sequentially items with values sampled independently from known distributions, and must decide at each new observation to either stop and gain the current value or reject it irrevocably and move to the next step. This model is often too pessimistic and does not adequately represent real-world online selection processes. Potentially, rejected items can be revisited and a fraction of their value can be recovered. To analyze this problem, we consider general decay functions D_1, D_2, \dots , quantifying the value to be recovered from a rejected item, depending on how far it has been observed

in the past. We analyze how lookback improves, or not, the competitive ratio in prophet inequalities in different order models. We show that, under mild monotonicity assumptions on the decay functions, the problem can be reduced to the case where all the decay functions are equal to the same function $x \mapsto \gamma x$, where $\gamma = \inf_{x>0} \inf_{j \geq 1} D_j(x)/x$. Consequently, we focus on this setting and refine the analyses of the competitive ratios, with upper and lower bounds expressed as increasing functions of γ .

In [14], we consider a hiring process with candidates coming from different universities. It is easy to order candidates with the same background, yet it can be challenging to compare them otherwise. The latter case requires additional costly assessments, leading to a potentially high total cost for the hiring organization. Given an assigned budget, what would be an optimal strategy to select the most qualified candidate? We model the above problem as a multicolor secretary problem, allowing comparisons between candidates from distinct groups at a fixed cost. Our study explores how the allocated budget enhances the success probability in such settings.

In [22], we study online selection problems in both the prophet and secretary settings, when arriving agents have interdependent values. In the interdependent values model, introduced in the seminal work of Milgrom and Weber [1982], each agent has a private signal and the value of an agent is a function of the signals held by all agents. Results in online selection crucially rely on some degree of independence of values, which is conceptually at odds with the interdependent values model. For prophet and secretary models under the standard independent values assumption, prior works provide constant factor approximations to the welfare. On the other hand, when agents have interdependent values, prior works in Economics and Computer Science provide truthful mechanisms that obtain optimal and approximately optimal welfare under certain assumptions on the valuation functions. We bring together these two important lines of work and provide the first constant factor approximations for prophet and secretary problems with interdependent values. We consider both the algorithmic setting, where agents are non-strategic (but have interdependent values), and the mechanism design setting with strategic agents. All our results are constructive and use simple stopping rules.

6.4.2 Algorithms with prediction and scheduling

In [15], we address the problem of learning-augmented priority queue. Priority queues are one of the most fundamental and widely used data structures in computer science. Their primary objective is to efficiently support the insertion of new elements with assigned priorities and the extraction of the highest priority element. In this study, we investigate the design of priority queues within the learning-augmented framework, where algorithms use potentially inaccurate predictions to enhance their worst-case performance. We examine three prediction models spanning different use cases, and we show how the predictions can be leveraged to enhance the performance of priority queue operations. Moreover, we demonstrate the optimality of our solution and discuss some possible applications.

In [16], we consider the non-clairvoyant scheduling problem. The non-clairvoyant scheduling problem has gained new interest within learning-augmented algorithms, where the decision-maker is equipped with predictions without any quality guarantees. In practical settings, access to predictions may be reduced to specific instances, due to cost or data limitations. Our investigation focuses on scenarios where predictions for only B job sizes out of n are available to the algorithm. We first establish near-optimal lower bounds and algorithms in the case of perfect predictions. Subsequently, we present a learning-augmented algorithm satisfying the robustness, consistency, and smoothness criteria, and revealing a novel tradeoff between consistency and smoothness inherent in the scenario with a restricted number of predictions.

In [21], we consider the question: How should an expert send forecasts to maximize her utility subject to passing a calibration test? We consider a dynamic game where an expert sends probabilistic forecasts to a decision maker. The decision maker uses a calibration test based on past outcomes to verify the expert's forecasts. We characterize the optimal forecasting strategy by reducing the dynamic game to a static persuasion problem. A distribution of forecasts is implementable by a calibrated strategy if and only if it is a mean-preserving contraction of the distribution of conditionals (honest forecasts). We characterize the value of information by comparing what an informed and uninformed expert can attain. Moreover, we consider a decision maker who uses regret minimization, instead of the calibration test, to take actions. We show that the expert can achieve the same payoff against a regret minimizer as under the calibration test, and in some instances, she can achieve strictly more.

In [10], we consider algorithms for scheduling deadline-sensitive malleable tasks. Due to the ubiquity of batch data processing, the related problems of scheduling malleable batch tasks have received significant attention. We consider a fundamental model where a set of tasks is to be processed on multiple identical machines and each task is specified by a value, a workload, a deadline and a parallelism bound. Within the parallelism bound, the number of machines assigned to a task can vary over time without affecting its workload. In this paper, we identify a boundary condition and prove by construction that a set of malleable tasks with deadlines can be finished by their deadlines if and only if it satisfies the boundary condition. This core result plays a key role in the design and analysis of scheduling algorithms: (i) when several typical objectives are considered, such as social welfare maximization, machine minimization, and minimizing the maximum weighted completion time, and, (ii) when the algorithmic design techniques such as greedy and dynamic programming are applied to the social welfare maximization problem. As a result, we give four new or improved algorithms for the above problems.

7 Partnerships and cooperations

7.1 International research visitors

7.1.1 Visits of international scientists

Other international visits to the team

H. Stein

Status: PhD

Institution of origin: Heidelberg University

Country: Germany

Dates: May (1week), 2024

Context of the visit: Research stay

J. Correa

Status: Professor

Institution of origin: Universidad de Chile

Country: Chile

Dates: November (1week), 2024

Context of the visit: Research stay

7.1.2 Visits to international teams

Research stays abroad

Cristina Butucea

Visited institution: Rostock University

Country: Germany

Dates: September 1-5, 2024

Context of the visit: Research, discussions

Visited institution: Freiburg University

Country: Germany

Dates: March (1 week), 2024

Context of the visit: Research, discussions

7.2 National initiatives

Foundry (PEPR IA)

Participants: Patrick Loiseau.

Title: Foundry: Foundation of robustness and reliability in AI

Partner Institution(s):

- Inria
- CNRS
- Université Paris Dauphine
- Institut Mines Telecom
- ENS de Lyon

Date/Duration: 2023-2027 (4 years)

Additional info/keywords: PEPR IA projet cible, 245k euros. Fairness, matching, auctions.

FairPlay (ANR JCJC)

Participants: Patrick Loiseau.

Title: FairPlay: Fair algorithms via game theory and sequential learning

Partner Institution(s):

- Inria

Date/Duration: 2021-2025 (4 years)

Additional info/keywords: ANR JCJC project, 245k euros. Fairness, matching, auctions.

BOLD (ANR)

Participants: Vianney Perchet.

Title: BOLD: Beyond Online Learning for Better Decisions

Partner Institution(s):

- Crest, Genes

Date/Duration: 2019-2024 (4.5 years)

Additional info/keywords: ANR project, 270k euros. online learning, optimization, bandits.

DOOM (ANR)

Participants: Vianney Perchet.

Title: DOOM: Design of Optimal Online Matching Markets

Partner Institution(s): • Crest, Genes

Date/Duration: 2014-2028 (4 years)

Additional info/keywords: ANR project, 456k euros. online markets, learning, matching.

8 Dissemination**8.1 Promoting scientific activities****8.1.1 Scientific events: organisation****General chair, scientific chair**

Participants: Cristina Butucea.

Title: Asymptotically equivalent representations of non-stationary processes.

Partner Institution(s): • Crest, Genes

Date/Duration: March 4-15

Location: CIRM, Marseille

Member of the organizing committees

Participants: Vianney Perchet.

Title: Member of the scientific advisory committee of the Hi! Paris summer school

Date/Duration: July (1 week)

Location: HEC campus

8.1.2 Scientific events: selection**Chair of conference program committees**

Participants: Patrick Loiseau.

Title: Program chair of the Hi! Paris summer school

Date/Duration: July (1 week)

Location: HEC campus

Member of the conference program committees**Patrick Loiseau:** NeurIPS, ECML-PKDD (Area Chair)**Vianney Perchet:** NeurIPS, ICLR, ICML, COLT, ALT**Simon Mauras:** WINE, EC.**Reviewer****Hugo Richard:** NeurIPS, ICML, AISTATS**Clément Calauzènes:** ICML, NeurIPS**Benjamin Heymann:** AISTATS**Maxime Vono:** NeurIPS**Simon Mauras:** ESA, SODA, STOC**8.1.3 Journal****Member of the editorial boards****Vianney Perchet:** Foundations and Trends in Machine Learning, Operation Research, Operation Research Letters, Journal of Machine Learning Research, Journal of Dynamics and Games,**Cristina Butucea:** Annals of Statistics, Bernoulli**Reviewer - reviewing activities****Patrick Loiseau:** Dynamic games and applications, Operation Research letters**Vianney Perchet:** Annals of Statistics, Mathematics of Operation Research, Journal of the ACM**Cristina Butucea:** Annals of Statistics, Bernoulli, IEEE Transactions on Information Theory**Matthieu Lerasle:** Annals of statistics, Journal of the European Mathematical Society, Probability and Related Fields, Journal of Machine Learning Research, Journal of the American Statistical Association.**Marc Abeille:** Journal of Machine Learning Research**Maxime Vono:** Journal of Computational and Graphical Statistics**Benjamin Heymann:** SIAM Control and Optimization, Mathematical Reviews**Simon Mauras:** Annals of Mathematics and Artificial Intelligence, Mathematics of Operations Research, Operations Research Letters, Games and Economic Behavior**8.1.4 Invited talks****Cristina Butucea:** Anniversary Conference V. Koltchinskii (Bertinoro, Italy), ISNP (Braga, Portugal), MFOberwolfach Germany, Journées de la SMAI (Poitiers), University of Nanjing (China)**Hugo Richard:** PGMO Days (France)**Simon Mauras:** Junior Workshop on Mathematical Game Theory (Rome), Algorithms, Learning, and Games (Sicily), Alpine Game Theory Symposium (Grenoble), Game theory seminar (Paris), Optimization seminar (Paris)

8.1.5 Scientific expertise

Vianney Perchet: Expert for the evaluation of the LABEX MME:DII

Patrick Loiseau: External expert for the ANR and the Royal Society

Cristina Butucea: ANR evaluation committee, hiring committees France and Germany,

8.2 Teaching - Supervision - Juries

8.2.1 Supervision

Patrick Loiseau: PhD students: Rémi Castera, Mathieu Molina, Reda Jalal, Minrui Xu, Marie Generali, Melissa Tamine; postdocs: Felipe Garrido Lucero, Simon Finster, Denis Sokolov

Vianney Perchet: PhD students: Come Fiegel, Maria Cherifa, Mathieu Molina, Ziyad Benomar, Mike Liu, Hafedh El Ferchichi, Matilde Tullii, Giovanni Montanari. postdocs: Felipe Garrido Lucero, Nadav Merlis, Solenne Gaucher, Dorian Baudry

Matthieu Lerasle PhD Students: Clara Carlier, Hugo Chardon, Hafedh El Ferchichi.

Cristina Butucea PhD students: Nayel Bettache, Henning Stein, Antoine Schoonaert

Marc Abeille: Ahmed Ben Yahmed

Clément Calauzènes: Morgane Goibert, Maria Cherifa

Benjamin Heymann: Mélissa Tamine

Maxime Vono: Mélissa Tamine

Simon Mauras: PhD student: Minrui Xu.

8.3 Teaching

ENSAE:

Advanced Optimization Third year, lectures

Theoretical Foundations of Machine Learning Second year, lectures

Stopping time and online algorithms Third year, lectures

Statistics (ML) 1st and second year

Nonparametric Statistics 3rd year, M2

Mathematical Foundations of Probabilities 1st year

Ecole Polytechnique:

INF421: design and analysis of algorithms (Patrick Loiseau and Simon Mauras). Second-year level, PCs.

INF581: Advanced Machine Learning and Autonomous Agents (Patrick Loiseau). Third-year/M1 level, lectures and labs.

MAP433: Statistics (ML). First-year cycle polytechnicien, PCs.

MAP576: Learning Theory (ML). Second-year cycle polytechnicien, Lecture.

INF471S: Algorithms and Advanced Programming (Simon Mauras) Bachelor and cycle polytechnicien, lectures and tutorials.

Université Paris-Saclay: High Dimensional Probability (ML). Master 2

Stopping Time and Random Algorithm (ML). Master 2

PSL: Introduction to machine learning (Hugo Richard). L3 level, Lectures and labs.

Master IASD: Recommender Systems (Clément Calauzènes). Master 2, Lectures.

8.4 Popularization

8.4.1 Specific official responsibilities in science outreach structures

- Simon Mauras:**
- Future co-responsible médiation at INRIA Saclay.
 - Volunteer and treasurer of association France-IOI: training and selecting the french team for the International Olympiads in Informatics; organizing summer camps to teach algorithmic to groups of ~20 highschool students, creating online contests and educational content about informatics (Concours castor, concours Alkindi, Concours Algoréa, online content for SNT option in highschool)

9 Scientific production

9.1 Major publications

- [1] C. Fiegel, P. Ménard, T. Kozuno, R. Munos, V. Perchet and M. Valko. ‘Local and adaptive mirror descents in extensive-form games’. In: 2023 International Conference on Machine Learning. New Orleans, United States, 1st Sept. 2023. URL: <https://hal.science/hal-04416177>.
- [2] M. Molina, N. Gast, P. Loiseau and V. Perchet. ‘Trading-off price for data quality to achieve fair online allocation’. In: NeurIPS 2023 - 37th Conference on Neural Information Processing Systems. New orleans, USA, United States, 2023, pp. 1–43. URL: <https://inria.hal.science/hal-04360656>.
- [3] N. Verzelen, M. Fromont, M. Lerasle and P. Reynaud-Bouret. ‘Optimal Change-Point Detection and Localization’. In: *Annals of Statistics* 51.4 (2023), pp. 1586–1610. DOI: [10.1214/23-AOS2297](https://doi.org/10.1214/23-AOS2297). URL: <https://hal.science/hal-04251607>.

9.2 Publications of the year

International journals

- [4] M. Bompaire, A. Désir and B. Heymann. ‘Fixed Point Label Attribution for Real-Time Bidding’. In: *Manufacturing and Service Operations Management* 26.3 (May 2024), pp. 1043–1061. DOI: [10.1287/msom.2021.0611](https://doi.org/10.1287/msom.2021.0611). URL: <https://hal.science/hal-04722331> (cit. on p. 15).
- [5] E. Boursier and V. Perchet. ‘A survey on multi-player bandits’. In: *Journal of Machine Learning Research* (Jan. 2024). URL: <https://inria.hal.science/hal-03941302> (cit. on p. 18).
- [6] M. Chessa and P. Loiseau. ‘Enhancing voluntary contributions in a public goods economy via a minimum individual contribution level’. In: *Public Choice* 201.1-2 (12th Apr. 2024), pp. 237–261. DOI: [10.1007/s11127-024-01165-1](https://doi.org/10.1007/s11127-024-01165-1). URL: <https://inria.hal.science/hal-04943298> (cit. on p. 16).
- [7] C. Giraud, Y. Issartel, L. Lehericy and M. Lerasle. ‘Pair-Matching: Link Prediction with Adaptive Queries’. In: *Mathematical Statistics and Learning* (5th Mar. 2024). URL: <https://hal.science/hal-04578273> (cit. on p. 17).
- [8] B. Heymann and A. Jofré. ‘Procurement auctions with losses’. In: *Computational Management Science* 21.2 (9th July 2024), p. 38. DOI: [10.1007/s10287-024-00513-2](https://doi.org/10.1007/s10287-024-00513-2). URL: <https://hal.science/hal-04722332> (cit. on p. 15).
- [9] B. Heymann and A. Jofré. ‘Reverse auctions with transportation and convex costs’. In: *Computational Management Science* 21.1 (5th Feb. 2024), p. 20. DOI: [10.1007/s10287-023-00498-4](https://doi.org/10.1007/s10287-023-00498-4). URL: <https://hal.science/hal-04442630> (cit. on p. 15).
- [10] X. Wu and P. Loiseau. ‘Algorithms for Scheduling Deadline-Sensitive Malleable Tasks’. In: *SN Operations Research Forum* 5.2 (1st Apr. 2024), p. 30. DOI: [10.1007/s43069-024-00300-4](https://doi.org/10.1007/s43069-024-00300-4). URL: <https://inria.hal.science/hal-04589822> (cit. on p. 20).

International peer-reviewed conferences

- [11] D. Baudry, H. Richard, M. Cherifa, C. Calauzènes and V. Perchet. ‘Optimizing the coalition gain in Online Auctions with Greedy Structured Bandits’. In: *NeurIPS 2024*. NeurIPS 2024 - 38th Conference on Neural Information Processing Systems. Vancouver (BC), Canada, 10th Dec. 2024. URL: <https://hal.science/hal-04896482> (cit. on p. 15).
- [12] A. Ben Yahmed, C. Calauzènes and V. Perchet. ‘Strategic Multi-Armed Bandit Problems Under Debt-Free Reporting’. In: *NeurIPS 2024 - 38th Annual Conference on Neural Information Processing Systems*. Vancouver, Canada, 10th Dec. 2024. URL: <https://hal.science/hal-04912557> (cit. on p. 18).
- [13] Z. Benomar, D. Baudry and V. Perchet. ‘Lookback Prophet Inequalities’. In: *NeurIPS 2024 - 38th Conference on Neural Information Processing Systems*. Vancouver (CA), Canada, 10th Dec. 2024. URL: <https://hal.science/hal-04854065> (cit. on p. 18).
- [14] Z. Benomar, E. Chzhen, N. Schreuder and V. Perchet. ‘Addressing Bias in Online Selection with Limited Budget of Comparisons’. In: *PMLR*. NeurIPS 2024 - 38th Conference on Neural Information Processing Systems. Vancouver (CA), Canada, 10th Dec. 2024. URL: <https://hal.science/hal-04854073> (cit. on p. 19).
- [15] Z. Benomar and C. Coester. ‘Learning-Augmented Priority Queues’. In: *NeurIPS 2024 - 38th Conference on Neural Information Processing Systems*. Vancouver (CA), Canada, 10th Dec. 2024. URL: <https://hal.science/hal-04854076> (cit. on p. 19).
- [16] Z. Benomar and V. Perchet. ‘Non-clairvoyant Scheduling with Partial Predictions’. In: *Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria*. PMLR 235, 2024. ICML 2024 - The 41st International Conference on Machine Learning. Vienna, Austria, 21st July 2024. URL: <https://hal.science/hal-04709576> (cit. on p. 19).
- [17] A. Eden, M. Feldman, S. Mauras and D. Mohan. ‘Private Interdependent Valuations: New Bounds for Single-Item Auctions and Matroids’. In: *EC '24: Proceedings of the 25th ACM Conference on Economics and Computation*. EC 2024 - 25th ACM Conference on Economics and Computation. New Haven, United States: ACM, 8th July 2024, pp. 448–464. DOI: [10.1145/3670865.3673581](https://doi.org/10.1145/3670865.3673581). URL: <https://hal.science/hal-04923236> (cit. on p. 16).
- [18] H. El Ferchichi, M. Lerasle and V. Perchet. ‘Active Ranking and Matchmaking, with Perfect Matchings’. In: *Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria*. PMLR 235, 2024. ICML 2024 - The Forty-First International Conference on Machine Learning. Vienne, Austria, 2nd May 2024. URL: <https://hal.science/hal-04891449> (cit. on p. 17).
- [19] M. Feldman, S. Mauras, V. V. Narayan and T. Ponitka. ‘Breaking the Envy Cycle: Best-of-Both-Worlds Guarantees for Subadditive Valuations’. In: *EC '24: Proceedings of the 25th ACM Conference on Economics and Computation*. EC 2024 - 25th ACM Conference on Economics and Computation. New Haven, United States: ACM, 8th July 2024, pp. 1236–1266. DOI: [10.1145/3670865.3673592](https://doi.org/10.1145/3670865.3673592). URL: <https://hal.science/hal-04923223> (cit. on p. 16).
- [20] F. Garrido-Lucero, B. Heymann, M. Vono, P. Loiseau and V. Perchet. ‘DU-Shapley: A Shapley Value Proxy for Efficient Dataset Valuation’. In: *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*. Advances in Neural Information Processing Systems 37 (NeurIPS 2024). Vancouver, Canada, 2024. URL: <https://hal.science/hal-04420286> (cit. on p. 17).
- [21] A. Jain and V. Perchet. ‘Calibrated Forecasting and Persuasion’. In: *EC 2024 - The Twenty-Fifth ACM Conference on Economics and Computation*. New haven, United States, 21st June 2024. URL: <https://hal.science/hal-04912547> (cit. on p. 19).
- [22] S. Mauras, D. Mohan and R. Reiffenhäuser. ‘Optimal Stopping with Interdependent Values’. In: *EC '24: Proceedings of the 25th ACM Conference on Economics and Computation*. EC 2024 - 25th ACM Conference on Economics and Computation. New Haven, United States: ACM, 8th July 2024, pp. 246–265. DOI: [10.1145/3670865.3673628](https://doi.org/10.1145/3670865.3673628). URL: <https://hal.science/hal-04923199> (cit. on p. 19).

- [23] N. Merlis. ‘Reinforcement Learning with Lookahead Information’. In: NeurIPS 2024 - 38th Conference on Neural Information Processing Systems. Vancouver, Canada, 10th Dec. 2024. URL: <https://hal.science/hal-04894082> (cit. on p. 18).
- [24] N. Merlis, D. Baudry and V. Perchet. ‘The Value of Reward Lookahead in Reinforcement Learning’. In: NeurIPS 2024 - 38th Conference on Neural Information Processing Systems. Vancouver, Canada, 10th Dec. 2024. URL: <https://hal.science/hal-04893988> (cit. on p. 18).
- [25] M. Potfer, D. Baudry, H. Richard, V. Perchet and C. Wang. ‘Improved learning rates in multi-unit uniform price auctions’. In: *NeurIPS 2024*. NeurIPS 2024 - 38th Conference on Neural Information Processing Systems. Vancouver (BC), Canada, 10th Dec. 2024. URL: <https://hal.science/hal-04896481> (cit. on p. 16).
- [26] M. Tullii, S. Gaucher, N. Merlis and V. Perchet. ‘Improved Algorithms for Contextual Dynamic Pricing’. In: *NeurIPS 2024 - 38th Conference on Neural Information Processing Systems*. NeurIPS 2024 - The 38th Annual Conference on Neural Information Processing Systems. Vancouver (BC), Canada: arXiv, 2024. URL: <https://hal.science/hal-04887635> (cit. on p. 16).

Doctoral dissertations and habilitation theses

- [27] R. Castera. ‘Efficiency and Fairness in Matching Problems’. UGA (Université Grenoble Alpes), 16th Oct. 2024. URL: <https://hal.science/tel-04874485>.

Reports & preprints

- [28] C. Arnal, V. Cabannes and V. Perchet. *Mode Estimation with Partial Feedback*. 2024. DOI: 10.48550/arXiv.2402.13079. URL: <https://inria.hal.science/hal-04471461>.
- [29] D. Baudry, N. Merlis, M. Molina, H. Richard and V. Perchet. *Multi-Armed Bandits with Guaranteed Revenue per Arm*. 26th Jan. 2024. URL: <https://hal.science/hal-04420013>.
- [30] N. Bettache and C. Butucea. *Two-sided Matrix Regression*. 26th Jan. 2024. URL: <https://hal.science/hal-04419650>.
- [31] C. Butucea, J.-F. Delmas, A. Dutfoy and C. Hardy. *Off-the-grid prediction and testing for linear combination of translated features*. 18th July 2024. URL: <https://hal.science/hal-03880134>.
- [32] C. Butucea, J.-F. Delmas, A. Dutfoy and C. Hardy. *Simultaneous off-the-grid learning of mixtures issued from a continuous dictionary*. 21st Feb. 2024. URL: <https://hal.science/hal-03831208>.
- [33] R. Castera, P. Loiseau and B. Pradelski. *Correlation of Rankings in Matching Markets*. Feb. 2024. URL: <https://hal.science/hal-03672270>.
- [34] V. Divol and S. Gaucher. *Demographic parity in regression and classification within the unawareness framework*. 3rd Sept. 2024. URL: <https://hal.science/hal-04684789>.
- [35] S. Gaucher, G. Blanchard and F. Chazal. *Supervised Contamination Detection, with Flow Cytometry Application*. 5th Apr. 2024. URL: <https://hal.science/hal-04535142>.
- [36] B. Heymann, R. Chan-Renous-Legoubin and A. Gilotte. *A pragmatic policy learning approach to account for users’ fatigue in repeated auctions*. 12th July 2024. URL: <https://hal.science/hal-04646638>.
- [37] B. Heymann and M. Lanctot. *Learning in Games with progressive hiding*. 5th Sept. 2024. URL: <https://hal.science/hal-04728740>.
- [38] M. Tullii, S. Gaucher, H. Richard, E. Diemert, V. Perchet, A. Rakotomamonjy, C. Calauzènes and M. Vono. *Position Paper: Open Research Challenges for Private Advertising Systems under Local Differential Privacy*. 5th Feb. 2024. URL: <https://hal.science/hal-04438186>.

9.3 Cited publications

- [39] ‘24 CFR § 100.75 - Discriminatory advertisements, statements and notices.’ <https://www.law.cornell.edu/cfr/text/24/100.75> (cit. on p. 5).
- [40] E. Abbe. ‘Community detection and stochastic block models: recent developments’. In: *The Journal of Machine Learning Research* 18.1 (2017), pp. 6446–6531 (cit. on p. 12).
- [41] M. Ali, P. Sapiezynski, M. Bogen, A. Korolova, A. Mislove and A. Rieke. ‘Discrimination through optimization: How Facebook’s ad delivery can lead to skewed outcomes’. In: CSCW. 2019 (cit. on p. 5).
- [42] N. Alon, N. Cesa-Bianchi, O. Dekel and T. Koren. ‘Online learning with feedback graphs: Beyond bandits’. In: *Annual Conference on Learning Theory*. Vol. 40. Microtome Publishing. 2015 (cit. on p. 12).
- [43] I. Ashlagi, M. Burq, P. Jaillet and V. Manshadi. ‘On matching and thickness in heterogeneous dynamic markets’. In: *Operations Research* 67.4 (2019), pp. 927–949 (cit. on p. 11).
- [44] I. Ashlagi, P. Jaillet and V. H. Manshadi. *Kidney Exchange in Dynamic Sparse Heterogenous Pools*. 2013. arXiv: [1301.3509](https://arxiv.org/abs/1301.3509) [cs.DS] (cit. on p. 11).
- [45] P. Awasthi and T. Sandholm. ‘Online stochastic optimization in the large: Application to kidney exchange’. In: *Twenty-First International Joint Conference on Artificial Intelligence*. 2009 (cit. on p. 11).
- [46] B. Bahmani and M. Kapralov. ‘Improved bounds for online stochastic matching’. In: *European Symposium on Algorithms*. Springer. 2010, pp. 170–181 (cit. on p. 11).
- [47] E. del Barrio, F. Gamboa, P. Gordaliza and J.-M. Loubes. ‘Obtaining fairness using optimal transport theory’. In: *arXiv:1806.03195*. 2018, pp. 1–25 (cit. on p. 7).
- [48] B. Birnbaum and C. Mathieu. ‘On-line bipartite matching made simple’. In: *Acm Sigact News* 39.1 (2008), pp. 80–87 (cit. on p. 11).
- [49] B. Bollobas and G. Brightwell. ‘The width of random graph orders’. In: *The Mathematical Scientist* 20 (Jan. 1995) (cit. on p. 12).
- [50] C. Bordenave, M. Lelarge and J. Salez. *Matchings on infinite graphs*. 2011. arXiv: [1102.0712](https://arxiv.org/abs/1102.0712) [math.PR] (cit. on p. 12).
- [51] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998 (cit. on p. 6).
- [52] A. Borodin, C. Karavasilis and D. Pankratov. *An Experimental Study of Algorithms for Online Bipartite Matching*. 2018. arXiv: [1808.04863](https://arxiv.org/abs/1808.04863) [cs.DS] (cit. on p. 11).
- [53] E. Boursier and V. Perchet. ‘SIC-MMAB: Synchronisation Involves Communication in Multiplayer Multi-Armed Bandits’. In: *arXiv:1809.08151*. 2018, pp. 1–31 (cit. on p. 6).
- [54] E. Boursier and V. Perchet. ‘Utility/Privacy Trade-off through the lens of Optimal Transport’. In: *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*. Ed. by S. Chiappa and R. Calandra. Vol. 108. Proceedings of Machine Learning Research. Online: PMLR, Aug. 2020, pp. 591–601 (cit. on pp. 7, 11).
- [55] S. Bubeck and N. Cesa-Bianchi. ‘Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems’. In: *Machine Learning* 5.1 (2012), pp. 1–122 (cit. on p. 6).
- [56] S. Bubeck, V. Perchet and P. Rigollet. ‘Bounded regret in stochastic multi-armed bandits’. In: *Journal of Machine Learning Research: Workshop and Conference Proceedings (COLT)* 30 (2013), pp. 122–134 (cit. on p. 6).
- [57] N. Buchbinder, K. Jain and J. (Naor. ‘Online Primal-Dual Algorithms for Maximizing Ad-Auctions Revenue’. In: ed. by L. Arge, M. Hoffmann and E. Welzl. Berlin, Heidelberg, 2007, pp. 253–264 (cit. on p. 11).
- [58] C. Butucea, A. Dubois, M. Kroll, A. Saumard et al. ‘Local differential privacy: Elbow effect in optimal density estimation and adaptation over Besov ellipsoids’. In: *Bernoulli* 26.3 (2020), pp. 1727–1764 (cit. on p. 5).

- [59] C. Butucea, A. Rohde and L. Steinberger. ‘Interactive versus non-interactive locally, differentially private estimation: Two elbows for the quadratic functional’. In: *Annals of Stats* (2023). to appear (cit. on p. 5).
- [60] L. E. Celis, A. Mehrotra and N. K. Vishnoi. ‘Toward Controlling Discrimination in Online Ad Auctions’. In: *ICML*. 2019 (cit. on pp. 8, 13).
- [61] N. Cesa-Bianchi, C. Gentile and Y. Mansour. ‘Regret minimization for reserve prices in second-price auctions’. In: *Proceedings of SODA 2013*. 2013 (cit. on p. 12).
- [62] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006 (cit. on p. 6).
- [63] K. Chaudhuri, J. Imola and A. Machanavajjhala. ‘Capacity bounded differential privacy’. In: *Advances in Neural Information Processing Systems*. 2019, pp. 3469–3478 (cit. on p. 5).
- [64] S. Chawla and M. Jagadeesan. *Fairness in ad auctions through inverse proportionality*. arXiv:2003.13966. 2020 (cit. on pp. 8, 13).
- [65] I. R. Cohen and D. Wajc. ‘Randomized online matching in regular graphs’. In: *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2018, pp. 960–979 (cit. on p. 11).
- [66] V. Conitzer, C. Kroer, D. Panigrahi, O. Schrijvers, E. Sodomka, N. E. Stier-Moses and C. Wilkens. ‘Pacing equilibrium in first-price auction markets’. In: *EC*. 2019 (cit. on p. 8).
- [67] V. Conitzer, C. Kroer, E. Sodomka and N. E. Stier-Moses. ‘Multiplicative pacing equilibria in auction markets’. In: *WINE*. 2018 (cit. on p. 8).
- [68] J. Correa, P. Dütting, F. Fischer and K. Schewior. ‘Prophet inequalities for iid random variables from an unknown distribution’. In: *Proceedings of the 2019 ACM Conference on Economics and Computation*. 2019, pp. 3–17 (cit. on p. 6).
- [69] R. Cummings, V. Gupta, D. Kimpara and J. Morgenstern. ‘On the Compatibility of Privacy and Fairness’. In: *FairUMAP*. 2019 (cit. on p. 10).
- [70] J. Dastin. *Amazon scraps secret AI recruiting tool that showed bias against women*. Reuters, <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>. 2018 (cit. on p. 5).
- [71] C. Dwork. ‘Differential privacy’. In: *Encyclopedia of Cryptography and Security* (2011), pp. 338–340 (cit. on p. 5).
- [72] C. Dwork, M. Hardt, T. Pitassi, O. Reingold and R. Zemel. ‘Fairness through awareness’. In: *ITCS*. 2012 (cit. on p. 5).
- [73] C. Dwork, F. McSherry, K. Nissim and A. Smith. ‘Calibrating noise to sensitivity in private data analysis’. In: *Theory of cryptography conference*. Springer. 2006, pp. 265–284 (cit. on p. 5).
- [74] R. Eilat, K. Eliaz and X. Mu. *Optimal Privacy-Constrained Mechanisms*. Tech. rep. C.E.P.R. Discussion Papers, 2019 (cit. on p. 5).
- [75] V. Emelianov, N. Gast, K. P. Gummadi and P. Loiseau. ‘On the Effect of Positive Discrimination on Multistage Selection Problems in the Presence of Implicit Variance’. In: *EC*. 2020 (cit. on p. 9).
- [76] J. Feldman, A. Mehta, V. Mirrokni and S. Muthukrishnan. *Online Stochastic Matching: Beating 1-1/e*. 2009. arXiv: 0905.4100 [cs.DS] (cit. on p. 11).
- [77] K. Ferryman and M. Pitcan. *Fairness in Precision Medicine*. Tech. rep. Data & Society, 2018 (cit. on p. 5).
- [78] D. Fudenberg and J. Tirole. *Game Theory*. MIT press, 1991 (cit. on p. 6).
- [79] E. Garcelon, V. Perchet, C. Pike-Burke and M. Pirotta. ‘Local Differentially Private Regret Minimization in Reinforcement Learning’. In: *arXiv preprint arXiv:2010.07778* (2020) (cit. on p. 5).
- [80] N. Gast, S. Ioannidis, P. Loiseau and B. Roussillon. ‘Linear Regression from Strategic Data Sources’. In: *ACM Transactions on Economics and Computation* 8.2 (May 2020), 10:1–10:24 (cit. on p. 11).

- [81] N. Gast and B. Van Houdt. ‘A Refined Mean Field Approximation’. In: *SIGMETRICS*. 2017 (cit. on p. 8).
- [82] M. Gupta, A. Cotter, M. M. Fard and S. Wang. ‘Proxy fairness’. In: *arXiv preprint arXiv:1806.11212* (2018) (cit. on p. 10).
- [83] A. Hannák, C. Wagner, D. Garcia, A. Mislove, M. Strohmaier and C. Wilson. ‘Bias in online freelance marketplaces: Evidence from taskrabbit and fiverr’. In: *CSCW*. 2017 (cit. on p. 5).
- [84] M. Hardt, E. Price and N. Srebro. ‘Equality of Opportunity in Supervised Learning’. In: *NIPS*. 2016 (cit. on p. 5).
- [85] C.-J. Ho and J. W. Vaughan. ‘Online task assignment in crowdsourcing markets’. In: *Twenty-sixth AAAI conference on artificial intelligence*. 2012 (cit. on p. 11).
- [86] C. Ilvento, M. Jagadeesan and S. Chawla. ‘Multi-Category Fairness in Sponsored Search Auctions’. In: *FAT**. 2020 (cit. on pp. 8, 13).
- [87] K. Iyer, R. Johari and M. Sundararajan. ‘Mean field equilibria of dynamic auctions with learning’. In: *Management Science* 60.12 (2014), pp. 2949–2970 (cit. on p. 8).
- [88] P. Jaillet and X. Lu. ‘Online stochastic matching: New algorithms with better bounds’. In: *Mathematics of Operations Research* 39.3 (2014), pp. 624–646 (cit. on p. 11).
- [89] J. Josse, N. Prost, E. Scornet and G. Varoquaux. ‘On the consistency of supervised learning with missing values’. In: *arXiv preprint arXiv:1902.06931* (2019) (cit. on p. 10).
- [90] N. Kallus, X. Mao and A. Zhou. ‘Assessing algorithmic fairness with unobserved protected class using data combination’. In: *arXiv preprint arXiv:1906.00285* (2019) (cit. on p. 10).
- [91] C. Karande, A. Mehta and P. Tripathi. ‘Online bipartite matching with unknown distributions’. In: *Proceedings of the forty-third annual ACM symposium on Theory of computing*. 2011, pp. 587–596 (cit. on p. 11).
- [92] R. M. Karp, U. V. Vazirani and V. V. Vazirani. ‘An optimal algorithm for on-line bipartite matching’. In: *Proceedings of the twenty-second annual ACM symposium on Theory of computing*. 1990, pp. 352–358 (cit. on p. 11).
- [93] N. Kilbertus, M. R. Carulla, G. Parascandolo, M. Hardt, D. Janzing and B. Schölkopf. ‘Avoiding discrimination through causal reasoning’. In: *Advances in Neural Information Processing Systems*. 2017, pp. 656–666 (cit. on p. 11).
- [94] J. S. Kim, J. Chen and A. Talwalkar. ‘Model-Agnostic Characterization of Fairness Trade-offs’. In: *arXiv preprint arXiv:2004.03424* (2020) (cit. on pp. 10, 14).
- [95] D. Knuth. *Stable Marriage and Its Relation to Other Combinatorial Problems: An Introduction to the Mathematical Analysis of Algorithms*. English translation, (CRM Proceedings and Lecture Notes), American Mathematical Society, 1996 (cit. on p. 13).
- [96] V. Krishna. *Auction Theory*. 1st ed. Elsevier, 2009 (cit. on p. 13).
- [97] M. J. Kusner, J. Loftus, C. Russell and R. Silva. ‘Counterfactual fairness’. In: *Advances in neural information processing systems*. 2017, pp. 4066–4076 (cit. on p. 10).
- [98] M. J. Kusner and J. R. Loftus. *The long road to fairer algorithms*. 2020 (cit. on p. 10).
- [99] A. Lambrecht and C. Tucker. ‘Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads’. In: *Management Science* (2019) (cit. on p. 5).
- [100] J. Larson, S. Mattu, L. Kirchner and J. Angwin. *How We Analyzed the COMPAS Recidivism Algorithm*. ProPublica, <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>. 2016 (cit. on p. 5).
- [101] L. Liu, H. Mania and M. I. Jordan. ‘Competing Bandits in Matching Markets’. In: *arXiv:1906.05363*. 2019, pp. 1–15 (cit. on p. 6).
- [102] B. Maćkowiak and M. Wiederholt. ‘Business cycle dynamics under rational inattention’. In: *The Review of Economic Studies* 82.4 (2015), pp. 1502–1532 (cit. on p. 5).

- [103] M. Mahdian and Q. Yan. ‘Online bipartite matching with random arrivals: an approach based on strongly factor-revealing lps’. In: *Proceedings of the forty-third annual ACM symposium on Theory of computing*. 2011, pp. 597–606 (cit. on p. 11).
- [104] V. H. Manshadi, S. O. Gharan and A. Saberi. ‘Online stochastic matching: Online actions based on offline statistics’. In: *Mathematics of Operations Research* 37.4 (2012), pp. 559–573 (cit. on p. 11).
- [105] A. Mas-Colell, M. Whinston and J. Green. *Microeconomic Theory*. Oxford University Press, 1995 (cit. on p. 4).
- [106] A. Mastin and P. Jaillet. ‘Greedy online bipartite matching on random graphs’. In: *arXiv preprint arXiv:1307.2536* (2013) (cit. on p. 11).
- [107] F. Matějka and A. McKay. ‘Rational inattention to discrete choices: A new foundation for the multinomial logit model’. In: *American Economic Review* 105.1 (2015), pp. 272–98 (cit. on p. 5).
- [108] A. Mehta. ‘Online Matching and Ad Allocation’. In: *Found. Trends Theor. Comput. Sci.* 8.4 (Oct. 2013), pp. 265–368. DOI: [10.1561/04000000057](https://doi.org/10.1561/04000000057). URL: <https://doi.org/10.1561/04000000057> (cit. on p. 11).
- [109] J.-F. Mertens, S. Sorin and S. Zamir. *Repeated Games*. Econometric Society Monographs. Cambridge University Press, 2015 (cit. on p. 6).
- [110] J. Munkres. ‘Algorithms for the Assignment and Transportation Problems’. In: *Journal of the Society for Industrial and Applied Mathematics* 5.1 (1957), pp. 32–38 (cit. on p. 13).
- [111] M. Nasr and M. C. Tschantz. ‘Bidding Strategies with Gender Nondiscrimination Constraints for Online Ad Auctions’. In: *FAT**. 2020 (cit. on p. 13, 14).
- [112] T. Nedelec, M. Abeille, C. Calauzenes, B. Heymann, V. Perchet and N. El Karoui. ‘The bidder’s standpoint : a simple way to improve bidding strategies in revenue-maximizing auctions’. In: *Workshop on Learning in the Presence of Strategic Behavior (EC 2019)*. 2019 (cit. on p. 13).
- [113] T. Nedelec, N. El Karoui and V. Perchet. ‘Learning to bid in revenue-maximizing auctions’. In: *Proceedings of the 36th International Conference on Machine Learning (ICML 2019)*. 2019, pp. 4781–4789 (cit. on p. 13).
- [114] N. Nisan, T. Roughgarden, E. Tardos and V. Vazirani. *Algorithmic Game Theory*. New York, NY, USA: Cambridge University Press, 2007 (cit. on p. 4).
- [115] V. Perchet and M. Quincampoix. ‘A differential game on Wasserstein space. Application to weak approachability with partial monitoring’. In: *Journal of Dynamics and Games* 6 (2019), pp. 65–85 (cit. on p. 6).
- [116] V. Perchet and P. Rigollet. ‘The multi-armed bandit problem with covariates’. In: *Annals of Statistics* 41 (2013), pp. 693–721 (cit. on p. 6).
- [117] J. Peters, D. Janzing and B. Schölkopf. *Elements of causal inference*. The MIT Press, 2017 (cit. on p. 11).
- [118] G. Peyré and M. Cuturi. *Computational Optimal Transport*. ArXiv:1803.00567, 2018 (cit. on p. 6).
- [119] B. Rastegarpanah, M. Crovella and K. Gummadi. ‘Fair Inputs and Fair Outputs: The Incompatibility of Fairness in Privacy and Accuracy’. In: *FairUMAP*. 2020 (cit. on p. 10).
- [120] J. Reed and B. Pierce. ‘Distance makes the types grow stronger: a calculus for differential privacy’. In: *ACM Sigplan Notices*. Vol. 45. 2010, pp. 157–168 (cit. on p. 5).
- [121] A. Rohde, L. Steinberger et al. ‘Geometrizing rates of convergence under local differential privacy constraints’. In: *Annals of Statistics* 48.5 (2020), pp. 2646–2670 (cit. on p. 5).
- [122] A. Romanov, M. De-Arteaga, H. Wallach, J. Chayes, C. Borgs, A. Chouldechova, S. Geyik, K. Kenthapadi, A. Rumshisky and A. T. Kalai. ‘What’s in a Name? Reducing Bias in Bios without Access to Protected Attributes’. In: *arXiv preprint arXiv:1904.05233* (2019) (cit. on p. 10).
- [123] P. Sapiezynski, A. Gosh, L. Kaplan, A. Mislove and A. Rieke. ‘Algorithms that " Don’t See Color": Comparing Biases in Lookalike and Special Ad Audiences’. In: *arXiv preprint arXiv:1912.07579* (2019) (cit. on p. 10).

- [124] C. Sims. ‘Implications of rational inattention’. In: *Journal of monetary Economics* 50.3 (2003), pp. 665–690 (cit. on p. 5).
- [125] T. Speicher, M. Ali, G. Venkatadri, F. N. Ribeiro, G. Arvanitakis, F. Benevenuto, K. P. Gummadi, P. Loiseau and A. Mislove. ‘On the Potential for Discrimination in Online Targeted Advertising’. In: *ACM FAT**. 2018 (cit. on p. 5).
- [126] C. Villani. *Topics in optimal transportation*. Vol. 58. Graduate studies in Mathematics, AMS, 2003 (cit. on p. 6).
- [127] J. Weed, V. Perchet and P. Rigollet. ‘Online learning in repeated auctions’. In: *Proceedings of COLT 2016*. 2016 (cit. on pp. 12, 13).
- [128] B. Woodworth, S. Gunasekar, M. I. Ohannessian and N. Srebro. ‘Learning Non-Discriminatory Predictors’. In: *COLT*. 2017 (cit. on p. 5).
- [129] Y. Wu, L. Zhang and X. Wu. ‘Counterfactual Fairness: Unidentification, Bound and Algorithm.’ In: *IJCAI*. 2019, pp. 1438–1444 (cit. on p. 11).
- [130] Y. Wu, L. Zhang, X. Wu and H. Tong. ‘Pc-fairness: A unified framework for measuring causality-based fairness’. In: *Advances in Neural Information Processing Systems*. 2019, pp. 3404–3414 (cit. on p. 11).
- [131] M. B. Zafar, I. Valera, M. Gomez Rodriguez and K. P. Gummadi. ‘Fairness Beyond Disparate Treatment & Disparate Impact: Learning Classification Without Disparate Mistreatment’. In: *WWW*. 2017 (cit. on p. 5).
- [132] M. B. Zafar, I. Valera, M. Gomez Rodriguez, K. P. Gummadi and A. Weller. ‘From Parity to Preference-based Notions of Fairness in Classification’. In: *NIPS*. 2017 (cit. on p. 5).
- [133] R. Zemel, Y. Wu, K. Swersky, T. Pitassi and C. Dwork. ‘Learning Fair Representations’. In: *ICML*. 2013 (cit. on p. 5).