

RESEARCH CENTRE

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

IN PARTNERSHIP WITH:

Institut Polytechnique de Paris, Criteo

2022

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**

The Inria logo is a stylized, cursive script in red, positioned in the bottom right corner of the page.

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	3
2.1.2 Societal aspects and ethics	4
2.2 A large variety of tools and techniques	5
2.3 General objectives	6
3 Research program	7
3.1 Theme 1: Developing fair mechanisms for auctions and matching problems	7
3.1.1 Fairness in auction-based systems	7
3.1.2 Fairness in matching and selection problems	8
3.2 Theme 2: Reconciling, and enforcing privacy with fairness	9
3.3 Theme 3: Exploiting structure in online algorithms and learning problems	10
3.3.1 Leveraging structure in online matching	11
3.3.2 Exploiting side-information in sequential learning	12
4 Application domains	12
4.1 Typical problems and use-cases	12
4.2 Application areas	13
4.2.1 Online advertisement	13
4.2.2 “Physical matching”	14
5 Highlights of the year	14
6 New results	14
6.1 Matching, allocation, and auctions from a game theory perspective	14
6.2 Online learning	15
6.3 Privacy	16
6.4 Fairness	17
7 Partnerships and cooperations	18
7.1 International research visitors	18
7.1.1 Visits of international scientists	18
7.1.2 Visits to international teams	18
7.2 National initiatives	19
8 Dissemination	19
8.1 Promoting scientific activities	19
8.1.1 Scientific events: organisation	19
8.1.2 Scientific events: selection	20
8.1.3 Journal	20
8.1.4 Invited talks	21
8.1.5 Scientific expertise	21
8.2 Teaching - Supervision - Juries	21
8.2.1 Supervision	21
8.2.2 Juries	21
8.3 Teaching	22
9 Scientific production	22
9.1 Publications of the year	22
9.2 Cited publications	24

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, from Mar 2022, HDR]
- Marc Abeille [Criteo, Researcher, from Mar 2022]
- Louis Faury [Criteo, Researcher, from Mar 2022]
- Benjamin Heyman [Criteo, Researcher, from Mar 2022]
- Hugo Richard [Criteo, Researcher, from Mar 2022]
- Maxime Vono [Criteo, Researcher, from Mar 2022]

Faculty Members

- Vianney Perchet [Team leader, ENSAE & Criteo, Professor, from Mar 2022, HDR]
- Cristina Butucea [ENSAE, Professor, from Mar 2022, HDR]
- Julien Combe [Ecole Polytechnique, Associate Professor, from Mar 2022]
- Matthieu Lerasle [ENSAE, Professor, from Mar 2022, HDR]

Post-Doctoral Fellows

- Simon Finster [CNRS, from Nov 2022]
- Nadav Merlis [ENSAE, from Oct 2022]

PhD Students

- Ziyad Benomar [ENSAE, from Oct 2022]
- Nayel Bettache [ENSAE, from Mar 2022]
- Remi Castera [Univ. Grenoble Alpes, from Mar 2022]
- Hugo Chardon [ENSAE, from Mar 2022]
- Maria Cherifa [Criteo, CIFRE, from Mar 2022]
- Lorenzo Croissant [Criteo, CIFRE, from Mar 2022]
- Hafedh El Ferchichi [ENSAE, from Oct 2022]
- Côme Fiegel [ENSAE, from Oct 2022]
- Mike Liu [ENSAE, from Oct 2022]
- Mathieu Molina [Inria, from Mar 2022]
- Corentin Odic [ENSAE, from Mar 2022]
- Flore Sentenac [ENSAE, from Mar 2022]

Technical Staff

- Renaud Bauvin [Criteo, Engineer, from Mar 2022]
- Felipe Garrido Lucero [Inria, Engineer, from Oct 2022]

External Collaborators

- Clément Calauzènes [Criteo, from Mar 2022, scientific industrial advisor]
- Sophie Compagnon [Criteo, from Oct 2022, legal advisor]

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 [89]. 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” [98]: 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 [83, 109, 23, 67, 83, 25], but also in many other applications such as hiring [54], data-driven healthcare [61], or justice [84]. 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” [56], 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) [68], 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 [117], in-processing [115, 116, 112], and post-processing [68]. 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 [55, 57, 104] 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 [47]. 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 [108, 91, 86]. It was only recently that this approach was considered to measure “privacy losses” in economic mechanisms [58]. 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 [105, 42, 43] but is sometimes more adapted to handle user-generated data [63].

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 [62] 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 [46, 93].

Online and reinforcement learning [39] In online learning (a.k.a. multi-armed bandit [40, 100]), 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 [37, 85].

Online algorithms [35] 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 [52],

Optimal transport [110] 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 [99] thanks to its application to images and to new techniques to compute approximate matchings in a tractable way [102]. Even more unexpected applications of optimal transport have been discovered: to protect privacy [38], fairness [31], 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 of 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 [71] 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 [65] 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 [51, 50], but using different multipliers for the different groups (i.e., different values of the sensitive attribute).

Following recent theoretical work on auction fairness [44, 70, 48] (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 [59], 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) [103, 53].

A first option to solve this problem would be (when it is possible) to build proxies [66, 106] 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 [74], 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 [73].

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 [78]. 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 [107]—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 [81], which has many possible different level of assumptions and was popularized in 2020 [82]. The notion of counterfactual is key in causality [101]. 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 [77, 114, 113]. 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 [64, 38], 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.

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.

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 [92, 69, 29, 28]. 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 [92, 41]. 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 [69].

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 [92]. The seminal work [76] introduced an optimal algorithm for the adversarial case [32], 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} [72, 30, 60, 75, 87, 88]. 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 [88]. 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 [36]. A few results along these lines show that is a promising path. For instance, [41] 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, [49] designed a $1 - O(\sqrt{\log d}/\sqrt{d})$ competitive algorithm. [90] showed that greedy algorithms were highly efficient on Erdős-Renyi random graphs, with a competitive ratio of 0.837 in the worst case. [27] 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.

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 [24] 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 [33]. The second one involves the rich theory of graph weak local convergence [34]. 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 loose they learn that they would have lost with any lower bid [111, 45]. 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 [26]. 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 [80] 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 ask 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 [96]), 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 [111] 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 [97]. 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 [44, 95, 48, 70].

Matching [79, 94] 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 [78]. 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 [95], 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 CNAE, 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

The team was officially created on March 1st, 2022. This led a joint [press release](#) broadcasted on the side of both Criteo and Inria as the team is the first (and main, for now) element of the broader Inria-Criteo collaboration.

6 New results

6.1 Matching, allocation, and auctions from a game theory perspective

Participants: Vianney Perchet, Julien Combe, Clément Calauzènes.

A vast part of the Internet economy is powered by advertising, much of which is sold at auction. A key question for sellers is how to optimize the auction mechanism they use. Bidders, conversely, try to optimize their bidding strategy. Incentive compatible auctions are a sweet spot: theory predicts that it is in the bidders’ interest to bid their values, making it relatively easy for them to bid optimally. However, as they learn bidders’ value distributions, sellers can progressively optimize their mechanism and extract more revenue from bidders. In [5], we show that, in sharp contrast with most results in the academic

literature, bidders should not be bidding their value in incentive compatible auctions when there is no commitment from the seller about using a fixed auction. We provide a mix of theoretical and numerical results and practical methods that can easily be deployed in practice.

In [15], we describe an efficient algorithm to compute solutions for the general two-player Blotto game on n battlefields with heterogeneous values. While explicit constructions for such solutions have been limited to specific, largely symmetric or homogeneous, setups, this algorithmic resolution covers the most general situation to date: value-asymmetric game with asymmetric budget. The proposed algorithm rests on recent theoretical advances regarding Sinkhorn iterations for matrix and tensor scaling. An important case which had been out of reach of previous attempts is that of heterogeneous but symmetric battlefield values with asymmetric budget. In this case, the Blotto game is constant-sum so optimal solutions exist, and our algorithm samples from an ϵ -optimal solution in time $\tilde{O}(n^2 + \epsilon^{-4})$, independently of budgets and battlefield values. In the case of asymmetric values where optimal solutions need not exist but Nash equilibria do, our algorithm samples from an ϵ -Nash equilibrium with similar complexity but where implicit constants depend on various parameters of the game such as battlefield values.

In [2], we investigate the problem of reallocation with priorities where one has to assign objects or positions to individuals. Agents can have an initial ownership over an object. Each object has a priority ordering over the agents. In this framework, there is no mechanism that is both individually rational (IR) and stable, i.e. has no blocking pairs. Given this impossibility, an alternative approach is to compare mechanisms based on the blocking pairs they generate. A mechanism has minimal envy within a set of mechanisms if there is no other mechanism in the set that always leads to a set of blocking pairs included in the one of the former mechanism. Our main result shows that the modified Deferred Acceptance mechanism (Guillen and Kesten in *Int Econ Rev* 53(3):1027–1046, 2012), is a minimal envy mechanism in the set of IR and strategy-proof mechanisms. We also show that an extension of the Top Trading Cycle (Karakaya et al. in *J Econ Theory* 184:104948, 2019) mechanism is a minimal envy mechanism in the set of IR, strategy-proof and Pareto-efficient mechanisms. These two results extend the existing ones in school choice.

In the interdisciplinary paper [3], we discuss (in french) the perspectives of the new bioethics law from 2021, which opens up new ways of cross-donation of kidneys.

6.2 Online learning

Participants: Patrick Loiseau, Vianney Perchet.

In [4], we describe an approximate dynamic programming (ADP) approach to compute approximations of the optimal strategies and of the minimal losses that can be guaranteed in discounted repeated games with vector-valued losses. Among other applications, such vector-valued games prominently arise in the analysis of worst-case regret in repeated decision making in unknown environments, also known as the adversarial online learning framework. At the core of our approach is a characterization of the lower Pareto frontier of the set of expected losses that a player can guarantee in these games as the unique fixed point of a set-valued dynamic programming operator. When applied to the problem of worst-case regret minimization with discounted losses, our approach yields algorithms that achieve markedly improved performance bounds compared with off-the-shelf online learning algorithms like Hedge. These results thus suggest the significant potential of ADP-based approaches in adversarial online learning.

The workhorse of machine learning is stochastic gradient descent. To access stochastic gradients, it is common to consider iteratively input/output pairs of a training dataset. Interestingly, it appears that one does not need full supervision to access stochastic gradients, which is the main motivation of this paper. After formalizing the "active labeling" problem, which focuses on active learning with partial supervision, we provide in [7] a streaming technique that provably minimizes the ratio of generalization error over the number of samples. We illustrate our technique in depth for robust regression.

Potential buyers of a product or service, before making their decisions, tend to read reviews written by previous consumers. In [6], we consider Bayesian consumers with heterogeneous preferences, who sequentially decide whether to buy an item of unknown quality, based on previous buyers' reviews. The quality is multi-dimensional and may occasionally vary over time; the reviews are also multi-dimensional.

In the simple uni-dimensional and static setting, beliefs about the quality are known to converge to its true value. Our paper extends this result in several ways. First, a multi-dimensional quality is considered, second, rates of convergence are provided, third, a dynamical Markovian model with varying quality is studied. In this dynamical setting the cost of learning is shown to be small.

6.3 Privacy

Participants: Vianney Perchet, Patrick Loiseau.

Strategic information is valuable either by remaining private (for instance if it is sensitive) or, on the other hand, by being used publicly to increase some utility. These two objectives are antagonistic and leaking this information by taking full advantage of it might be more rewarding than concealing it. Unlike classical solutions that focus on the first point, in [1], we consider instead agents that optimize a natural trade-off between both objectives. We formalize this as an optimization problem where the objective mapping is regularized by the amount of information revealed to the adversary (measured as a divergence between the prior and posterior on the private knowledge). Quite surprisingly, when combined with the entropic regularization, the Sinkhorn loss naturally emerges in the optimization objective, making it efficiently solvable via better adapted optimization schemes. We empirically compare these different techniques on a toy example and apply them to preserve some privacy in online repeated auctions.

Contextual bandit algorithms are widely used in domains where it is desirable to provide a personalized service by leveraging contextual information, that may contain sensitive information that needs to be protected. Inspired by this scenario, we study in [10] the contextual linear bandit problem with differential privacy (DP) constraints. While the literature has focused on either centralized (joint DP) or local (local DP) privacy, we consider the shuffle model of privacy and we show that it is possible to achieve a privacy/utility trade-off between JDP and LDP. By leveraging shuffling from privacy and batching from bandits, we present an algorithm with regret bound $\tilde{O}(T^{2/3}/\epsilon^{1/3})$, while guaranteeing both central (joint) and local privacy. Our result shows that it is possible to obtain a trade-off between JDP and LDP by leveraging the shuffle model while preserving local privacy.

Contextual bandit is a general framework for online learning in sequential decision-making problems that has found application in a wide range of domains, including recommendation systems, online advertising, and clinical trials. A critical aspect of bandit methods is that they require to observe the contexts –i.e., individual or group-level data– and rewards in order to solve the sequential problem. The large deployment in industrial applications has increased interest in methods that preserve the users' privacy. In [11], we introduce a privacy-preserving bandit framework based on homomorphic encryption **which allows computations using encrypted data**. The algorithm *only* observes encrypted information (contexts and rewards) and has no ability to decrypt it. Leveraging the properties of homomorphic encryption, we show that despite the complexity of the setting, it is possible to solve linear contextual bandits over encrypted data with a $\tilde{O}(d\sqrt{T})$ regret bound in any linear contextual bandit problem, while keeping data encrypted.

In [16], we consider the problem of linear regression from strategic data sources with a public good component, i.e., when data is provided by strategic agents who seek to minimize an individual provision cost for increasing their data's precision while benefiting from the model's overall precision. In contrast to previous works, our model tackles the case where there is uncertainty on the attributes characterizing the agents' data – a critical aspect of the problem when the number of agents is large. We provide a characterization of the game's equilibrium, which reveals an interesting connection with optimal design. Subsequently, we focus on the asymptotic behavior of the covariance of the linear regression parameters estimated via generalized least squares as the number of data sources becomes large. We provide upper and lower bounds for this covariance matrix and we show that, when the agents' provision costs are superlinear, the model's covariance converges to zero but at a slower rate relative to virtually all learning problems with exogenous data. On the other hand, if the agents' provision costs are linear, this covariance fails to converge. This shows that even the basic property of consistency of generalized least squares estimators is compromised when the data sources are strategic.

6.4 Fairness

Participants: Patrick Loiseau.

To better understand discriminations and the effect of affirmative actions in selection problems (e.g., college admission or hiring), a recent line of research proposed a model based on differential variance. This model assumes that the decision-maker has a noisy estimate of each candidate's quality and puts forward the difference in the noise variances between different demographic groups as a key factor to explain discrimination. The literature on differential variance, however, does not consider the strategic behavior of candidates who can react to the selection procedure to improve their outcome, which is well-known to happen in many domains. In [9], we study how the strategic aspect affects fairness in selection problems. We propose to model selection problems with strategic candidates as a contest game: A population of rational candidates compete by choosing an effort level to increase their quality. They incur a cost-of-effort but get a (random) quality whose expectation equals the chosen effort. A Bayesian decision-maker observes a noisy estimate of the quality of each candidate (with differential variance) and selects the fraction α of best candidates based on their posterior expected quality; each selected candidate receives a reward S . We characterize the (unique) equilibrium of this game in the different parameters' regimes, both when the decision-maker is unconstrained and when they are constrained to respect the fairness notion of demographic parity. Our results reveal important impacts of the strategic behavior on the discrimination observed at equilibrium and allow us to understand the effect of imposing demographic parity in this context. In particular, we find that, in many cases, the results contrast with the non-strategic setting. We also find that, when the cost-of-effort depends on the demographic group (which is reasonable in many cases), then it entirely governs the observed discrimination (i.e., the noise becomes a second-order effect that does not have any impact on discrimination). Finally we find that imposing demographic parity can sometimes increase the quality of the selection at equilibrium; which surprisingly contrasts with the optimality of the Bayesian decision-maker in the non-strategic case. Our results give a new perspective on fairness in selection problems, relevant in many domains where strategic behavior is a reality.

In recent years, it has become clear that rankings delivered in many areas need not only be useful to the users but also respect fairness of exposure for the item producers. We consider the problem of finding ranking policies that achieve a Pareto-optimal tradeoff between these two aspects. Several methods were proposed to solve it; for instance a popular one is to use linear programming with a Birkhoffvon Neumann decomposition. These methods, however, are based on a classical Position Based exposure Model (PBM), which assumes independence between the items (hence the exposure only depends on the rank). In many applications, this assumption is unrealistic and the community increasingly moves towards considering other models that include dependences, such as the Dynamic Bayesian Network (DBN) exposure model. For such models, computing (exact) optimal fair ranking policies remains an open question. In [13], we answer this question by leveraging a new geometrical method based on the so-called expohedron proposed recently for the PBM (Kletti et al., WSDM'22). We lay out the structure of a new geometrical object (the DBN-expohedron), and propose for it a Carathéodory decomposition algorithm of complexity $O(n^3)$, where n is the number of documents to rank. Such an algorithm enables expressing any feasible expected exposure vector as a distribution over at most n rankings; furthermore we show that we can compute the whole set of Pareto-optimal expected exposure vectors with the same complexity $O(n^3)$. Our work constitutes the first exact algorithm able to efficiently find a Pareto-optimal distribution of rankings. It is applicable to a broad range of fairness notions, including classical notions of meritocratic and demographic fairness. We empirically evaluate our method on the TREC2020 and MSLR datasets and compare it to several baselines in terms of Paretooptimality and speed.

Statistical discrimination results when a decision-maker observes an imperfect estimate of the quality of each candidate dependent on which demographic group they belong to. Prior literature is limited to simple selection problems with a single decision-maker. In [8], we initiate the study of statistical discrimination in matching, where multiple decision-makers are simultaneously facing selection problems from the same pool of candidates (e.g., colleges admitting students). We propose a model where two colleges observe noisy estimates of each candidate's quality. The estimation noise controls a new key feature of the problem, namely the correlation between the estimates of the two colleges: if the noise is high, the

correlation is low and vice-versa. We consider stable matchings in an infinite population of students. We show that a lower correlation (i.e., higher estimation noise) for one of the groups worsens the outcome for all groups. Further, the probability that a candidate is assigned to their first choice is independent of their group. In contrast, the probability that a candidate is assigned to a college at all depends on their group, revealing the presence of discrimination coming from the correlation effect alone. Somewhat counter-intuitively the group that is subjected to more noise is better off.

Discrimination in machine learning often arises along multiple dimensions (a.k.a. protected attributes); it is then desirable to ensure intersectional fairness-i.e., that no subgroup is discriminated against. It is known that ensuring marginal fairness for every dimension independently is not sufficient in general. Due to the exponential number of subgroups, however, directly measuring intersectional fairness from data is impossible. In [14], our primary goal is to understand in detail the relationship between marginal and intersectional fairness through statistical analysis. We first identify a set of sufficient conditions under which an exact relationship can be obtained. Then, we prove bounds (easily computable through marginal fairness and other meaningful statistical quantities) in highprobability on intersectional fairness in the general case. Beyond their descriptive value, we show that these theoretical bounds can be leveraged to derive a heuristic improving the approximation and bounds of intersectional fairness by choosing, in a relevant manner, protected attributes for which we describe intersectional subgroups. Finally, we test the performance of our approximations and bounds on real and synthetic data-sets.

7 Partnerships and cooperations

7.1 International research visitors

7.1.1 Visits of international scientists

Other international visits to the team

A. Rohde

Status Professor

Institution of origin: Freiburg University

Country: Germany

Dates: Octobre 23-27, 2022

Context of the visit: Seminar talk and discussions

7.1.2 Visits to international teams

Research stays abroad

Cristina Butucea

Visited institution: Nottingham University

Country: UK

Dates: July 3-8, 2022

Context of the visit: Research, discussions

7.2 National initiatives

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.

Explainable and Responsible AI (MIAI chair)

Participants: Patrick Loiseau.

Title: Explainable and Responsible AI chair of the MIAI @ Grenoble Alpes institute

Partner Institution(s): • Univ. Grenoble Alpes

Date/Duration: 2019-2023 (4 years)

Additional info/keywords: Chair of the MIAI @ Grenoble Alpes institute co-held by Patrick Loiseau. Fairness, privacy.

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.

8 Dissemination

8.1 Promoting scientific activities

8.1.1 Scientific events: organisation

General chair, scientific chair

Participants: Vianney Perchet.

Title: Hi!Paris Symposium: Hi! PARIS Symposium on AI and Society

Partner Institution(s): • Crest, Genes

Date/Duration: June 2022

Member of the organizing committees

Participants: Vianney Perchet.

Title: ALT: Algorithmic Learning Theory

Partner Institution(s): • Crest, Genes

Date/Duration: April 2022

Participants: Cristina Butucea.

Title: Meetings in Mathematical Statistics, Luminy, Marseille

Partner Institution(s): • Crest, Genes; Ecole Centrale Marseille, CIRM

Date/Duration: December 2022

8.1.2 Scientific events: selection**Member of the conference program committees**

Patrick Loiseau: ICML, ECML-PKDD, The Web Conf, DE

Vianney Perchet: NeurIPS, ICLR, ICML, COLT, ALT, IJCAI

8.1.3 Journal**Member of the editorial boards**

Vianney Perchet: Operation Research, Operation Research Letters, Journal of Machine Learning Research, Journal of Dynamics and Games, SIAM data-science

Cristina Butucea: Annals of Statistics, Bernoulli

Reviewer - reviewing activities

Patrick Loiseau: Lato Sensu : Revue de la Société de philosophie des sciences, IEEE Open Journal of the Communications Society, Games and Economics Behavior, Statistics and Probability Letters

Julien Combe: Econometrica, Management Science, American Economic Review, Journal of Political Economy, Games and Economic Behavior, Operations Research Forum, Theoretical Economics, Revue d'Economie Politique

Vianney Perchet: Annals of Statistics, Mathematics of Operation Research, Journal of the ACM

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.

8.1.4 Invited talks

Vianney Perchet: IP Paris Optimization Meeting, Statistical inference and convex optimization, FILO-FOCS 2022, Dynamic Matching and Queueing Workshop, Summer Research Institute 2022 Learning: Optimization and Stochastics, Economics and Computation, Thirty-ninth International Conference on Machine Learning, Journées MAS 2022, Learning in Luminy Workshop, Statistic Seminar of Sorbonne University,

Julien Combe: Boston College Micro seminar (online), Lausanne Matching and Market Design Workshop 2022 at University of Lausanne, 12th Conference Economic Design at University of Padova, EEA-ESEM 2022 at University of Bocconi, 2022 Transatlantic Theory Workshop at Kellogg Business School, Stanford GSB theory seminar

Matthieu Lerasle Journées Statistiques du Sud, Journées ALEA, Colloquium Université Rouen, Séminaire Université du Luxembourg, Séminaire Geneva School of Economics and Management.

Cristina Butucea Université Paul Sabatier Toulouse, Collegio Carlo Alberto Torino, IMS London, ICSCA 2021 postponed to 2022 (visio) Conference Xian University China, Inverse Problems: From experimental data to models and back University of Potsdam, Workshop Mathematical Statistics in the Information Age Rostock.

Hugo Richard Séminaire Laboratoire Jean Kuntzmann (Grenoble).

8.1.5 Scientific expertise

Vianney Perchet: Expert for the Tenure committee of Tufts University

Patrick Loiseau: External expert for the evaluation of startups to enter in incubator Agoranov, Expert for the FR.S.-FNRS, Belgium

Cristina Butucea: Reviewer for tenure committee Hamburg University, hiring committee Professor Université de Toulouse

8.2 Teaching - Supervision - Juries

8.2.1 Supervision

Patrick Loiseau: PhD students: Rémi Castera, Mathieu Molina, Till Kletti, Vitalii Emelianov; postdocs: Felipe Garrido Lucero, Simon Finster

Vianney Perchet: PhD students: Sasila Ilandarideva, Flore Sentenac, Corentin Odic, Come Fiegel, Maria Cherifa, Mathieu Molina, Ziyad Benomar, Mike Liu, Hamed El Ferchichi. postdocs: Felipe Garrido Lucero, Hugo Richard, Nadav Merlis

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

Cristina Butucea PhD students: Clément Hardy, Nayel Bettache, Postdoc: Y. Issartel

8.2.2 Juries

Patrick Loiseau: PhD jury V. Grari (reviewer), mid-term Omar Boufous, Carlos Pinzon

Vianney Perchet: PhD jury: J. Achddou (reviewer), D. Beaudry, A. Bismuth, C.-S. Gauthier (reviewer), H. Dakdouk (reviewer), S. Gaucher, F. Hu, G. Rizk, P. Muller, HDR jury: A. Simonetto (reviewer)

Matthieu Lerasle PhD Jury: S. Arradi-Alaoui (reviewer), J. Cheng. HDR Jury: P. Mozharovskiy, A. Sabourin (internal reviewer).

Cristina Butucea PhD Jury: O. Collier; HDR Jury (reviewer) N. Verzelen

8.3 Teaching

ENSAE: Advanced Machine Learning (Vianney Perchet) Third year, lectures

Theoretical Foundations of Machine Learning (Vianney Perchet) Second year, lectures

Stopping time and online algorithms (Vianney Perchet) Third year, lectures

Statistics (Matthieu Lerasle) 1st and second year

Nonparametric Statistics (Cristina Butucea) 3rd year, M2

Mathematical Foundations of Probabilities (Cristina Butucea) 1st year

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

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

ECO301: Advanced microeconomics (Julien Combe). Third-year bachelor of science in Mathematics and Economics, Lecture and tutorials.

MIE65: Advanced Microeconomics : design and study of markets (Julien Combe). M2 PhD track level, lecture.

ECO361: Introduction to economics (Julien Combe). First-year cycle polytechnicien, PCs

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

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

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

ENPC: High Dimensional statistics (Hugo Richard). Third-year/M1 level, Lectures and labs.

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

9 Scientific production

9.1 Publications of the year

International journals

- [1] E. Boursier and V. Perchet. ‘Utility/privacy trade-off as regularized optimal transport’. In: *Mathematical Programming* (2022). DOI: [10.1007/s10107-022-01811-w](https://doi.org/10.1007/s10107-022-01811-w). URL: <https://hal.science/hal-03951419>.
- [2] J. Combe. ‘Reallocation with priorities and minimal envy mechanisms’. In: *Economic Theory* (7th Feb. 2022). DOI: [10.1007/s00199-022-01465-x](https://doi.org/10.1007/s00199-022-01465-x). URL: <https://hal.inria.fr/hal-03951678>.
- [3] J. Combe, V. Hiller, O. Tercieux, B. Audry, J. Baudet, G. Malaquin, F. Kerbaul, C. Antoine, M.-A. Macher, C. Jacquelinet, O. Bastien and M. Pastural. ‘Perspectives for future development of the kidney paired donation programme in France’. In: *Néphrologie & Thérapeutique* 18.4 (July 2022), pp. 270–277. DOI: [10.1016/j.nephro.2022.02.001](https://doi.org/10.1016/j.nephro.2022.02.001). URL: <https://shs.hal.science/halshs-03761536>.
- [4] V. Kamble, P. Loiseau and J. Walrand. ‘An Approximate Dynamic Programming Approach to Repeated Games with Vector Losses’. In: *Operations Research* (2022). DOI: [10.1287/opre.2022.2334](https://doi.org/10.1287/opre.2022.2334). URL: <https://hal.inria.fr/hal-03951270>.
- [5] T. Nedelec, C. Calauzènes, V. Perchet and N. El Karoui. ‘Revenue-Maximizing Auctions: A Bidder’s Standpoint’. In: *Operations Research* 70.5 (Sept. 2022), pp. 2767–2783. DOI: [10.1287/opre.2022.2316](https://doi.org/10.1287/opre.2022.2316). URL: <https://hal.science/hal-03951434>.

International peer-reviewed conferences

- [6] E. Boursier, V. Perchet and M. Scarsini. ‘Social Learning in Non-Stationary Environments’. In: alt2022 - 33rd International Conference on Algorithmic Learning Theory. Paris, France, 2022. URL: <https://hal.science/hal-03089798>.
- [7] V. Cabannes, F. S. Bach, V. Perchet and A. Rudi. ‘Active Labeling: Streaming Stochastic Gradients’. In: NeurIPS 2022 - 36th Conference on Neural Information Processing Systems. New Orleans, United States, 28th Nov. 2022. URL: <https://hal.inria.fr/hal-03806666>.
- [8] R. Castera, P. Loiseau and B. Pradelski. ‘Statistical Discrimination in Stable Matchings’. In: ACM Conference on Economics and Computation (EC’22). Boulder, Colorado, United States, 11th July 2022. DOI: [10.1145/3490486.3538364](https://doi.org/10.1145/3490486.3538364). URL: <https://hal.science/hal-03672270>.
- [9] V. Emelianov, N. Gast and P. Loiseau. ‘Fairness in Selection Problems with Strategic Candidates’. In: EC 2022 - ACM Conference on Economics and Computation. Boulder, Colorado, United States: ACM, 11th July 2022, pp. 1–29. DOI: [10.1145/3490486.3538287](https://doi.org/10.1145/3490486.3538287). URL: <https://hal.inria.fr/hal-03677966>.
- [10] E. Garcelon, K. Chaudhuri, V. Perchet and M. Pirodda. ‘Privacy Amplification via Shuffling for Linear Contextual Bandits’. In: The 33rd International Conference on Algorithmic Learning Theory. Paris, France, 11th Dec. 2021. URL: <https://hal.science/hal-03951300>.
- [11] E. Garcelon, V. Perchet and M. Pirodda. ‘Encrypted Linear Contextual Bandit’. In: 25th International Conference on Artificial Intelligence and Statistics. Valence, Spain, 2022. URL: <https://hal.science/hal-03951440>.
- [12] E. Gkiouzepi, A. Andreou, O. Goga and P. Loiseau. ‘Collaborative Ad Transparency: Promises and Limitations’. In: 44th IEEE Symposium on Security and Privacy. San Francisco, United States, 2023. URL: <https://hal.inria.fr/hal-03916393>.
- [13] T. Kletti, J.-M. Renders and P. Loiseau. ‘Pareto-Optimal Fairness-Utility Amortizations in Rankings with a DBN Exposure Model’. In: SIGIR 2022 - 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. Madrid, Spain: ACM, 11th July 2022, pp. 1–12. DOI: [10.1145/3477495.3532036](https://doi.org/10.1145/3477495.3532036). URL: <https://hal.inria.fr/hal-03691743>.
- [14] M. Molina and P. Loiseau. ‘Bounding and Approximating Intersectional Fairness through Marginal Fairness’. In: NeurIPS 2022 - 36th Conference on Neural Information Processing Systems. New Orleans, United States, 28th Nov. 2022, pp. 1–32. URL: <https://hal.inria.fr/hal-03827777>.
- [15] V. Perchet, P. Rigollet and T. L. Gouic. ‘An algorithmic solution to the Blotto game using multi-marginal couplings’. In: The Twenty-Third ACM Conference on Economics and Computation. Boulder (CO), France, 11th July 2022. URL: <https://hal.inria.fr/hal-03941399>.
- [16] B. Roussillon, N. Gast, P. Loiseau and P. Mertikopoulos. ‘Asymptotic Degradation of Linear Regression Estimates with Strategic Data Sources’. In: ALT 2022 - 33rd International Conference on Algorithmic Learning Theory. Paris, France, 29th Mar. 2022, pp. 1–31. URL: <https://hal.univ-grenoble-alpes.fr/hal-03593516>.

Reports & preprints

- [17] E. Boursier and V. Perchet. *A survey on multi-player bandits*. 29th Nov. 2022. URL: <https://hal.inria.fr/hal-03941302>.
- [18] C. Butucea, J.-F. Delmas, A. Dutfoy and C. Hardy. *OFF-THE-GRID LEARNING OF SPARSE MIXTURES FROM A CONTINUOUS DICTIONARY*. 28th June 2022. URL: <https://hal.science/hal-03707465>.
- [19] C. Butucea, J.-F. Delmas, A. Dutfoy and C. Hardy. *OFF-THE-GRID PREDICTION AND TESTING FOR MIXTURES OF TRANSLATED FEATURES*. 2nd Dec. 2022. URL: <https://hal.science/hal-03880134>.
- [20] C. Butucea, J.-F. Delmas, A. Dutfoy and C. Hardy. *SIMULTANEOUS OFF-THE-GRID LEARNING OF MIXTURES ISSUED FROM A CONTINUOUS DICTIONARY*. 26th Oct. 2022. URL: <https://hal.science/hal-03831208>.

- [21] C. Fiegel, P. Ménard, T. Kozuno, R. Munos, V. Perchet and M. Valko. *Adapting to game trees in zero-sum imperfect information games*. 23rd Dec. 2022. URL: <https://hal.inria.fr/hal-03941364>.
- [22] H. Richard, F. Sentenac, C. Odic, M. Molina and V. Perchet. *Static Scheduling with Predictions Learned through Efficient Exploration*. 31st May 2022. URL: <https://hal.inria.fr/hal-03941374>.

9.2 Cited publications

- [23] ‘24 CFR § 100.75 - Discriminatory advertisements, statements and notices.’ <https://www.law.cornell.edu/cfr/text/24/100.75>.
- [24] E. Abbe. ‘Community detection and stochastic block models: recent developments’. In: *The Journal of Machine Learning Research* 18.1 (2017), pp. 6446–6531.
- [25] 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.
- [26] 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.
- [27] 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.
- [28] 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].
- [29] 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.
- [30] B. Bahmani and M. Kapralov. ‘Improved bounds for online stochastic matching’. In: *European Symposium on Algorithms*. Springer. 2010, pp. 170–181.
- [31] 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.
- [32] B. Birnbaum and C. Mathieu. ‘On-line bipartite matching made simple’. In: *Acm Sigact News* 39.1 (2008), pp. 80–87.
- [33] B. Bollobas and G. Brightwell. ‘The width of random graph orders’. In: *The Mathematical Scientist* 20 (Jan. 1995).
- [34] C. Bordenave, M. Lelarge and J. Salez. *Matchings on infinite graphs*. 2011. arXiv: [1102.0712](https://arxiv.org/abs/1102.0712) [math.PR].
- [35] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
- [36] 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].
- [37] E. Boursier and V. Perchet. ‘SIC-MMAB: Synchronisation Involves Communication in Multiplayer Multi-Armed Bandits’. In: *arXiv:1809.08151*. 2018, pp. 1–31.
- [38] 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.
- [39] 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.
- [40] 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.

- [41] N. Buchbinder, K. Jain and J. (Naor. ‘Online Primal-Dual Algorithms for Maximizing Ad-Auctions Revenue’. In: *Algorithms – ESA 2007*. Ed. by L. Arge, M. Hoffmann and E. Welzl. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 253–264.
- [42] 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.
- [43] 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.
- [44] L. E. Celis, A. Mehrotra and N. K. Vishnoi. ‘Toward Controlling Discrimination in Online Ad Auctions’. In: *ICML*. 2019.
- [45] N. Cesa-Bianchi, C. Gentile and Y. Mansour. ‘Regret minimization for reserve prices in second-price auctions’. In: *Proceedings of SODA 2013*. 2013.
- [46] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [47] K. Chaudhuri, J. Imola and A. Machanavajjhala. ‘Capacity bounded differential privacy’. In: *Advances in Neural Information Processing Systems*. 2019, pp. 3469–3478.
- [48] S. Chawla and M. Jagadeesan. *Fairness in ad auctions through inverse proportionality*. arXiv:2003.13966. 2020.
- [49] 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.
- [50] 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.
- [51] V. Conitzer, C. Kroer, E. Sodomka and N. E. Stier-Moses. ‘Multiplicative pacing equilibria in auction markets’. In: *WINE*. 2018.
- [52] 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.
- [53] R. Cummings, V. Gupta, D. Kimpara and J. Morgenstern. ‘On the Compatibility of Privacy and Fairness’. In: *FairUMAP*. 2019.
- [54] J. Destin. *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.
- [55] C. Dwork. ‘Differential privacy’. In: *Encyclopedia of Cryptography and Security* (2011), pp. 338–340.
- [56] C. Dwork, M. Hardt, T. Pitassi, O. Reingold and R. Zemel. ‘Fairness through awareness’. In: *ITCS*. 2012.
- [57] 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.
- [58] R. Eilat, K. Eliaz and X. Mu. *Optimal Privacy-Constrained Mechanisms*. Tech. rep. C.E.P.R. Discussion Papers, 2019.
- [59] 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.
- [60] J. Feldman, A. Mehta, V. Mirrokni and S. Muthukrishnan. *Online Stochastic Matching: Beating 1-1/e*. 2009. arXiv: 0905.4100 [cs.DS].
- [61] K. Ferryman and M. Pitcan. *Fairness in Precision Medicine*. Tech. rep. Data & Society, 2018.
- [62] D. Fudenberg and J. Tirole. *Game Theory*. MIT press, 1991.
- [63] E. Garcelon, V. Perchet, C. Pike-Burke and M. Pirota. ‘Local Differentially Private Regret Minimization in Reinforcement Learning’. In: *arXiv preprint arXiv:2010.07778* (2020).

- [64] 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.
- [65] N. Gast and B. Van Houdt. ‘A Refined Mean Field Approximation’. In: *SIGMETRICS*. 2017.
- [66] M. Gupta, A. Cotter, M. M. Fard and S. Wang. ‘Proxy fairness’. In: *arXiv preprint arXiv:1806.11212* (2018).
- [67] 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.
- [68] M. Hardt, E. Price and N. Srebro. ‘Equality of Opportunity in Supervised Learning’. In: *NIPS*. 2016.
- [69] C.-J. Ho and J. W. Vaughan. ‘Online task assignment in crowdsourcing markets’. In: *Twenty-sixth AAAI conference on artificial intelligence*. 2012.
- [70] C. Ilvento, M. Jagadeesan and S. Chawla. ‘Multi-Category Fairness in Sponsored Search Auctions’. In: *FAT**. 2020.
- [71] K. Iyer, R. Johari and M. Sundararajan. ‘Mean field equilibria of dynamic auctions with learning’. In: *Management Science* 60.12 (2014), pp. 2949–2970.
- [72] P. Jaillet and X. Lu. ‘Online stochastic matching: New algorithms with better bounds’. In: *Mathematics of Operations Research* 39.3 (2014), pp. 624–646.
- [73] 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).
- [74] N. Kallus, X. Mao and A. Zhou. ‘Assessing algorithmic fairness with unobserved protected class using data combination’. In: *arXiv preprint arXiv:1906.00285* (2019).
- [75] 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.
- [76] 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.
- [77] 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.
- [78] J. S. Kim, J. Chen and A. Talwalkar. ‘Model-Agnostic Characterization of Fairness Trade-offs’. In: *arXiv preprint arXiv:2004.03424* (2020).
- [79] 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.
- [80] V. Krishna. *Auction Theory*. 1st ed. Elsevier, 2009.
- [81] M. J. Kusner, J. Loftus, C. Russell and R. Silva. ‘Counterfactual fairness’. In: *Advances in neural information processing systems*. 2017, pp. 4066–4076.
- [82] M. J. Kusner and J. R. Loftus. *The long road to fairer algorithms*. 2020.
- [83] 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).
- [84] 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.
- [85] L. Liu, H. Mania and M. I. Jordan. ‘Competing Bandits in Matching Markets’. In: *arXiv:1906.05363*. 2019, pp. 1–15.
- [86] B. Maćkowiak and M. Wiederholt. ‘Business cycle dynamics under rational inattention’. In: *The Review of Economic Studies* 82.4 (2015), pp. 1502–1532.

- [87] 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.
- [88] 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.
- [89] A. Mas-Colell, M. Whinston and J. Green. *Microeconomic Theory*. New York: Oxford University Press, 1995.
- [90] A. Mastin and P. Jaillet. ‘Greedy online bipartite matching on random graphs’. In: *arXiv preprint arXiv:1307.2536* (2013).
- [91] 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.
- [92] 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>.
- [93] J.-F. Mertens, S. Sorin and S. Zamir. *Repeated Games*. Econometric Society Monographs. Cambridge University Press, 2015.
- [94] 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.
- [95] M. Nasr and M. C. Tschantz. ‘Bidding Strategies with Gender Nondiscrimination Constraints for Online Ad Auctions’. In: *FAT**. 2020.
- [96] 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.
- [97] 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.
- [98] N. Nisan, T. Roughgarden, E. Tardos and V. Vazirani. *Algorithmic Game Theory*. New York, NY, USA: Cambridge University Press, 2007.
- [99] 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.
- [100] V. Perchet and P. Rigollet. ‘The multi-armed bandit problem with covariates’. In: *Annals of Statistics* 41 (2013), pp. 693–721.
- [101] J. Peters, D. Janzing and B. Schölkopf. *Elements of causal inference*. The MIT Press, 2017.
- [102] G. Peyré and M. Cuturi. *Computational Optimal Transport*. ArXiv:1803.00567, 2018.
- [103] B. Rastegarpanah, M. Crovella and K. Gummadi. ‘Fair Inputs and Fair Outputs: The Incompatibility of Fairness in Privacy and Accuracy’. In: *FairUMAP*. 2020.
- [104] 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.
- [105] 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.
- [106] 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).
- [107] 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).
- [108] C. Sims. ‘Implications of rational inattention’. In: *Journal of monetary Economics* 50.3 (2003), pp. 665–690.

-
- [109] 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.
 - [110] C. Villani. *Topics in optimal transportation*. Vol. 58. Graduate studies in Mathematics, AMS, 2003.
 - [111] J. Weed, V. Perchet and P. Rigollet. ‘Online learning in repeated auctions’. In: *Proceedings of COLT 2016*. 2016.
 - [112] B. Woodworth, S. Gunasekar, M. I. Ohanessian and N. Srebro. ‘Learning Non-Discriminatory Predictors’. In: *COLT*. 2017.
 - [113] Y. Wu, L. Zhang and X. Wu. ‘Counterfactual Fairness: Unidentification, Bound and Algorithm.’ In: *IJCAI*. 2019, pp. 1438–1444.
 - [114] 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.
 - [115] 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.
 - [116] 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.
 - [117] R. Zemel, Y. Wu, K. Swersky, T. Pitassi and C. Dwork. ‘Learning Fair Representations’. In: *ICML*. 2013.