



HABILITATION À DIRIGER DES RECHERCHES
UNIVERSITÉ PIERRE ET MARIE CURIE — PARIS 6

**Combining game theory and statistical learning
for security, privacy and network systems**

Patrick Loiseau

EURECOM

Soutenu le 8 Décembre 2016 devant un jury composé de :

Eitan Altman, DR Inria	Rapporteur
Tamer Başar, Professeur UIUC	Rapporteur
Gérard Biau, Professeur UPMC	Examineur
Rainer Böhme, Professeur University of Innsbruck	Rapporteur
Bruno Gaujal, DR Inria	Examineur
Refik Molva, Professeur EURECOM	Examineur
Vianney Perchet, Professeur ENS Paris Saclay	Examineur

C'est ce que nous pensons déjà connaître qui nous empêche souvent d'apprendre.
Claude Bernard.

ACKNOWLEDGEMENTS

10-25-2016

I remember reading a book whose Preface read something like: “For my previous book, I started by writing a very nice Preface and ended up never finishing the actual book. [...]” In an effort to avoid this, I have decided not to write acknowledgments before the manuscript itself is finished; even though I do expect that this will by far be the most interesting part.

12-08-2016

Finally, the time has come to write the acknowledgements. I might not be the strongest supporter of HDR as a useful diploma to evaluate the ability of a researcher, but on a personal level, I must admit that it had two very valuable benefits.

First, it forced a number (7) of famous researchers who are otherwise very busy to go through my work and get to know it better. I want to thank Eitan Altman, Tamer Başar and Rainer Böhme for accepting to review my manuscript, and Gérard Biau, Bruno Gaujal, Refik Molva and Vianney Perchet for accepting to be in the jury of my HDR. I was deeply honored that such distinguished researchers accepted to spend time evaluating my work, and I hope that they found some interesting things in it so that this was not a complete waste of their time.

Second, it gives me a rare occasion to thank the many people who have contributed to my work and to my life as a researcher. I would like to start by thanking EURECOM, its director Ulrich Finger, and my two successive department heads Refik Mokva and Pietro Michiardi for giving me the conditions and freedom to develop my work and for supporting my various projects and constraints.

Although the very exercise of writing an HDR manuscript seems to consist in finding a unifying logic in my work, it would be silly to deny that in reality, much of my research was influenced by people that I had the chance to meet. I would like to first thank John Musacchio who trusted me to join him as a postdoc to work on network economics even though I had no clue at the time what game theory even was. I have never regretted deciding to join him and to move to research on game theory. I also thank Galina Schwartz for patiently introducing me to basic economics and guiding me through my first steps. I hope our collaboration has later become more mutually fruitful. I thank Jean Walrand for warmly welcoming me at Berkeley and always supporting me and advising me kindly. Finally, I want to thank Krishna Gummadi for hosting me at MPI-SWS, for sharing so many ideas, and for always taking the time to cope with my problems, as silly as they can be. The times I spend at MPI-SWS always bring me a fantastic combination of excitement and peace of mind that make me love research so much. All these persons have offered me strong support at times where I often wasn't too sure that I deserved it, and this has deeply affected my career. I hope that I was worthy of their efforts and that I will be as benevolent to future generations as they were to me. Of course, I also remember my earlier advisors Paulo Gonçalves and Julien Barral, whose influence I can sometimes still see in my research, and I thank Ravi Mazumdar for his constant support at various important moments.

Research is always the fruit of a collaborative effort and I have been lucky on that to work with great people. I deeply thank all my collaborators, in particular Saurabh Amin, Ernst Biersack, Jens Grossklags, Stratis Ioannidis and Giovanni Neglia. Of course, much of the hard work was often done by interns, students, and postdocs and I would like to thank them all sincerely for their commitment and efforts: Athanasios Andreou, Alberto Benegiamo, Michela Chessa, Lemonia Dritsoula, Nina Grgić-Hlača, Hadrien Hours, Vijay Kamble, Amine Lahouel, Yifan Pi, Yannick Terme, Quan Dong Vu, Xiaohu Wu.

I thank my colleagues and some members of my community at large, with whom I had interactions much beyond the context of research that made my work more pleasant on a daily basis. I also thank my friends and family for making my life happier.

Finally, I want to say a few words for my wife, although I know it is impossible to acknowledge all she brings to me within a few lines. Oana has always been there for me. She is the one who had to support all my disappointments and the one who helped me get back on my feet so many times. She showed me how to be strong, trust in my ideas and fight for them even in hard times. But she is also a big part of why I like research. Her energy, excitement, and constant flow of fascinating ideas are inspiring and motivating for me. More importantly, Oana is simply the most incredible person I know and being with her has made me an infinitely happier and better man. I also thank our son Luca, for having the good idea of waiting that I had sent my manuscript to the reviewers before showing up, but also mostly for bringing us so much joy. I don't know if he will enjoy research but if he does, I hope he is as good as his mother at it.

CONTENTS

Acknowledgements	iii
Abstract	vii
1 Introduction	1
1.1 Context and motivation	1
1.2 Organization of the manuscript	2
2 Game theory and statistical learning for security	5
2.1 Adversarial classification	5
2.1.1 Context	5
2.1.2 Our contributions	7
2.2 Security resource allocation	12
2.2.1 Context	13
2.2.2 Our contributions	13
2.3 Adversarial online decision making with discounted losses	16
2.3.1 Context	16
2.3.2 Our contributions	17
3 Game theory and statistical learning for privacy	23
3.1 Learning from personal data provided by privacy-conscious users	23
3.1.1 Context	23
3.1.2 Our contributions	25
3.2 Learning and estimation of the privacy risks from public data	28
3.2.1 Context	28
3.2.2 Our contributions	29
4 Game theory and statistical learning for network systems	33
4.1 Practical causal analysis for network performance prediction	33
4.1.1 Context	33
4.1.2 Our contributions	34
4.2 Incentives in networks with congestion	38
4.2.1 Context	38
4.2.2 Our contributions	40
4.3 Cloud resources allocation	44
4.3.1 Context	45

4.3.2 Our contributions	45
5 Conclusion and perspectives	47
5.1 Learning from strategic data	47
5.2 Human-friendly learning algorithms	49
A Publications	51
B Summary of my professional activities and advising experience and complete CV	55
References	71

ABSTRACT

The Internet is now omnipresent in our lives. We use user-centric services, that is digital services leveraging our personal data to provide personalized services of high value, for almost everything from social networking to shopping, banking, or entertainment. With great utility, however, digital user-centric services also brought very important security and privacy problems that threaten our well-being, and the growth and sustainability of digital services. Our increased dependence on online services also reinforced the need for improving the network infrastructure that support them.

I argue that tackling these essential questions requires a combination of methods from game theory and statistical learning. Game theory because the security, privacy, and performance of user-centric services ultimately depend on the behavior of humans who respond to the incentives provided by the system's design, and game theory is the natural tool to model such strategic interactions. Statistical learning because it is at the core of user-centric services, both to secure the system and to exploit personal data.

This manuscript synthesizes my research efforts on game theory and statistical learning for security, privacy and network systems. I first focus on the security aspects and describe my work on developing and using game-theoretic models to design classification, resource allocation and sequential learning methods in adversarial environment. Then I focus on the privacy aspects and describe my work on developing and studying algorithms to learn from personal data and analyzing their impact on privacy. Finally, I focus on the network systems aspects and describe my work on analyzing and improving the infrastructure's performance. I conclude by describing the perspectives of my research, summarized as the study of 'humans *versus* machine learning' and containing two main directions: (i) developing algorithms to learn from data generated or provided by strategic human agents for security and privacy (using game theory), and (ii) studying how machine learning algorithms inconspicuously affect humans in their daily lives and how to make them more 'human-friendly'.

INTRODUCTION

1.1 Context and motivation

The digital world has become central in our lives and economies. Over the last two decades, the Internet enabled the emergence of an ecosystem of user-centric services that collect large amounts of personal data of users and use it to provide high-value services through personalization. Today, we use such digital user-centric services for almost everything from social networking to shopping, banking, or entertainment.

Despite their success, digital user-centric services also brought very important problems of security and privacy. On the security side, user-centric services are plagued with cyber-criminality issues that have consequences on users at various levels, from the incessant presence of illegitimate/undesired content that reduces the service utility to more severe problems such as identity theft and data breaches that can seriously damage users integrity. On the privacy side, concerns have been growing exponentially in recent times. Users are increasingly worried about the loss of privacy incurred when revealing their personal data. They do not feel appropriately rewarded for the data they give and this is amplified by the lack of transparency from services on how the data is used and shared.

As we become more dependent of digital services, we naturally become more sensitive to their security and privacy weaknesses; which reinforces the urge to improve them. We also become increasingly dependent of the network systems and infrastructures (communication network and cloud) that support digital services; which also reinforces the need to improve their performance.

Designing a secure, private or well-performing user-centric service cannot be done based on technical engineering considerations only. Indeed, all three aspects are strongly dependent on how human users use the system and will therefore be governed to a large extent by incentives and strategic considerations. For security, attacks are generated by strategic human agents that can adapt and one needs to take that into account in the defense's design. For privacy, data is revealed by human users who attempt to protect their privacy and one needs to take that into account in the design of how data is collected and used. For the infrastructure, its performance is determined by the load which is generated by human users so that reducing the load often necessitates to act on the users demand.

In all three cases, we see that taking into account the incentives of the different parties at stake is key in the design of the system, so that the system eventually works well. The research field that does that broadly is called network economics. Taking incentives into account is typically done using *game theory* as the natural tool to model the strategic agents and find the optimal system design given agents' incentives. The field of network economics has addressed with success a few categories of systems in the past (e.g., peer-to-peer networks), leading to the design of systems that indeed performed much better in practice. However, many questions remain open, in particular in

the area of security and privacy of user-centric services.

Another key aspect of user-centric services is that they heavily rely on *statistical learning* for at least two purposes: to exploit users personal data to predict their preference and increase the service value, and to detect attacks on their services. Machine learning is a well-established field with many recent advances. However, the context of security and privacy of user-centric services is special and yields many new challenges in statistical learning, in particular because one needs to take into account the presence of strategic agents (attackers or privacy-conscious users) who can alter the data, but also because it raises issues related to the very large scale of the problem due to the number of users of user-centric services.

My research efforts in the past years have been dedicated to working on game theory and statistical learning for tackling problems of security, privacy and network systems performance. Although my work was often motivated by applications that guide the models that I consider, I often worked on the methods themselves (solving new games, developing new statistical learning methods or studying existing ones in a new game-theoretic context).

1.2 Organization of the manuscript

This manuscript synthesizes my research contributions on game theory and statistical learning for security, privacy and network systems. It is organized in three chapters.

Chapter 2 summarizes my work on game theory and statistical learning for security. At a high level in this work, I develop and use game-theoretic models to design learning and resource allocation methods for security that work well in adversarial environment. I first study classification in adversarial environment using a new nonzero-sum game model to find attack detection methods adapted to the strategic setting. Then I investigate the famous Colonel Blotto game for allocation of security resources in adversarial environment. Finally, I tackle the question of sequential learning in adversarial environment under discounted losses.

Chapter 3 summarizes my work on game theory and statistical learning for privacy. At a high level in this work, I study how algorithms to learn from personal data affect privacy. I first investigate the problem of learning from personal data revealed by privacy-conscious users who may choose the precision of the data revealed, using a new game that models the public good nature of the learning outcome. Then I investigate the possibilities of learning from personal data already publicly available (in particular matching identities across multiple online sources at very large scale) and how this affects users privacy.

Chapter 4 summarizes my work on game theory and statistical learning for network systems. Here, I focus on the infrastructure's performance evaluation and improvement. I first study how causal analysis can be applied for communication networks performance analysis. Then I investigate new incentive schemes to reduce network congestion and flatten the demand curve. Finally, I look at how to allocate resources in a cloud.

Finally, in Chapter 5, I conclude by presenting the perspectives that my past work has opened for my future research.

Biographical note

The work presented in this manuscript constitutes the main part of my research since the end of my PhD. Some of my works have been omitted in the interest of thematic consistency. My complete set of contributions is represented by my papers listed in Appendix A.

Prior to starting working in network economics, I did my PhD in the area of network traffic modeling. There I looked at techniques from probability and stochastic processes, in particular long-range dependence and large deviations. I then spent one year as a postdoc working on applications of large deviations in the medical area, before moving to game theory and network economics during my second year of postdoc. Although the works presented in this manuscript are in a different area than my PhD and first year of postdoc, my background in probability has been of great use and, as I am starting to dig deeper into the side of statistical learning and its interaction with game theory, I feel that this background will become even more important.

GAME THEORY AND STATISTICAL LEARNING FOR SECURITY

This chapter covers the work in papers [8, 9], [30, 32, 33]. This work was done in collaboration with John Musacchio (UC Santa Cruz), Shankar Sastry (UC Berkeley), Galina Schwartz (UC Berkeley), Jean Walrand (UC Berkeley); and with the following students: LEMONIA DRITSOULA (PhD student at UC Santa Cruz whom I informally advised for the work presented here), Nina Grgić-Hlača (intern at EURECOM under my supervision), Vijay Kamble (PhD student at UC Berkeley who came to EURECOM as an intern under my supervision during his PhD). The ideas developed here also led to the PhD of Quan Vu which will start in Dec. 2016 under my supervision and in collaboration with Alonso Silva (Nokia Bell Labs).

This chapter summarizes my work relating to game theory and statistical learning for security. Throughout this work, my main focus was to work on developing defense methods that take into account the strategic nature of attackers, which naturally leads to game-theoretic analysis. I do believe that using in security standard methods (e.g., learning methods) that neglect the strategic nature of the attacker is one of the key reasons for underperformance; and that it is therefore crucial to rethink those methods carefully. I also focused more on developing and analyzing generic models that lead to simple and generally applicable insights rather than looking at very specific models for given security situations. This is mainly justified by the fact that many methods are common to various security applications; e.g., classification for attack detection is used similarly for simple spam detection or for life-threatening intrusions.

In the remainder of the chapter, I sequentially describe my contributions to classification in adversarial environment for attack detection, security resource allocation in adversarial environment, and online decision making with adversarial losses.

2.1 Adversarial classification

In many areas relating to digital systems (and in many more areas) security relies on *detecting attacks*, which is outsourced to a machine learning algorithm (classification or anomaly detection), see e.g., [Taylor et al., 2007, Tsai and Yu, 2009, Guzella and Caminhas, 2009, Song et al., 2010, Caruana and Li, 2012]. In this section, we investigate that problem of detecting attacks.

2.1.1 Context

Recent works from the security community mostly focused on finding new features that are good for the learning algorithm performance, and at implementing the systems using standard learning

algorithms developed for non-adversarial scenarios (see, e.g., [Stringhini et al., 2012] or the references above). It was shown, however, on the example of SpamBayes, that a spammer controlling only 1% of the training set can make the filter useless in two ways: by making it unable to detect spam or by making it mark as spam other messages [Nelson et al., 2009]. It is therefore of crucial importance to rethink the learning algorithms that are used for security applications, which are by essence adversarial scenarios.

The study of classification algorithms in adversarial environment was pioneered by [Dalvi et al., 2004] and a significant literature followed [Lowd and Meek, 2005, Globerson and Roweis, 2006, Barreno et al., 2010, Laskov and Lippmann, 2010, Nelson et al., 2010, Huang et al., 2011, Zhou and Kantarcioglu, 2014, Li and Vorobeychik, 2015], especially intensified in recent years where experiments showed in practice the loss of performance of algorithms due to strategic attackers [Nelson et al., 2009, Sommer and Paxson, 2010, Thomas et al., 2013, Wang et al., 2014]. This literature is essentially divided in two parts. The first part studies ‘poisoning attacks’, where the attacker can alter the training set [Globerson and Roweis, 2006, Barreno et al., 2010, Laskov and Lippmann, 2010, Huang et al., 2011, Zhou and Kantarcioglu, 2014]. All of these studies make assumptions on the attacker knowledge and capabilities and propose defenses that are either based on simple ideas such as using robust statistics (which is inefficient against an adaptive adversary) or based on optimization against a worst-case attack (which is very pessimistic and leads to poorly performing algorithms in practice). The second part studies ‘evasion attacks’, where the attacker cannot alter the training set and instead needs to reverse engineer the fixed classifier in order to find a negative instance of minimal cost [Lowd and Meek, 2005, Nelson et al., 2010, Li and Vorobeychik, 2015]. This literature shows that reverse engineering a linear or convex-inducing classifier is ‘algorithmically easy’. To mitigate this, it proposes as an intuitive defense to use random classifiers; but no formal justification is given that would help defining the set of classifiers to use and their probabilities.

Interestingly, the idea of randomizing the defense in order to be less predictable appears very naturally and with formal justifications in the literature on games for security that uses game theory to study security problems. This literature was pioneered in 2003 in the context of intrusion detection [Alpcan and Başar, 2003], and later received a large attention, see for instance [Chen and Leneutre, 2009, Alpcan and Başar, 2010, Tambe, 2011, Manshaei et al., 2013] and the many reference therein. Typically, in these works, the defender has limited resources to defend several assets and decides on where to allocate his resources while the attacker decides on where to attack. Then, by computing the Nash equilibrium, one finds that the defender must randomize to avoid predictability (mixed strategies), and that the equilibrium strategy of the defender depends only on the attacker’s payoff. On the contrary, if the attacker was naive (i.e., using a fixed strategy regardless of the defender’s action), the optimal defender’s strategy would depend only on his own payoff. This illustrates the sharp difference between decision making by optimization (where the attacker is considered naive) and the game-theoretic approach which assumes that the attacker is strategic and adapts to the defender’s action. With this idea in mind, a few data-mining papers applied game theory to learning in adversarial classification scenarios, but using zero-sum games which corresponds to a worst-case assumption [Kantarcioglu et al., 2011, Zhou et al., 2012]. The only exception to our knowledge is the work of [Brückner and Scheffer, 2011, Brückner et al., 2012], but it assumes restrictions on the possible classifiers that do not enable optimizing the classification.

2.1.2 Our contributions

In this context, we proposed a new model to study strategic classification in a nonzero-sum game, in a more flexible framework (not assuming restrictions and more flexible than worst case) [32,33], [9], and provided a complete analysis of the model that reveals intuitive messages on how to perform classification in the presence of an attacker. We present here the model and the main results derived from it.

Model

We model the interaction as a game between a defender who chooses a classifier to distinguish between attacks and normal behavior based on a set of observed features and an attacker who chooses his attack features (class 1 data). Normal behavior (class 0 data) is random and exogenous. The attacker's objective balances the benefit from attacks and the cost of being detected while the defender's objective balances the benefit of a correct attack detection and the cost of false alarm.

Specifically, consider a set $\mathcal{V} \subseteq \mathbb{R}^d$, $d \geq 1$, of possible feature vectors that we assume finite. The set of all possible classifiers is then the set of all possible mappings from the observed attack vector to a classification decision 0 or 1: $C = \{c : \mathcal{V} \rightarrow \{0, 1\}\} = 2^{\mathcal{V}}$. We assume that there is a prior probability $p \in (0, 1)$ of facing an attacker (class 1) and that the non-attacker (class 0) has a fixed distribution P_N on \mathcal{V} . We define the one-shot complete information game, denoted

$$\mathcal{G} = (\mathcal{V}, C, p, c_d, c_{fa}, P_N),$$

as the game between the attacker and the defender where the set of pure actions of the attacker is \mathcal{V} , the set of pure actions of the defender is C , and the payoffs are:

$$U^A(v, c) = R(v) - c_d \mathbb{1}_{c(v)=1}, \quad (2.1)$$

where $R : \mathcal{V} \rightarrow \mathbb{R}_+$ is the ‘‘reward function’’ that describes the gain of the attacker when choosing feature vector v , and c_d is a cost in case of detection; and

$$U^D(v, c) = -R(v) + c_d(v) \mathbb{1}_{c(v)=1} - \frac{1-p}{p} c_{fa} \sum_{v' \in \mathcal{V}} P_N(v') \mathbb{1}_{c(v')=1}, \quad (2.2)$$

where c_{fa} is a cost of false alarm. The defender's payoff can be interpreted as follows. The first component captures the expected loss to an attacker. We assume that this part equals what is gained by the attacker. Then, since the defender interacts with an attacker with chance p , the expected loss is $-pU^A(v, c)$. The second component captures the expected loss due to false alarms. Since the non-attacker is present with chance $1 - p$, the expected false alarm cost is $1 - p$ times the chance that a non-attacker would pick a v that gets classified as an attacker. Finally, the whole payoff function is scaled by the constant $1/p$ for the convenience of having the term $U^A(v, c)$ appear unscaled in the payoff. This scaling does not affect the Nash equilibrium strategies of the game.

In our work, we mainly focus on the Nash equilibrium of game \mathcal{G} . Before describing our results, let us mention that the game is clearly best-response equivalent (as defined in [Rosenthal, 1974]) to a zero-sum game where the defender's payoff is unchanged and the attacker's payoff is $-U^D$. This is because the false alarm part in (2.2), $-\frac{1-p}{p} c_{fa} \sum_{v' \in \mathcal{V}} P_N(v') \mathbb{1}_{c(v')=1}$ does not depend on the attacker's strategy, hence adding it to the attacker's payoff does not change his best response and the Nash equilibrium strategies are therefore the same in the zero-sum game and in the original

game. We will use this equivalence to derive our results (see Theorem 2.1.3 below). It is known that, numerically, zero-sum games can be solved in polynomial time using Linear Programming techniques. Yet, two main problems subsist. First, the size of the defender's action set is very large (of the order of $2^{|\mathcal{V}|}$) since it contains all functions from \mathcal{V} to $\{0, 1\}$, so that a direct numerical solution is not possible in reasonable cases.¹ Second, applying known algorithms to numerically solve zero-sum games does not give intuition on the structure of the equilibrium, whereas our ambition here is to derive an analytical characterization of the equilibrium strategies. The results presented in the next section address both problems.

Results

In [32, 33], we analyzed the model above only for the case where $\mathcal{V} = \mathbb{R}$, that is where the classification is based only on a scalar feature (in [32], the payoff structure is a bit more general than (2.1)). In [9], we give the general solution for any finite \mathcal{V} with utilities defined as in (2.1) and (2.2). We present these general results here and discuss some extensions at the end.

We consider mixed strategies, that is a probability distribution α on \mathcal{V} and β on C , with the standard bilinear extension of the payoffs and definition of the Nash equilibrium (hereafter shortened NE) [Fudenberg and Tirole, 1991].

A first problem when looking for the Nash equilibrium of \mathcal{G} is that the action space C of the defender is very large. Our first result establishes that we can restrict to a much smaller set. Before stating the result, we define the probability of detection function as the probability of class 1 classification (or detection) given the attack vector v and the defender's strategy β :

$$\pi_d^\beta(v) = \sum_{c \in C} \beta_c \mathbb{1}_{c(v)=1}, \quad \forall v \in \mathcal{V}. \quad (2.3)$$

We define the set of threshold classifiers

$$C^T = \{c \in C : c(v) = \mathbb{1}_{R(v) \geq t}, \forall v \in \mathcal{V} \text{ for some } t \in \mathbb{R}\},$$

and assume that $C^T \subseteq C$, which holds for any reasonable C , in particular for $C = 2^{|\mathcal{V}|}$. Threshold classifiers are simple and intuitive classifiers where the defender compares what the attack reward would have been from the observed attack vector to a threshold instead of computing a mapping from any possible attack vector to a detection probability. In our model, we established that it is sufficient to consider only threshold classifiers:

Theorem 2.1.1. *For any NE (α, β) of $\mathcal{G} = (\mathcal{V}, C, p, c_d, c_{fa}, P_N)$, there exists a NE of $\mathcal{G}^T = (\mathcal{V}, C^T, p, c_d, c_{fa}, P_N)$ with the same α and equilibrium payoff pair and the same π_d in the support of the non-attacker's distribution.*

Theorem 2.1.1 shows in particular that, when restricted to using only threshold classifiers, the defender achieves the same equilibrium payoff. Hence, although there may exist Nash equilibria where the defender uses other classifiers, he does not lose anything by using only threshold classifiers. The proof goes through multiple steps detailed in [9]. We first show that the payoffs depend on the defender's strategy β only through the probability of detection function π_d . Then we show that at a NE, the probability of detection is increasing in $R(v)$, that is, a more highly rewarding

¹Note that a direct reduction to an action set $\{0, 1\}$ for the defender (where one would consider functions of v through the maxmin strategy which is known by the maxmin theorem to yield the same value as the minmax) is not possible because, for given actions v and c , the payoffs depend on the entire function c and not only on $c(v)$.

vector is classifier as attacker with a higher probability. Finally, we show that every probability of detection function that is increasing in R can be achieved with a strategy β putting positive weight only on classifiers in C^T .

Theorem 2.1.1 is important because it reduces the strategy space of the defender considerably (making the game solvable), but also because it reveals an intuitive message on how classification should be performed in this adversarial setting: the defender should use classifiers that correspond to thresholds on the attacker's reward. This is in contrast with standard classifiers such as logistic regression or SVM which use a fixed shape of the decision boundary (in some space) and are therefore bound to be suboptimal (unless by luck the reward function happens to have the pre-defined shape). Instead here, the decision boundary should be adapted to the attacker's reward function.

To complete the analysis, we provided an efficient algorithm to compute all Nash equilibria and a compact characterization of the possible forms of a Nash equilibrium. We first showed that we can also reduce the attacker's strategy space to the set

$$\mathcal{V}^R = \{r \in \mathbb{R}_+ : r = R(v), \text{ for some } v \in \mathcal{V}\} \quad (2.4)$$

of all unique reward values. We denote by $r_i, i \in \{1, \dots, |\mathcal{V}^R|\}$ the elements of this set and assume that they are ordered. Defining the non-attacker's reduced probability measure on \mathcal{V}^R by

$$P_N^R(r) = \sum_{v'} P_N(v' \in \mathcal{V}) \mathbb{1}_{R(v')=r}, \quad (2.5)$$

the result can be stated as follows.

Proposition 2.1.2. *If (α, β) is a NE of $\mathcal{G}^T = (\mathcal{V}, C^T, p, c_d, c_{fa}, P_N)$, then (α^*, β) is a NE of $\mathcal{G}^{R,T} = (\mathcal{V}^R, C^T, p, c_d, c_{fa}, P_N^R)$ with the same equilibrium payoff pair where $\alpha_{r_i}^* = \sum_{v_j \in \mathcal{V}, R(v_j)=r_i} \alpha_{v_j}, \forall r_i \in \mathcal{V}^R$.*

Note that, although \mathcal{V}^R is not rigorously a subset of \mathcal{V} , \mathcal{V}^R is a reduced strategy space in the sense that R is clearly a surjection from \mathcal{V} to \mathcal{V}^R . Moreover, since β is a probability on C^T , any two attack vectors with the same reward have the same probability of detection, so that by abuse of notation we can define the probability of detection function as a function of the reward by $\pi_d(r) := \pi_d(v)$, where $r = R(v)$. Then, utilities in $\mathcal{G}^{R,T}$ are defined by adapting (2.1)-(2.2) in the obvious way.

The proof of Proposition 2.1.2 is given in [9] where we also provide an easy way from the NE (α^*, β) of $\mathcal{G}^{R,T}$ to recover a NE (α, β) of \mathcal{G}^T . We can therefore focus on the simpler problem of computing the NE of $\mathcal{G}^{R,T}$, which is given in our last main result:

Theorem 2.1.3. *Algorithm 1 finds all NE of the classification game $\mathcal{G}^{R,T}$. Moreover, if (α, β) is a NE, then, there exists $k \in \{1, \dots, |\mathcal{V}^R|\}$ such that*

$$\begin{aligned} \beta &= (0, \dots, 0, \beta_k, \dots, \beta_{|\mathcal{V}^R|}, \beta_{|\mathcal{V}^R|+1}), \\ \alpha &= (0, \dots, 0, \alpha_k, \dots, \alpha_{|\mathcal{V}^R|}), \end{aligned}$$

where

$$\beta_i = \frac{r_i - r_{i-1}}{c_d}, \quad \forall i \in \{k+1, \dots, |\mathcal{V}^R|\}, \quad (2.6)$$

$$\alpha_i = \frac{1-p}{p} \frac{c_{fa}}{c_d} P_N^R(r_i), \quad \forall i \in \{k+1, \dots, |\mathcal{V}^R| - 1\}, \quad (2.7)$$

and $\beta_k, \beta_{|\mathcal{V}^R|+1} \geq 0$ and $\alpha_k, \alpha_{|\mathcal{V}^R|} \geq 0$ are such that

- (i) $\beta_k \in (0, \frac{r_k - r_{k-1}}{c_d})$, $\beta_{|V^R|+1} = 0$, and α_k satisfies (2.7), $\alpha_{|V^R|} > 0$; or
- (ii) $\beta_k = 0$, $\beta_{|V^R|+1} > 0$, and $\alpha_k \in (0, \frac{1-p}{p} \frac{c_{fa}}{c_d} P_N^R(r_k))$, $\alpha_{|V^R|}$ satisfies (2.7); or
- (iii) $\beta_k = 0$, $\beta_{|V^R|+1} = 0$, and $\alpha_k \in [0, \min(\frac{1-p}{p} \frac{c_{fa}}{c_d} P_N^R(r_k), 1 - \sum_{i=k+1}^{|V^R|-1} \alpha_i)]$, $\alpha_{|V^R|} \geq 0$; or
- (iv) $\beta_k \in [0, \frac{r_k - r_{k-1}}{c_d}]$, $\beta_{|V^R|+1} \geq 0$, and $\alpha_k, \alpha_{|V^R|}$ satisfy (2.7).

Theorem 2.1.3 provides both an algorithm (Algorithm 1) that finds all NE and a compact characterization of the restricted number of possible forms that a NE can have. Interestingly, we observe that the defender assigns a weight to a reward r_i that is positive and proportional to the marginal reward increase at that point, on a support that goes until the highest reward $r_{|V^R|}$. This is somewhat counter-intuitive as it implies that the defender includes at NE with positive weights classifiers that almost never classify as attacker even for a high reward and even if the probability that a non-attacker uses this reward is arbitrarily small. We also observe in Theorem 2.1.3 the fact that the attacker mimics the non-attacker's distribution (proportionally) on a support that corresponds to the defender's support.

The proof of Theorem 2.1.3 relies on the key observation mentioned above that the game $\mathcal{G}^{R,T}$ is best-response equivalent (see [Rosenthal, 1974]) to a zero-sum game where the defender's payoff is unchanged and the attacker's payoff is $-U^D$. This is because the false alarm part in (2.2), $-\frac{1-p}{p} c_{fa} \sum_{v' \in \mathcal{V}} P_N(v') \mathbb{1}_{c(v')=1}$ does not depend on the attacker's strategy, hence adding it to the attacker's payoff does not change his best response. Then, the NE strategies of the defender and attacker are solutions of dual Linear Programs and we compute them by looking at the extreme points of a well defined polyhedron. The multiplicity of NE observed in special cases in Theorem 2.1.3 is also in accordance with the relationship between degeneracy and multiplicity of the primal and the dual optimal solutions of a Linear Program. Note finally that there are known algorithms to compute the solution of a Linear Program in polynomial time. Our algorithm (Algorithm 1) also runs in polynomial time but its advantage lies in the intuition about the solutions that it provides (i.e., the analytical characterization).

We also explored qualitatively and quantitatively the impact of the non-attacker and underlying parameters on the equilibrium strategies. In particular, we discuss in [9] how to compute using the NE found above the benefit of acquiring new features to do the detection (e.g., installing a new sensor), which should eventually be compared to the investment cost.

Extensions

The model presented above was first introduced in [32] in the case of a single feature (scalar attack vector). In [33], we also studied the case of a single feature, but with a more general payoff model where the cost of detection c_d can depend on v . In this paper, we also restrict the defender to threshold classifiers, but without proof at that time. Then we show that, under conditions of discrete concavity of $R(v) - c_d(v)$ and of $R(v)$, we still can compute the NE using similar LP based methods and the defender's NE strategy still has a contiguous block of non-zero weights (actually, of tight inequality constraints) until the highest threshold.

Algorithm 1: How to compute the NE (α, β)

```

for  $type = 1, 2$  do
  | construct  $\beta$  for  $s \in \{1, \dots, |\mathcal{V}^R|\}$  using Algorithm 2
  | find  $(\beta_{1,2}, s_{1,2}^*)$  that maximize  $U_{1,2}^D$ 
if  $U_1^D > U_2^D$  then
  |  $\beta \leftarrow \text{compute-}\beta(s_1^*, 1)$ ;  $\alpha \leftarrow \text{compute-}\alpha(s_1^*)$ 
if  $U_1^D < U_2^D$  then
  | if  $s_2^*$  is unique then
  | |  $\beta \leftarrow \text{compute-}\beta(s_2^*, 2)$ ;  $\alpha \leftarrow \text{compute-}\alpha(s_2^*)$ 
  | else
  | | // denote  $s_{2a}^*$  and  $s_{2b}^* = s_{2a}^* + 1$  the 2 solutions
  | |  $\beta_a \leftarrow \text{compute-}\beta(s_{2a}^*, 2)$ ;  $\beta_b \leftarrow \text{compute-}\beta(s_{2b}^*, 2)$ 
  | |  $\beta \leftarrow \text{convex hull of } \beta_a, \beta_b$ 
  | |  $\alpha \leftarrow \text{compute-}\alpha(s_{2b}^*)$ 
if  $U_1^D = U_2^D$  then
  | if  $s_2^*$  is unique then
  | | //  $s_2^* = s_1^*$ 
  | |  $\beta_1 \leftarrow \text{compute-}\beta(s_1^*, 1)$ ;  $\beta_2 \leftarrow \text{compute-}\beta(s_1^*, 2)$ 
  | | if  $\beta_1 \neq \beta_2$  then
  | | |  $\beta \leftarrow \text{convex hull of } \beta_1, \beta_2$ 
  | | else
  | | |  $\beta \leftarrow \beta_1$ 
  | |  $\alpha \leftarrow \text{compute-}\alpha(s_1^*)$ 
  | else
  | | // denote  $s_{2a}^*$  and  $s_{2b}^* = s_{2a}^* + 1$  the 2 solutions
  | | // the type I and type IIa solutions are identical
  | |  $\beta_a \leftarrow \text{compute-}\beta(s_{2a}^*, 1)$ ;  $\beta_b \leftarrow \text{compute-}\beta(s_{2b}^*, 2)$ 
  | |  $\beta \leftarrow \text{convex hull of } \beta_a, \beta_b$ 
  | |  $\alpha \leftarrow \text{compute-}\alpha(s_{2b}^*)$ 

```

Algorithm 2: Compute- $\beta(s, type)$

```

for  $i = 1$  to  $s - 1$  do
  |  $\beta_i \leftarrow 0$ 
for  $i = s + 1$  to  $|\mathcal{V}^R|$  do
  |  $\beta_i \leftarrow \frac{r_i - r_{i-1}}{c_d}$ 
 $\beta_s \leftarrow \mathbf{1}_{type=1}(1 - \sum_{s+1}^{|\mathcal{V}^R|} \beta_i)$ 
 $\beta_{|\mathcal{V}^R|+1} \leftarrow \mathbf{1}_{type=2}(1 - \sum_{s+1}^{|\mathcal{V}^R|} \beta_i)$ 
 $U_{type}^D \leftarrow \min[\Lambda\beta] - \mu' \beta$ 

```

Algorithm 3: Compute- $\alpha(\beta, k)$

```

//  $\beta$  is the set of all convex combinations if multiple
for  $i = 1$  to  $k - 1$  do
   $\alpha_i \leftarrow 0$ 
for  $i = k + 1$  to  $|\mathcal{V}^R| - 1$  do
   $\alpha_i \leftarrow \frac{1-p}{p} \frac{c_{fa}}{c_d} P_N^R(r_i)$ 
if  $\beta_k > 0$  for some  $\beta$  then
   $\alpha_k \leftarrow \frac{1-p}{p} \frac{c_{fa}}{c_d} P_N^R(r_k)$ 
   $\alpha_{|\mathcal{V}^R|} \leftarrow 1 - \sum_{i=k}^{|\mathcal{V}^R|-1} \alpha_i$ 
else if  $\beta_{|\mathcal{V}^R|+1} > 0$  for some  $\beta$  then
   $\alpha_{|\mathcal{V}^R|} \leftarrow \frac{1-p}{p} \frac{c_{da}}{c_d} P_N^R(r_{|\mathcal{V}^R|})$ 
   $\alpha_k \leftarrow 1 - \sum_{i=k+1}^{|\mathcal{V}^R|} \alpha_i$ ;
else if  $\beta_k = \beta_{|\mathcal{V}^R|+1} = 0$  then
  for  $i = 1$  to  $k - 1$  and  $i = k + 1$  to  $|\mathcal{V}^R| - 1$  do
     $\alpha_i^1 \leftarrow \alpha_i, \alpha_i^2 \leftarrow \alpha_i$ 
   $\alpha_k^1 \leftarrow 0$ 
   $\alpha_k^2 \leftarrow \min\left(\frac{1-p}{p} \frac{c_{fa}}{c_d} P_N^R(r_k), 1 - \sum_{i=k+1}^{|\mathcal{V}^R|-1} \alpha_i\right)$ 
   $\alpha_{|\mathcal{V}^R|}^{1,2} \leftarrow 1 - \sum_{i=k}^{|\mathcal{V}^R|-1} \alpha_i^{1,2}$ 
   $\alpha \leftarrow$  convex hull of  $\alpha^1, \alpha^2$ 

```

Finally, in recent work (not yet published), we extended the model to a Bayesian game that takes into account the incomplete information of the defender on the attacker's payoff. We assumed that the attacker has a type $\theta_a \in \Theta_A$ which defines his reward function $R_{\theta_a}(v)$. Then, the payoffs are defined as before. There is a probability P_{Θ_A} on the attacker's type that is common knowledge. For this game, we looked for the Bayesian Nash equilibrium (BNE). On the attacker's side, we showed that at BNE, the expected strategy across all types is the same as found in the complete information game of before. We are currently analyzing the defender sides to see if we can identify a small set of strategies which are sufficient as in the complete information case.

Summary of our contribution to adversarial classification:

We propose a nonzero-sum game model between an attacker choosing class 1 feature values and a defender choosing the classifier. We compute all NE and give a compact characterization revealing that classification should be done in a new way: by using a threshold on the attacker's reward. We qualitatively and quantitatively explore the system and propose extensions to more general payoffs and to incomplete information.

2.2 Security resource allocation

Even before the problem of classification, i.e., making the correct decision upon observation of an event, security often relies on the appropriate allocation of a finite amount of defense resources

(be it human or computer resources) to monitor potential security attacks. We now focus on this question of security resource allocation.

2.2.1 Context

Approaches to defense resource allocations for intrusion detection have been mentioned in the previous section, but they all consider the objective of detecting a single intrusion. For more complex systems that contains several elements, a more global approach is needed to minimize the number of compromised elements. Such approaches have been proposed in different contexts such as airport security, see for instance the survey [Tambe, 2011]. Those, however, consider Stackelberg games where the defender's strategy is known to the attacker, which does not apply well to certain security scenarios such as cybersecurity. A key problem in this area is that all strategic resource allocation games in simultaneous moves are derivative of the famous Colonel Blotto game introduced by Borel in 1921 [Borel, 1921]. In the Colonel Blotto game, two players allocate an exogenously given amount of resources to a fixed number of battlefields with given values (corresponding to elements of the infrastructure). Each battlefield is then won by the player who allocated more resources to it, and each player maximizes the aggregate value of battlefields he wins.

Despite its apparent simplicity, the Colonel Blotto game is very intricate. The first solution was given by Borel and Ville in 1938 [Borel and Ville, 1938] for the case of two players with symmetric resources and two battlefields. In 1950, Gross and Wagner [Gross and Wagner, 1950] extended this solution to the game with two players with asymmetric (unequal) resources and two battlefields, and to the game with two player with symmetric resources and three battlefields. The case of asymmetric player resources and an arbitrary number of battlefields remained unsolved until 2006. Only then, a complete Nash equilibrium solution was given for the case of asymmetric player resources and an arbitrary number of battlefields by Roberson [Roberson, 2006], but only if all battlefields have the same value, which is not realistic in some applications.

Since Roberson's seminal contribution, a number of variants of the Colonel Blotto game have been proposed and studied, e.g., versions that consider sequential moves instead of simultaneous moves [Powell, 2009, Rinott et al., 2012], coalitional versions [Kovenock and Roberson, 2012a], versions with endogeneous resources [Kovenock et al., 2010] or graph dependences [Masucci and Silva, 2014], etc. (see also [Kvasov, 2007, Arce et al., 2012] and a survey in [Roberson, 2010]). Yet, with the exception of [Hortala-Vallve and Llorente-Saguer, 2012] which only considers very partial solutions where a pure equilibrium exist, all studies with asymmetric player resources and an arbitrary number of battlefields have been restricted to the case of identical battlefields.

2.2.2 Our contributions

In this context, we propose a solution of the Colonel Blotto game for asymmetric players and an arbitrary number of battlefields, under a restriction on the minimum number of battlefields with every unique value. We describe below the main results; all details can be found in [30].

Model and formal definitions

We denote by A and B the two players and by $X_A \in \mathbb{R}_+$ and $X_B \in \mathbb{R}_+$ their respective resources. Without loss of generality, we assume that $X_A \leq X_B$. Let n be an integer denoting the number of battlefields. Each battlefield $j \in \{1, \dots, n\}$ is endowed with a value $v^j \in \mathbb{R}_+$. We denote by

$\mathbf{v} = (v^1, \dots, v^n)$ the vector of battlefield values, and by $V = \sum_{j=1}^n v^j$ the aggregate value of all battlefields.

Players choose how to allocate their resources between the battlefields, i.e., each player chooses how to distribute his (given) resources between the battlefields. Formally, a pure strategy of player $p \in \{A, B\}$ is a vector $\mathbf{x}_p = (x_p^j)_{j \in \{1, \dots, n\}} \in \mathbb{R}_+^n$ satisfying the budget constraint $\sum_{j=1}^n x_p^j \leq X_p$, where $x_p^j \in \mathbb{R}_+$ denotes the amount of resources allocated to battlefield j by player p . We let S_p denote the set of pure strategies for player p :

$$S_p = \left\{ \mathbf{x}_p \in \mathbb{R}_+^n : \sum_{j=1}^n x_p^j \leq X_p \right\}, \quad (p \in \{A, B\}).$$

For each battlefield $j \in \{1, \dots, n\}$, the player who dedicates the highest amount of resources wins this entire battlefield. In case of a tie, player B (the player with higher total amount of resources) wins the battlefield. Hence, for each battlefield $j \in \{1, \dots, n\}$, if $x_A^j > x_B^j$ then player A wins the battlefield, and if $x_A^j \leq x_B^j$ then player B wins it. Each player's payoff equals to the sum of the values of the battlefields that he wins.

Let $\mathfrak{B}(X_A, X_B, \mathbf{v})$ denote the above presented Blotto game. The game \mathfrak{B} is a one-shot game in which players A and B simultaneously choose their allocation of forces to the battlefields to maximize the total value of the battlefields they win. The game \mathfrak{B} is a complete information game. The parameters of the game (X_A, X_B, n and \mathbf{v}), players' action spaces and objectives are common knowledge.

In most cases of interest, that is when for each player the expected payoff is strictly positive, no pure strategy Nash equilibrium exists. We therefore will focus on mixed strategy equilibria. A mixed strategy for player $p \in \{A, B\}$ is an n -variate distribution $P_p : \mathbb{R}_+^n \rightarrow [0, 1]$ whose support is contained in S_p . For a given n -variate distribution, we denote by $F_p^j : \mathbb{R}_+ \rightarrow [0, 1]$ the univariate marginal distribution of resources allocated by player p to battlefield j .

Our results will require that we can make groups of battlefields of the same value with a large enough number of battlefield in each group. To precisely state these conditions and establish our framework, we introduce the following definitions. Let k denote the number of different battlefield values, and let $\{w_1, \dots, w_k\}$ denote the corresponding set of unique battlefield values. Formally we have $w_{i_1} \neq w_{i_2}$ for all $i_1 \neq i_2$ and, for all $j \in \{1, \dots, n\}$ there exists $i \in \{1, \dots, k\}$ such that $v^j = w_i$. For $i \in \{1, \dots, k\}$, define $C(i)$ as the set of battlefields with value w_i , i.e., $C(i) = \{j \in \{1, \dots, n\} : v^j = w_i\}$. Define, for each $i \in \{1, \dots, k\}$, $n_i = \#C(i)$ the number of battlefields of value w_i (again, $n_i \geq 2$ for all i) and $V_i = \sum_{j \in C(i)} v^j = n_i w_i$ the aggregate value of all battlefields of value w_i . Note that $\sum_{i=1}^k V_i = V$.

Results

With the notation introduced above, we can state our results. We start by characterizing the unique equilibrium marginal distributions for the game $\mathfrak{B}(X_A, X_B, \mathbf{v})$:

Theorem 2.2.1. *Assume that, for all groups of battlefields $i \in \{1, \dots, k\}$, we have*

$$\frac{2}{n_i} < \frac{X_A}{X_B} \leq 1. \quad (2.8)$$

Then, in equilibrium, each player allocates resources with the following unique univariate marginal distribution functions $\forall j \in \{1, \dots, n\}$:

(i) For player A:

$$F_A^j(x) = \left(1 - \frac{X_A}{X_B}\right) + \frac{x}{\frac{2v^j}{V}X_B} \left(\frac{X_A}{X_B}\right), \quad x \in [0, \frac{2v^j}{V}X_B]; \quad (2.9a)$$

(ii) For player B:

$$F_B^j(x) = \frac{x}{\frac{2v^j}{V}X_B}, \quad x \in [0, \frac{2v^j}{V}X_B]. \quad (2.9b)$$

The proof works by establishing a one-to-one correspondence between the solution of the Blotto game and the solution of n independent all-pay auctions, with the later established in [Hillman and Riley, 1989, Baye et al., 1996]. Theorem 2.2.1 (and the subsequent results) requires that condition (2.8) holds. First note that it implies that $n_i \geq 3$, i.e., each group of battlefields has at least three battlefields of the same value. This ensures that we will be able to construct an n -variate distribution with the marginals in Theorem 2.2.1 (see Theorem 2.2.3 below and the discussion that follows); which is necessary to prove Theorem 2.2.1. Condition (2.8) also restricts the disparity of players resources. However, even for a large asymmetry, this condition will be satisfied as soon as the number of battlefields is large enough in each group. We note that such an assumption is often made in the literature to restrict the complexity of the analysis, for instance in [Kovenock and Roberson, 2012b]. The cases in which condition (2.8) is not met can be studied separately but are of limited practical interest.

Under the condition of Theorem 2.2.1, we have unique equilibrium marginals. These marginals are uniform, as in the game with identical battlefield values, but now the marginal's support is proportional to the battlefields value. Theorem 2.2.1 allows us to obtain the equilibrium player payoffs:

Corollary 2.2.2. *Under condition (2.8) of Theorem 2.2.1, in equilibrium, player A and B expected payoffs are $V \frac{X_A}{2X_B}$ and $V \left(1 - \frac{X_A}{2X_B}\right)$, respectively.*

Remarkably, Corollary 2.2.2 yields equilibrium payoffs identical the payoffs in the game with equal battlefield values.

Theorem 2.2.1 describes only the marginal distributions. So far, we merely hypothesized the existence of some n -variate distribution with such marginals that respects player resource constraints. Next, we establish the existence of such n -variate distribution.

Theorem 2.2.3. *Under condition (2.8) of Theorem 2.2.1, for each player $p \in \{A, B\}$, there exists an n -variate distribution with support contained in S_p such that the marginals are given by (2.9a) - (2.9b) for all battlefields $j \in \{1, \dots, n\}$.*

The proof is done by constructing an n -variate distribution with the correct marginals that respects the players budget constraints. Roughly speaking, the construction consists of two steps. First, we make a deterministic allocation to each group of battlefields with the same value. The amount of allocated resources is proportional to the aggregate value of the group. Second, within each group the randomization is done as in Roberson [Roberson, 2006]. This is possible because, by our assumption, each group has enough battlefields. Our construction allows to obtain an equilibrium n -variate distribution with the correct marginals respecting the budget constraints. Still, there may exist other solutions that randomize the global resource allocation between the different groups of battlefields with a common value.

We are currently considering various extensions of this work, in particular to remove the restriction, to find easier solutions in the limit of a large number of battlefields, and to analyze the case where the value of a battlefield is not the same to each player (non constant-sum).

Summary of our contribution to security resource allocation:

We give a solution of the Colonel Blotto game for asymmetric players and an arbitrary number of battlefields, under a restriction on the minimum number of battlefields of each unique value.

2.3 Adversarial online decision making with discounted losses

Finally, in many real-world situations, the defender is uncertain about the payoff or the behavior of the attacker she is facing. Many of these situations can be modeled as online decision making problems where the defender sequentially makes decisions based on previous observations. This raises the question of how to select actions which we tackle here in the specific context of discounted losses.

2.3.1 Context

In many scenarios involving repeated decision-making in uncertain environments, one desires robust performance guarantees. The well-known regret minimization paradigm captures this objective by formulating the problem as a repeated game between the decision-maker and the environment, which is modeled as an adversary. The goal of the decision-maker in this game is to design an adaptive and possibly randomized algorithm for choosing actions (henceforth, algorithm/strategy/policy) so as to minimize *regret*, which is defined as the difference between the expected loss incurred by the algorithm, and the loss incurred by the best fixed action that could have been chosen in hindsight against the sequence of actions chosen by the adversary. In this setting, several *no-regret* algorithms are now known, which ensure that the time-averaged regret asymptotically vanishes in the long run, not just in expectation, but with probability 1, irrespective of the sequence of actions chosen by the adversary. The first study of regret minimization in repeated games dates back to the pioneering work of Hannan [Hannan, 1957], who introduced the notion of regret optimality in repeated games and proposed the earliest known no-regret algorithm. Since then, numerous other such algorithms have been proposed, particularly for the problem of prediction using expert advice, see [Littlestone and Warmuth, 1994, Vovk, 1990, Cesa-Bianchi et al., 1997, Freund and Schapire, 1999], one particularly well-known being the multiplicative weights update algorithm, also known as Hedge. Other settings with limited feedback have been considered, most notably the multi-armed bandit setting [Auer et al., 2002, Bubeck and Cesa-Bianchi, 2012]. Stronger notions of regret such as internal regret, have also been studied [Foster and Vohra, 1997, Cesa-Bianchi and Lugosi, 2003, Blum and Mansour, 2005, Stoltz and Lugosi, 2005].

In many realistic cases, however (e.g., when decision-making horizons are finite), the minimal expected regret one can achieve in the worst case over all the sequence of actions chosen by the adversary (henceforth, just optimal regret) is non-zero. This is also the case when present losses are more important than future one (modeled as discounted losses), which is a realistic case on which we focus in our work. In such cases, standard no-regret algorithms can perform quite poorly compared to the optimum. Indeed, regret minimization with non-uniformly weighted losses, of which the discounted loss is a special case, has been considered before. While the average regret goes to zero if the weights satisfy a non-summability condition, lower bounds exist ([Cesa-Bianchi and Lugosi, 2003], Thm 2.7) that show that the optimal regret is bounded away

from 0 if the weights are summable, which is the case with discounting. Natural extensions of no-regret algorithms incur a regret of $O(\sqrt{1-\beta})$ in this case (where β , in this section, denotes the discount factor), for instance see [Cesa-Bianchi and Lugosi, 2003], Thm 2.8 and [Perchet, 2014], Prop. 2.22. ([Chernov and Zhdanov, 2010] derive better bounds for the case where future losses are given a higher weight than current ones.)

In most cases with finite horizon or discounted losses, the exact optimal regret and strategies are unknown, except for a few special cases. In an early work, [Cover, 1966] gave the optimal algorithm for the problem of prediction using expert advice over any finite horizon T , for the case of 2 experts, and where the losses are $\{0, 1\}$. [Gravin et al., 2014] recently extended the result to the case of 3 experts for both the finite horizon and geometrically distributed random horizon problems.² [Abernethy et al., 2008] considered a related problem, where a gambler places bets from a finite budget repeatedly on a fixed menu of events, the outcomes of which are adversarially chosen from $\{0, 1\}$ (you win or you lose), and characterized the minmax optimal strategies for the gambler and the adversary. [Luo and Schapire, 2013] considered a similar repeated decision-making problem where an adversary is restricted to pick loss vectors (i.e., a loss for each action of the decision-maker in a stage) from a set of basis vectors, and characterized the minmax optimal strategy for the decision-maker under both, a fixed and an unknown horizon. Most of the approaches in these works are specific to their settings, and exploit the assumptions on the structure of the loss vectors.³ All of these are examples where games with finite action spaces (called matrix games) are repeated, which is the setting that we are concerned with. There are also works that consider exact minmax optimality in repeated games with continuous action spaces, with specific types of loss functions, see [Koolen et al., 2014, Bartlett et al., 2015, Koolen et al., 2015] and references therein. In general, however, if the loss vectors are arbitrary, none of these approaches can be extended and indeed it is recognized that characterizing the optimal regret and algorithm is difficult, cf. [Luo and Schapire, 2013].

2.3.2 Our contributions

In this context, our main contribution was the proposal of a systematic set-valued dynamic programming approach for designing regret-optimal strategies in repeated games with discounted losses. In these games, the loss criterion is the weighted sum of per-stage losses, with the loss at stage $t \geq 1$ weighted by β^{t-1} , where $\beta \in (0, 1)$ is the discount factor. Such discounting is natural in practice, where minimizing current losses is more important than the ones in the future. As mentioned above, in this case, since the losses incurred in the initial stages have a non-vanishing contribution to the cumulative loss as the number of stages increases, the optimal long-run average discounted regret for any fixed $\beta \in (0, 1)$ is non-zero (see [Cesa-Bianchi and Lugosi, 2003], Thm 2.7).⁴ Several known no-regret algorithms guarantee an average regret of $O(\sqrt{1-\beta})$ asymptotically in this setting, but the performance of these algorithms can be far from optimal for a fixed β . Our approach on the other hand gives a procedure to compute ϵ -regret-optimal strategies in these games for any $\epsilon > 0$. These strategies are extremely simple to implement and require a

²Although a geometric time horizon model seems to be related to the infinite horizon model with discounted losses, the two problem formulations define regret differently, and thus lead to different optimal regrets.

³Many of these works rely on a particular nice property of these settings, which is that the optimal strategy of the adversary is a controlled random walk that makes any algorithm incur the same regret. If the losses are simple, for instance if they are the basis vectors, then this random walk can be exactly analyzed to compute the optimal regret. Knowing the optimal regret then simplifies the computation of the optimal strategy of the decision-maker.

⁴By average discounted regret we mean that the weights are normalized by multiplying with $(1-\beta)$, which ensures that the sum of the weights over an infinite horizon is 1.

finite memory (which grows as ϵ decreases). For instance, using our approach, we are able to design a provably near-optimal algorithm for the problem of prediction using expert advice with discounted losses, for the case of 2 experts. To the best of our knowledge, no such algorithm was known before.

Our solution begins with a standard approach in regret minimization, of transforming the repeated game into a repeated game with vector losses. In this game, the number of vector components is the number of actions available to the decision-maker, where each component keeps track of the additional loss incurred relative to the loss incurred if the corresponding action was always chosen in the past. The goal of regret minimization in the original game now translates to the goal of *simultaneously* minimizing the worst-case expected losses on all the components in this vector-valued game. In fact there is a tradeoff here: a better guarantee on one component implies a worse guarantee on some other, and we consider the more ambitious objective of characterizing the entire Pareto-frontier of the minimal losses that can be simultaneously guaranteed across the different components. Our main technical contribution is an effective characterization of this set as the fixed point of a set-valued dynamic programming operator, which simultaneously also characterizes the strategies that achieve the different points on it. This characterization then allows us to design an iterative scheme to approximate this set and compute approximately optimal strategies for arbitrarily low approximation error.

In what follows, we describe the main elements of the model, approach and results. All details can be found in [8].

Model

Let G be a two player game with m actions $A = \{1, \dots, m\}$ for player 1, who is assumed to be the minimizer and who we will call Alice (the decision-maker), and n actions $B = \{1, \dots, n\}$ for player 2, who is the adversary and who we will call Bob. For each pair of actions $a \in A$ and $b \in B$, the corresponding loss for Alice is $l(a, b) \in \mathbb{R}$. The losses for different pairs of actions are known to Alice. The game G is played repeatedly for T stages $t = 1, 2, \dots, T$. In each stage, both Alice and Bob simultaneously pick their actions $a_t \in A$ and $b_t \in B$ and Alice incurs the corresponding loss $l(a_t, b_t)$. The loss of the repeated game is defined to be the total discounted loss given by $\sum_{t=1}^T \beta^{t-1} l(a_t, b_t)$, where $\beta \in (0, 1)$. We define the total discounted regret of Alice as:

$$\sum_{t=1}^T \beta^{t-1} l(a_t, b_t) - \min_{a \in A} \sum_{t=1}^T \beta^{t-1} l(a, b_t), \quad (2.10)$$

which is the difference between her actual discounted loss, and the loss corresponding to the single best action that could have been chosen against the sequence of actions chosen by Bob in hindsight. An adaptive randomized strategy π_A for Alice specifies for each stage t , a mapping from the set of observations till stage t , i.e., $H_t = (a_1, b_1, \dots, a_{t-1}, b_{t-1})$, to a probability distribution on the action set A , denoted by $\Delta(A)$. Let Π_A be the set of all such policies of Alice.

The adversary Bob is assumed to choose a deterministic oblivious strategy, i.e., his choice is simply a sequence of actions $\pi_B = (b_1, b_2, b_3, \dots, b^T)$ chosen before the start of the game. Let Π_B be the set of all such sequences.⁵ We would like to compute the worst case or minmax *expected*

⁵Having an oblivious adversary is a standard assumption in regret-minimization literature [Cesa-Bianchi and Lugosi, 2003] and in fact it is known that in this case, an oblivious adversary is as powerful as a non-oblivious (adaptive) adversary.

discounted regret which is defined as:

$$\min_{\pi_A \in \Pi_A} \max_{\pi_B \in \Pi_B} \left(E_{\pi_A} \left[\sum_{t=1}^T \beta^{t-1} l(a_t, b_t) \right] - \min_{a \in A} \sum_{t=1}^T \beta^{t-1} l(a, b_t) \right), \quad (2.11)$$

and the strategy for Alice that guarantees this value. Here the expectation is over the randomness in Alice's strategy. Here one can see that there is no loss of generality in assuming that the adversary is deterministic. Indeed if Π_B is allowed to be the set of possible randomizations over T-length sequences of Bob's actions, the optimal policy of Bob in the problem

$$\max_{\pi_B \in \Pi_B} E_{\pi_A, \pi_B} \left[\sum_{t=1}^T \beta^{t-1} l(a_t, b_t) - \min_{a \in A} \sum_{t=1}^T \beta^{t-1} l(a, b_t) \right]$$

is a deterministic sequence.

We can now equivalently write (2.11) as:

$$\min_{\pi_A \in \Pi_A} \max_{\pi_B \in \Pi_B} \max_{a \in A} E_{\pi_A} \left[\sum_{t=1}^T \beta^{t-1} (l(a_t, b_t) - l(a, b_t)) \right]. \quad (2.12)$$

In order to address this objective, it is convenient to define a vector-valued game \mathbb{G} , in which, for a pair of actions $a \in A$ and $b \in B$, the vector of losses is $\mathbf{r}(a, b)$ with m components (recall that $|A| = m$), where

$$r_k(a, b) = l(a, b) - l(k, b) \quad (2.13)$$

for $k = 1, \dots, m$. $r_k(a, b)$ is the single-stage additional loss that Alice bears by choosing action a instead of action k , when Bob chooses b : the so called single-stage regret with respect to action k . For a choice of strategies $\pi_A \in \Pi_A$ and $\pi_B \in \Pi_B$ of the two players, the expected loss on component k in this vector-valued repeated game over horizon T is given by

$$R_k^T(\pi_A, \pi_B) = E_{\pi_A} \left[\sum_{t=1}^T \beta^{t-1} r_k(a_t, b_t) \right], \quad (2.14)$$

where the expectation is over the randomness in Alice's strategy. Now observe that by playing a fixed policy $\pi_A \in \Pi_A$, irrespective of the strategy chosen by Bob, Alice guarantees that the total expected loss on component k is no more than $\max_{\pi_B^k \in \Pi_B} R_k^T(\pi_A, \pi_B^k)$. Suppose that we determine the set of all *simultaneous guarantees* that correspond to *all* the strategies $\pi_A \in \Pi_A$, defined as:

$$\mathcal{W}^T \triangleq \left\{ \left(\max_{\pi_B^k \in \Pi_B} R_k^T(\pi_A, \pi_B^k) \right)_{k=1, \dots, m} : \pi_A \in \Pi_A \right\}. \quad (2.15)$$

Then it is clear that

$$\min_{\pi_A \in \Pi_A} \max_{\pi_B \in \Pi_B} \max_{a \in A} E_{\pi_A} \left[\sum_{t=1}^T \beta^{t-1} (l(a_t, b_t) - l(a, b_t)) \right] = \min_{\mathbf{x} \in \mathcal{W}^T} \max_k x_k.$$

In fact, we are only interested in finding the *minimal* points in the set \mathcal{W}^T , i.e., its *Lower Pareto frontier*, which is the defined as the set

$$\Lambda(\mathcal{W}^T) \triangleq \{ \mathbf{x} \in \mathcal{W}^T : \forall \mathbf{x}' \in \mathcal{W}^T \setminus \{ \mathbf{x} \}, \exists k \text{ s.t. } x_k < x'_k \}, \quad (2.16)$$

since all other points are strictly sub-optimal. Let this set be denoted as \mathcal{V}^T . Our goal in this work is to characterize and compute the set \mathcal{V}^∞ that can be achieved in the infinite horizon game and compute policies for Alice in Π^A that guarantee different points in it.

In all of this work, we consider in fact a more general model for a vector-valued repeated game \mathbb{G} with m actions for Alice and n actions for Bob, where m and n are arbitrary but finite, but we assume that the number of components in the vector of losses $\mathbf{r}(a, b)$ is K , where K may be different from m . Our results will hold for any vector-valued repeated game, not just the one that arises from the regret minimization formulation discussed before.

Overview of the results

Our goal is to compute the set of minimal guarantees \mathcal{V}^∞ that Alice can achieve in the infinite horizon vector-valued discounted repeated game and characterize the policies that achieve it. Our main results are as follows:

1. We show that the set \mathcal{V}^∞ of *minimal* losses that Alice can simultaneously guarantee in an infinitely repeated vector-valued zero-sum repeated game with discounted losses is the fixed point of a set-valued dynamic programming operator defined on the space of lower Pareto frontiers of closed convex sets with an appropriately defined metric. We then show that the optimal policies that guarantee different points in this set are of the following form. \mathcal{V}^∞ can be parametrized so that each point corresponds to a parameter value, which can be thought of as an ‘‘information state’’. Each state is associated with an immediate optimal randomized action and a transition rule that depends on the observed action of the adversary. In order to attain a point in \mathcal{V}^∞ , Alice starts with the corresponding state, plays the associated randomized action, transitions into another state depending on Bob’s observed action as dictated by the rule, plays the randomized action associated with the new state and so on. In particular, the strategy does not depend on Alice’s past actions and it depends on Bob’s past actions only through this state that the minimizing player keeps track of.
2. For the case where Alice has only 2 actions, we give a value-iteration based procedure to approximate \mathcal{V}^∞ and to compute an approximately optimal policy that only uses a coarse finite quantization of the parameter space. This strategy can be simply implemented by a randomized finite-state automaton. Any desired diminishing approximation error can be attained by choosing the appropriate quantization granularity and number of iterations. Our procedure in principle can be extended to an arbitrary number of actions. We finally illustrate our theory and the approximation procedure on a simple model of prediction with expert advice with 2 experts.

Overview of the approach. Informally, our approach can be described as follows. Let \mathbb{G}^T denote the T -stage repeated game and let \mathbb{G}^∞ denote the infinitely repeated game. Let $\mathcal{V}^0 = \{(0, 0)\}$. We can show that one can obtain the set \mathcal{V}^{T+1} from the set \mathcal{V}^T , by decomposing Alice’s strategy in \mathbb{G}^{T+1} into a strategy for the 1st stage, and a continuation strategy for the remainder of the game from stage 2 onwards, as a function of the action chosen by both the players in the 1st stage. The induction results from the fact that the minimal guarantees that she can guarantee from stage 2 onwards are exactly the set \mathcal{V}^T . Suppose that at the start of \mathbb{G}^{T+1} , Alice fixes the following plan for the entire game: she will play a mixed strategy $\alpha \in \Delta(A)$ in stage 1. Then depending on her realized action a and Bob’s action b , from stage 2 onwards she will play a continuation strategy

that achieves the guarantee $\mathbf{R}(a, b) \in \mathcal{V}^T$ (she will choose one such point $\mathbf{R}(a, b)$ for every $a \in A$ and $b \in B$). Note that it is strictly sub-optimal for Alice to choose any points outside \mathcal{V}^T from stage 2 onwards. Now this plan for the entire game \mathbb{G}^{T+1} gives Alice the following expected simultaneous guarantees on the two components:

$$\left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a [r_1(a, b) + \beta R_1(a, b)] \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a [r_2(a, b) + \beta R_2(a, b)] \right\} \right).$$

By varying the choice of α and the map $\mathbf{R}(a, b)$ we can obtain the set of all the simultaneous guarantees that Alice can achieve in the $(T + 1)$ -stage game. The Lower Pareto frontier of this set is exactly \mathcal{V}^{T+1} . Thus there is an operator Φ , such that

$$\mathcal{V}^{T+1} = \Phi(\mathcal{V}^T)$$

for any $T \geq 0$, where \mathcal{V}^0 is defined to be the singleton set $\{(0, 0)\}$. We show that this operator is a contraction in the space of Lower Pareto frontiers of closed convex sets, with an appropriately defined metric. This space is shown to be complete, and thus the sequence \mathcal{V}^T converges in the metric d to a set \mathcal{V}^* , which is the unique fixed point of this operator Φ . As one would guess, this \mathcal{V}^* is indeed the set \mathcal{V}^∞ of minimal simultaneous guarantees that Alice can achieve in the infinitely repeated game \mathbb{G}^∞ .

Note that an important step in our approach is a reduction of the problem to a vector-valued repeated game. The study of vector-valued repeated games was pioneered by Blackwell [Blackwell, 1956a]. He gave sufficient conditions for a set to be *approachable* by a player, which means that there exists a strategy for a player that ensures that the average loss approaches this set regardless of the adversary's actions. Moreover, he explicitly defined an adaptive randomized strategy that ensures this. Later he also showed that this theory can be used to obtain no-regret strategies as formulated by Hannan [Hannan, 1957], using the transformation of the repeated game into a vector-valued game that we described earlier [Blackwell, 1956b]. This theory was subsequently extended in various ways [Vieille, 1992, Lehrer, 2003], and stronger connections with regret minimization and other learning problems like calibration were shown [Abernethy et al., 2011, Perchet, 2014]. But as far as we know, there has been no prior work on the discounted loss criterion. In this work, as a by-product of our analysis, we successfully bridge this important gap in the theory of vector-valued repeated games. As a result, this theory bridges significant gaps in other decision-making problems where Blackwell's approachability theory has found applications.⁶

Summary of our contribution to online decision making with adversarial losses:

We propose a systematic set-valued dynamic programming approach for designing regret-optimal strategies in repeated games with discounted losses.

⁶A notable example is the analysis of zero-sum repeated games with incomplete information [Zamir, 1992, Aumann and Maschler, 1995, Sorin, 2002]

GAME THEORY AND STATISTICAL LEARNING FOR PRIVACY

This chapter covers the work in papers [7], [25, 27, 29, 31]. This work was done in collaboration with Oana Goga and Krishna Gummadi (Max-Planck Institute for Software Systems), Jens Grossklags (PennState University), Stratis Ioannidis (Northeastern University, formerly Technicolor and Yahoo! labs), Robin Sommer (ICSI Berkeley), Renata Teixeira (Inria, formerly CNRS); and with the following students: Athanasios Andreou (PhD student at EURECOM under my supervision), Michela Chessa (postdoc at EURECOM under my supervision). The ideas developed here also led to the PhD of Amine Lahouel started in Sept. 2016 under my supervision and in collaboration with Cédric Hébert (SAP labs).

This chapter summarizes my work relating to game theory and statistical learning for online privacy. Throughout this work, my general focus was on investigating algorithms to learn from personal data and how they affect privacy. I first focus on the problem (from the service provider's side) of learning from personal data directly revealed by privacy-conscious users. Since such users tend to strategically obfuscate the data they reveal, this naturally leads to game-theoretic analysis. In a second time, I focus on learning from data that is publicly available, as a tool to evaluate the privacy risk of users given the data that they already revealed.

3.1 Learning from personal data provided by privacy-conscious users

Online services collect personal data of users and exploit it using learning algorithms for various purposes such as to infer what users like in order to provide personalized services (recommendation, targeted advertisement, etc.). This raises the question of what is the best algorithm to learn from personal data.

3.1.1 Context

Machine learning is an active field of research and many efficient algorithms have been developed in the last decades for applications such as computer vision or medicine, in particular for regression and classification problems [Hastie et al., 2009]. In essence, such algorithms automatically learn from data (called training data) and generalize to new examples (called testing data) in order to make predictions. Recommendation algorithms based on personal data were also well studied in the last decade [Ricci et al., 2011], using different techniques but with the same philosophy as for other learning problems that consists in assuming that data is generated independently from the learning algorithm.

Personal data of human being is different from other types of data such as nature's pictures because its revelation raises privacy issues. This led to a large interest in the question of how to protect the privacy of users providing data while being able to learn from the data. A solution was proposed in the computer science community, called 'differential privacy' [Dwork, 2006, Dwork and Roth, 2014, Kifer et al., 2012], that consists in adding noise to the output of an algorithm so that it is not possible to determine whether the data of a given user participated in the computation. Since its introduction in 2006, this solution received almost exclusive attention from the privacy and computer science community. It usually requires, however, that users can give their raw data to a trusted third party who computes the algorithm's output and adds noise to it. If this is not possible, then users can directly add noise to their data before revealing it. This solution, although originating from the older idea of randomized response [Warner, 1965], recently gained interest in the formal framework of 'local differential privacy' [Duchi et al., 2013, Kairouz et al., 2016]. Broadly, the literature on differential privacy and local differential privacy seeks obfuscation mechanisms to guarantee a given level of privacy, characterization of the accuracy of statistical estimation under privacy constraints, and mechanisms to optimize this privacy-utility tradeoff where utility is defined based on the accuracy of a model learned from the data (see e.g., [Dwork and Roth, 2014, Chapter 11] and [Kairouz et al., 2016] and references therein). Perturbing a dataset before submitting it as input to a data mining algorithm also has a long history in privacy-preserving data-mining (see, e.g., [Vaidya et al., 2006, Domingo-Ferrer, 2008]). In the early days of data sharing, researchers were interested in how companies or governments can perturb their data before sharing it with third parties to hide sensitive information independently of the algorithm [Traub et al., 1984, Duncan and Mukherjee, 2000]. For more accurate results, perturbations tailored to specific data mining tasks have then been proposed in the context of reconstructing the original distribution of the underlying data or building decision trees [Agrawal and Srikant, 2000], clustering [Oliveira and Zaiane, 2003], and association rule mining [Atallah et al., 1999].

The aforementioned papers consider only the relation between privacy and utility (under various definitions) for given obfuscation mechanisms. In reality, in many applications (including online services), personal data is revealed directly by the individuals who can then *choose* the obfuscation level (or how much data to reveal, or whether to reveal data at all) according to their privacy sensitivity. That introduces questions related to incentives to reveal data, which also generated a significant literature. Researches in behavioral economics [Huberman et al., 2005, Acquisti and Grossklags, 2012] attempted to quantify experimentally the private cost incurred by an agent when releasing personal data. The game-theoretic analysis of incentives in personal data was pioneered by [Kleinberg et al., 2001], who proposed fair compensation mechanisms for personal data based on cooperative game theory in a simplistic model. Recently, a significant thread of research started on designing mechanisms to buy data from potentially untruthful agents for various objectives [Ghosh and Roth, 2011, Dandekar et al., 2012, Ligett and Roth, 2012, Roth and Schoenebeck, 2012, Dwork and Roth, 2014, Cummings et al., 2015]. In all those works, the loss of privacy by releasing data is quantified using variants of differential privacy. [Riederer et al., 2011] also propose a mechanism called transactional privacy where users can sell access to their data through an unlimited supply auction. Finally, other models also considered the case where agents may choose their effort in providing data [Cai et al., 2015, Luo et al., 2015] or have heterogeneous costs of providing data [Abernethy et al., 2015]. In all these works, users are assumed to maximize the payment received (minus cost of effort). The data elicitation literature also considered related problems where one tries to incentivize an expert to truthfully reveal his prediction of an event, typically using scoring rules [Gneiting and Raftery, 2007, Chambers and Lambert, 2014, Frongillo et al., 2015] (see also the literature on incentives in crowdsourcing [Dasgupta and Ghosh, 2013]).

All the literature mentioned above about incentives and privacy considers agents who are maximizing only the payment received but are insensitive to the quality of the learning result; and looks at how to optimize the payments while *the learning algorithm is fixed*. In reality, individuals do choose to reveal data without being paid in many online services, simply because they have an interest in the result of the learning algorithm and therefore in helping to make it accurate (this is also the case in many data analytics projects of societal importance). Personal data is therefore a special kind of data provided by agents whose objectives depend on the learning algorithm itself. That brings a completely orthogonal question of how to *change the learning algorithm* to increase the accuracy without involving payments. Barely any study considered that question to date. The (relatively small) literature on strategy-proof learning considered cases where each agent is interested in maximizing the learning accuracy around his own data points [Perote and Perote-Pena, 2004, Dekel et al., 2008, Nix and Kantarcioglu, 2012, Meir et al., 2012], but it does not include privacy considerations or effort that users provide to give data. [Chorppath and Alpcan, 2013, Caragiannis et al., 2016] consider effort and privacy questions but in simple cases without learning task (averaging).

3.1.2 Our contributions

In this context, our first and main contribution was the proposition in [31] of the first model of privacy-conscious users that takes into account the interest of users in the learning outcome by modeling the precision of the information learned as a public good with externalities (i.e., the information revealed by an agent benefits other agents). Here, we summarize the model as well as our main results in the original setting [31] and in a simplified setting [27, 29].

Model

Our model is based on the following basic setting. A set $N = \{1, \dots, n\}$ of users each have a *public* variable $\mathbf{x}_i \in \mathbb{R}^d$, $d \geq 1$ (a set of features that are publicly available such as age), and a *private* variable $y_i \in \mathbb{R}$ (known only by user i , e.g., concentration of a substance in blood). We assume that the public and private variables are linked by a linear model:

$$y_i = \boldsymbol{\beta}^T \mathbf{x}_i + \epsilon_i, \quad (i \in N),$$

where $\boldsymbol{\beta} \in \mathbb{R}^d$ is the parameter of the linear model and $\epsilon_i \in \mathbb{R}$ is an i.i.d. random variable of zero mean and of finite variance σ^2 . Users do not reveal their private variable y_i , but rather a noisy version of it: $\tilde{y}_i = y_i + \tilde{\epsilon}_i$, where $\tilde{\epsilon}_i$ is an i.i.d. random variable of zero mean, independent of ϵ_i . We assume that, by his choice of the variance of $\tilde{\epsilon}_i$, user i chooses the precision (i.e., inverse variance) λ_i of \tilde{y}_i in an interval $[0, 1/\sigma^2]$. An analyst collects the variables x_i and \tilde{y}_i and uses them to obtain an estimate of $\boldsymbol{\beta}$ through generalized least-square regression (GLS):

$$\hat{\boldsymbol{\beta}}_{\text{GLS}} = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^d} \left(\sum_{i \in N} \lambda_i (\tilde{y}_i - \boldsymbol{\beta}^T \mathbf{x}_i)^2 \right) = (X^T \Lambda X)^{-1} X^T \Lambda \tilde{\mathbf{y}}, \quad (3.1)$$

where $\tilde{\mathbf{y}} = [\tilde{y}_i]_{i \in N}$ is the n -dimensional vector of perturbed variables, $X = [\mathbf{x}_i^T]_{i \in N} \in \mathbb{R}^{n \times d}$ the $n \times d$ matrix whose rows comprise the transposed feature vectors, $\boldsymbol{\lambda} = [\lambda_i]_{i \in N}$, and $\Lambda = \text{diag}(\boldsymbol{\lambda})$. GLS is known to give the smallest covariance, in the positive semi-definite sense, amongst linear unbiased estimators [Aitken, 1935]. Then, the precision of the linear model estimated is characterized by the covariance $V(\boldsymbol{\lambda}) = (X^T \Lambda X)^{-1}$ of $\hat{\boldsymbol{\beta}}_{\text{GLS}}$. Finally, we assume that each user minimizes a cost

$$J_i(\lambda_i, \lambda_{-i}) = F(V(\boldsymbol{\lambda})) + c_i(\lambda_i),$$

where F is a convex and increasing function (e.g., the trace function) that describes the *estimation cost*, i.e., the cost incurred by a user due to imprecision of the analyst's estimate (the public good part); and c_i is a convex and increasing function that describe the individual's *privacy cost*. Specifically, we will use the following assumptions:

Assumption 1. The privacy costs $c_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, $i \in N$, are twice continuously differentiable, non-negative, non-decreasing and strictly convex.

Assumption 2. The scalarization $F : S_{++}^d \rightarrow \mathbb{R}_+$ is twice continuously differentiable, non-negative, non-constant, non-decreasing in the positive semidefinite order, and convex.

We denote by Γ the complete information game defined above.

Results

For the model above, we first analyze the Nash equilibria. We show that the game is potential and, using the potential game structure of Γ , we derive the following result.

Theorem 3.1.1. *There exists a unique non-trivial equilibrium of the game Γ .*

Here, a non-trivial equilibrium is an equilibrium where each agent has a finite cost. Having a unique non-trivial equilibrium is an interesting property both for theoretical derivations and for applications. In [31], we also show the same result for variants of the game Γ where the analyst does not use GLS estimator but rather a generic linear estimator.

Then, we move to studying the efficiency of the Nash equilibrium, captured by the price of stability defined as

$$\text{PoS} = \frac{\min_{\lambda \in \mathcal{NE}} (nF(V(\lambda)) + \sum_{i \in N} c_i(\lambda_i))}{\min_{\lambda \in [0, 1/\sigma^2]^n} (nF(V(\lambda)) + \sum_{i \in N} c_i(\lambda_i))},$$

where $\mathcal{NE} \subset [0, 1/\sigma^2]^n$ is the set of Nash equilibria. Note that the price of stability is equal here to the price of anarchy that we would obtain after eliminating the trivial equilibria. Then we can show the following bounds.

Theorem 3.1.2. *Under Assumptions 1 and 2, $\text{PoS} \leq n$.*

The bound of Theorem 3.1.2 works for every privacy cost satisfying Assumption 1, but it is quite crude. By restricting the privacy costs, we can obtain better results. We begin by providing a bound on the price of stability when privacy costs are monomial functions:

Theorem 3.1.3. *Assume that the cost functions are given by $c_i(\lambda) = c_i \lambda^k$, where $c_i > 0$ and $k \geq 1$. If the estimation cost is given by $F_1(V) = \text{trace}(V)$, then $\text{PoS} \leq n^{\frac{1}{k+1}}$. If the estimation cost is given by $F_2(V) = \|V\|_F^2$, then $\text{PoS} \leq n^{\frac{2}{k+2}}$.*

The proof of Theorem 3.1.3 relies on characterizing explicitly the socially optimal profile under relaxed constraints, and showing it equals the Nash equilibrium λ^* multiplied by a scalar. Moreover, the theorem states that, among monomial privacy costs, the largest PoS is $n^{\frac{1}{2}}$ for $F = F_1$, and $n^{\frac{2}{3}}$ for $F = F_2$. Both are attained at linear privacy costs; in fact, the above ‘‘worst-case’’ bounds can be generalized to a class of functions beyond monomials.

Theorem 3.1.4. Assume that for every $i \in N$ the privacy cost functions $c_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ satisfy Assumption 1. If the estimation cost is the extended-value extension of $F_1(V) = \text{trace}(V)$, and the derivatives c'_i satisfy

$$nc'_i(\lambda) \leq c'_i(n^{\frac{1}{2}}\lambda) \quad (3.2)$$

then $\text{PoS} \leq n^{\frac{1}{2}}$. Similarly, if the estimation cost is the extended-value extension of $F_2(V) = \|V\|_F^2$, and the derivatives c'_i satisfy

$$nc'_i(\lambda) \leq c'_i(n^{\frac{1}{3}}\lambda) \quad (3.3)$$

then $\text{PoS} \leq n^{\frac{2}{3}}$.

Theorem 3.1.4 applies to privacy cost functions that have the “strong” convexity properties (3.2) and (3.3). Roughly speaking, such functions grow no slower than cubic and fourth-power monomials, respectively. In contrast to Theorem 3.1.3, in the case of Theorem 3.1.4, we cannot characterize the social optimum precisely; as a result, the proof relies on Brouwer’s fixed point theorem to relate λ^{opt} to the non-trivial Nash equilibrium λ^* .

All the efficiency bounds above apply to the case where the analyst uses the GLS estimator. However, for general privacy costs and even for $F = \text{trace}$, GLS is not guaranteed to give a minimum estimation cost at Nash equilibrium. To continue our analysis of the equilibrium estimation cost and how to improve it, we restrict in [27, 29] to a simpler setting where the analyst simply tries to compute the average of the private variables. That is, $d = 0$ (there is no public variable), and the analyst computes the generalized least square estimator which is simply the weighted average

$$\hat{y}_M(\lambda) = \frac{\sum_{i \in N} \lambda_i \tilde{y}_i}{\sum_{i \in N} \lambda_i}. \quad (3.4)$$

The covariance is then simply the variance

$$\sigma_M^2(\lambda) = \frac{1}{\sum_{i \in N} \lambda_i} \in [\sigma^2/n, +\infty]. \quad (3.5)$$

We also first assume that the population is homogeneous, that is that each agent has the same privacy cost. In that case, we establish two types of results:

Monotonicity and convergence: We show that, as the number n of agents increases, the equilibrium precision of each agent decreases and converges to zero (i.e., each agent gives less and less precise data hence having more and more privacy); but the equilibrium estimation variance σ_M^2 decreases and converges to zero (i.e., the aggregate learning precision increases).

Improvement of the equilibrium variance: We show that it is possible to improve the equilibrium variance simply by setting a minimum precision η . That is, each agent can either give no data at all, or give data with precision in $[\eta, 1/\sigma^2]$. We show that this strictly decreases the equilibrium variance for a well chosen value of η .

Finally, we show that those results can be extended to the case of an heterogeneous population. This provides a simple way, by restricting the action space of the players, to improve the quality of the model learned at equilibrium.

Summary of our contribution to learning from personal data disclosed by privacy-conscious users:

We propose the first model that takes into account the public good nature of the learning outcome, in the context of linear regression. We provide efficiency results for the general model and, for a simplified version, we provide a simple way to increase the learning precision.

3.2 Learning and estimation of the privacy risks from public data

Individuals publicly share large amounts of data about themselves on *social computing systems* such as Facebook, Twitter, LinkedIn, Reddit, IMDB, or Yelp. Although they receive great utility from those systems, users are also concerned that such data sharing negatively affects their *privacy*; but estimating the privacy risk from public data sharing is a hard challenge as it depends on how much one can learn from the available data.

3.2.1 Context

In the past decade, a large body of research has provided and evaluated methods that show that (hidden) sensitive information about a user such as ethnicity or political affiliation can be inferred by mining publicly available data within a single social computing system [Backstrom et al., 2010, Zheleva and Getoor, 2009, Mislove et al., 2010, Chang et al., 2010]. This type of information disclosure is called *attribute disclosure*: it consists in inferring the value of an attribute (e.g., ethnicity) that was hidden (i.e., not directly shared by the user). All these studies use either *friendship* or *user behavior* data (or both) and exploit the homophily property to make the inference (i.e., the fact that it is possible based on friendship or user behavior data to construct groups of users sharing similar attributes).

In parallel, many works appeared in recent years on *matching* identities across multiple social computing systems [Motoyama and Varghese, 2009, Perito et al., 2011, Malhotra et al., 2012, Paridhi Jain and Joshi, 2013, Acquisti et al., 2011, You et al., 2011, Vosecky et al., 2009, Raad et al., 2010, Northern and Nelson, 2011, Peled et al., 2013, Liu et al., 2013, Zafarani and Liu, 2013, Zafarani and Liu, 2009, Labitzke et al., 2011], that is on building algorithms to find, for a given identity in a social computing system, the identity in a second social computing system that belongs to the same individual (termed the matching identity). This type of information disclosure is called *identity disclosure*. The proposed matching algorithms typically use publicly available attributes (such as name and bio) and leverage the fact a individuals share attributes across social computing systems that might be unique enough to identify them.

However, in this space, we identify two missing stones:

- Existing matching schemes were only tested on very small datasets and we have no estimation of their reliability at a reasonable scale corresponding to real-world social computing systems. Here, reliability refers to the extent to which different profiles belonging to the same user can be matched across social networks, while avoiding mistakenly matching profiles belonging to different users. Matching schemes need to be highly reliable because incorrectly matched profiles communicate an inaccurate portrait of a user and could have seriously negative consequences for the user in many application scenarios. For example, Spokeo¹ has been recently sued over providing inaccurate information about a person which caused “actual harm” to the person employment prospects². This raises many questions such as *how to evaluate the reliability of a matching scheme at scale? how to build more reliable matching schemes? what is the reliability that can be achieved in real-world social computing systems?*
- Surprisingly, few studies have systematically considered the fact that, in addition to identity disclosure risks, considering multiple social computing systems also introduces significant new attribute disclosure risks due to the possibility of inferring a hidden attribute in a profile by looking at

¹<http://www.spokeo.com/>

²<http://www.ftc.gov/sites/default/files/documents/cases/2012/06/120612spokeocmpt.pdf>

another social computing system (either through homophily or by finding the matching identity). Such attribute disclosure is powerful because individuals reveal different pieces of information on different social computing systems [Chen et al., 2012] (e.g., personal life on Facebook, profession on LinkedIn, interests on Twitter).

Even more importantly, to the best of our knowledge, no study has *jointly* analyzed identity and attribute disclosure risks. Doing this joint analysis is particularly important because the research community recently gained interest in building defenses against privacy attacks (such as privacy advisors). Defenses were proposed separately in the context of attribute disclosure [Biega et al., 2016] (warning users when their behavior put them at risk of attribute disclosure, e.g., “liking this will reveal that about you”), and in the context of identity disclosure (advising users to blend into the crowd, that is to share information at a granularity that makes them less uniquely identifiable [Backes et al., 2016]). However, it is not clear that, in the context of multiple social computing systems where both risks are present, one type of defense also helps against the other type of risk. Intuitively indeed, while blending into the crowd might help against identity disclosure, it might also offer more opportunities to learn attributes and hence increase the attribute disclosure risk. This raises the key questions: *what is the link between the two disclosure risks? does a lower identity disclosure risk always correspond to a lower attribute disclosure risk? do defenses against identity disclosure risk also reduce the attribute disclosure risk?*

3.2.2 Our contributions

Our work tackles both sets of questions mentioned above. We start by describing our results on matching; all details of which are in [25]. Our first contribution lies in the definition of a framework consisting in a set of four properties for profile attributes—Availability, Consistency, non-Impersonability, and Discriminability (ACID)—that determine the reliability of a matching scheme and help us understand what the reliability of a matching scheme depends on.

To present the ACID framework, we first introduce a few notation. We consider that two profiles in two social computing systems match if they belong to/are managed by the same user. The profile matching problem is then: given a profile a^1 in one social computing system SN_1 , find all its *matching profiles* in another large social network SN_2 , if at least one exists. We will denote by a^2 generic profiles in SN_2 and by \hat{a}^2 matching profiles of a^1 . For conciseness, we will also write a^2 -match- a^1 if a^2 is a matching profile of a^1 and a^2 -non-match- a^1 otherwise. We investigate the extent to which we can match profiles by exploiting the *attributes* users publicly provide in their profiles such as their *real names*, *screen names*, *location*, *profile photos*, and *friends*. For profile a^1 (resp. a^2), we denote by v^1 (resp. v^2) the value of a considered attribute. From attribute values, we define a *feature* as the similarity between the values of profiles in SN_1 and SN_2 : $s(v^1, v^2)$.

For a given attribute, we propose the following four properties to help capture its quality to match profiles reliably:

Availability: At first, to enable finding the matching profile, an attribute should have its value available in both social computing systems. To formalize this notion, we model the attribute values of a^1 and each $a^2 \in SN_2$ as random variables and we define the availability of an attribute as:

$$A = Pr(v^1 \text{ and } v^2 \text{ available} | a^2\text{-match-}a^1).$$

Consistency: It is crucial that the selected attribute is consistent across matching profiles, i.e., users provide the same or similar attribute values across the different profiles they manage. For-

mally, we define the consistency of an attribute as:

$$C = Pr\left(s(v^1, v^2) > th \mid a^2\text{-match-}a^1, v^1 \text{ and } v^2 \text{ available}\right),$$

where th is a threshold parameter.

non-Impersonability: If an attribute can be easily impersonated, i.e., faked, then attackers can compromise the reliability of the matching by creating fake profiles that appear to be matching with the victim's profiles on other sites. Some public attributes like "name" and "profile photo" are easier to copy than others such as "friends". To formalize this notion, we introduce the notation a^2 -impersonate- a^1 to denote that profile a^2 has been created by an attacker impersonating profile a^1 . We denote the probability that there exists at least one profile a^2 impersonating a^1 by $p_I = Pr(a^1 \text{ is impersonated})$ and the probability that there is no profile impersonating a^1 by $p_{nI} = 1 - p_I$. The difficulty to manipulate an attribute is characterized by its non-Impersonability defined as:

$$nI = Pr\left(\max_{a^2: a^2\text{-impersonate-}a^1} s(v^1, v^2) < th\right).$$

Discriminability: Even without impersonations, in order to enable finding the matching profile, an attribute needs to uniquely identify a profile in SN_2 . A highly discriminating attribute would have a unique and different value for each profile, while a less discriminating attribute would have similar values for many profiles. For example, "name" is likely to be more discriminating than "gender". Formally, we define the discriminability of an attribute as:

$$D = Pr\left(\max_{a^2: a^2\text{-non-match-}a^1} s(v^1, v^2) < th \mid a^1 \text{ not impersonated}\right).$$

In practice, it is impossible to estimate D unless we are able to identify impersonating profiles. Instead, we estimate:

$$\tilde{D} = Pr\left(\max_{a^2: a^2\text{-non-match-}a^1} s(v^1, v^2) < th\right).$$

\tilde{D} represents the "effective discriminability" taking into account possible impersonations. Since impersonators create non-matching profiles as similar as possible to the original profile, it is reasonable to assume that $\tilde{D} \leq D$. Moreover, by application of Bayes formula, we can show that $D \leq \tilde{D}/p_{nI}$ so that, if p_I is not too large, \tilde{D} gives a good estimate of D . If we assume that the impersonating profiles are independent from the other non-matching profiles, we can also prove that $\tilde{D} = D \cdot (p_{nI} + nI \cdot p_I)$. This clearly shows that \tilde{D} is close to D if either the attribute is hard to impersonate (nI close to one) or the proportion of impersonator is small (p_I small).

The ACID properties are clear and intuitive properties that help understand the potential of an attribute to perform reliable matching, as the following theorem formalizes. The performance of the matching scheme is measured by the standard *precision* and *recall* quantities where precision is defined as the fraction of all pairs returned by the matching scheme which are true matches and recall as the fraction of matching profiles that are identified, that is:

$$precision = Pr\left(a^2\text{-match-}a^1 \mid s(v^1, v^2) > th\right) \quad \text{and} \quad recall = Pr\left(s(v^1, v^2) > th \mid a^2\text{-match-}a^1\right).$$

Theorem 3.2.1. Consider a classifier based on a given attribute that classifies as matching profiles if $s(v^1, v^2) > th$. The performance of the classifier is characterized by the following results.

(i) We have

$$recall = C \cdot A.$$

(ii) Assume that, for each profile $a^1 \in SN_1$, there is at most one matching profile in SN_2 . Then,

$$\text{precision} \leq \frac{\text{recall}}{\text{recall} + 1 - \bar{D}}.$$

(iii) Assume that $p_1 > 0$. Then, $\text{precision} = \text{recall} = 1$ iff $A = C = nI = D = 1$.

In Theorem 3.2.1, the threshold parameter th must be the same as the one in the definitions of C , nI and D . Theorem 3.2.1-(i) shows that the classifier’s recall is simply the product of consistency and availability. Theorem 3.2.1-(ii) gives a simple upper bound of the precision as a function of the effective discriminability (which itself is a function of the discriminability and of the impersonability, see above). This upper bound gives a good order of magnitude for the precision; moreover, for high precision (which is what we aim), given the small number of false positives, the true precision should be close to the bound. Finally, Theorem 3.2.1-(iii) confirms that a high value of all four ACID properties is *necessary* and *sufficient* to obtain high precision *and* recall.

Properties A , C and nI are independent of the network scale, but the discriminability very largely depends on the network scale since having more non-matching pairs decreases the probability that none of them has a high similarity score. This implies that we must estimate the precision and the recall of a matching scheme using datasets that accurately capture the ACID properties of profile attributes of the entire social network; otherwise the precision and the recall will be incorrect. This is typically what is done wrong by previous works which evaluate their schemes on random datasets, that is datasets that contain a random subset of profiles from SN_1 and their matching profiles from SN_2 . The discriminability for such datasets is very high and the performance appears good as shown in Figure 3.1a.

We proposed a way of emulating a very large-scale dataset, which preserves the discriminability. As shown in Figure 3.1b, the actual performance is then much lower than estimated on a random dataset. Further, we proposed optimization methods for the training that increase the performance (Figure 3.1c), as well as a new scheme that performs even better based on the assumption that each profile has at most one matching account (Figure 3.1c). At best, we obtain 30% recall for a 95% precision (whereas the erroneous evaluation of previous work predicted 90% recall). Finally, we perform tests using human workers and observe that even they only achieve a 40% recall for a 95% precision.

In our follow-up work [7] (not published yet), we tackled the question of the relationship between attribute and identity disclosure. First, while working on quantifying identity disclosure risks we realized that its definition is not completely straightforward; hence we proposed a new definition. Then, we showed that there is a trade-off between attribute and identity disclosure in some regime, that is, profiles with a lower identity disclosure risk (because they blend into the crowd) have a higher attribute disclosure risk (because there are more similar profiles to learn from).

Summary of our contribution to learning for privacy risk estimation:

We propose a framework (ACID) to understand and evaluate matching schemes at very large scale. We propose optimized matching schemes for very large scale and study the limits in current social computing systems. We also analyze the relationship between identity and attribute disclosure risks.

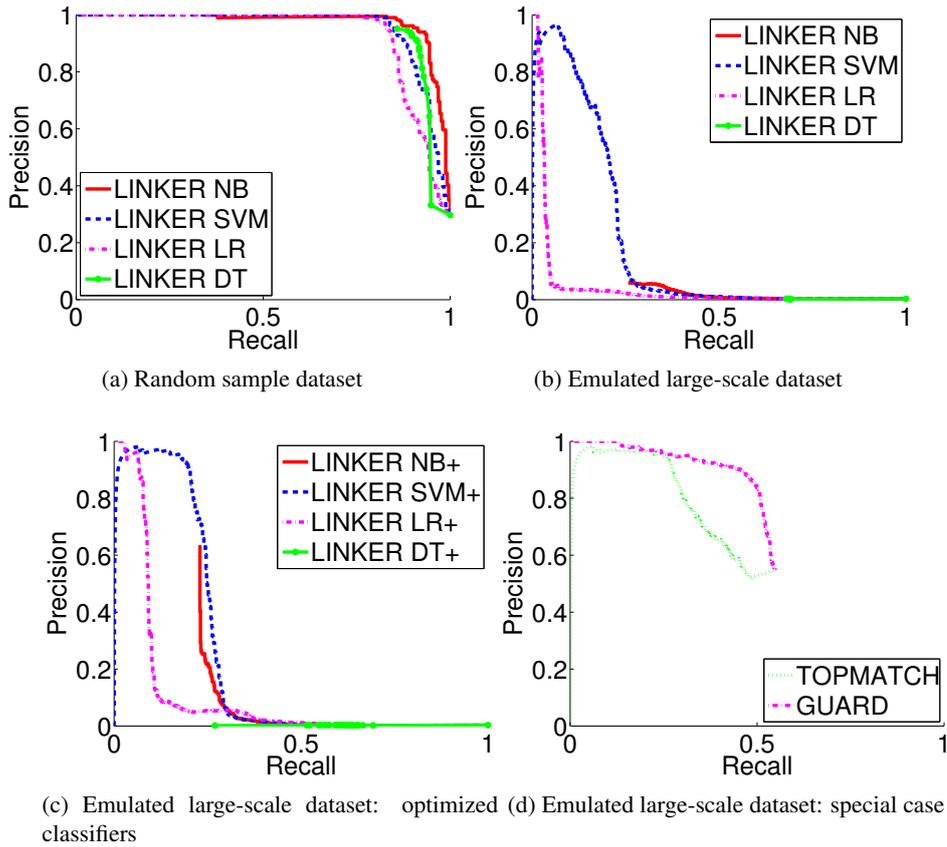


Figure 3.1: Precision and recall tradeoff for matching Twitter to Facebook profiles using different classifiers when evaluated over a random dataset and an emulated large-scale dataset. "Linker" refers to the simple binary classifier that classifies a pair as a match if its similarity score is larger than a threshold, using a standard classification method: Naive Bayes (NB), Support Vector Machine (SVM), Logistic Regression (LR) or Decision Tree (DT). The symbol "+" identifies the optimized training. "Topmatch" and "Guard" correspond to classifiers that optimize the performance by using the assumption that a profile in SN_1 can have only one matching in SN_2 —see details in [25].

GAME THEORY AND STATISTICAL LEARNING FOR NETWORK SYSTEMS

This chapter covers the work in papers [11], [13–15], [23, 24, 26, 34, 36], [40, 41]. This work was done in collaboration with Saurabh Amin (MIT, formerly UC Berkeley), Ernst Biersack (EURECOM), John Musacchio (UC Santa Cruz), Giovanni Neglia (Inria), Shankar Sastry (UC Berkeley), Galina Schwartz (UC Berkeley); and with the following students: Alberto Benegiamo (research assistant under co-supervision of Giovanni Neglia and myself), Hadrien Hours (PhD student at EURECOM under co-supervision of Ernst Biersack and myself), Xiaohu Wu (PhD student at EURECOM under my supervision).

This final chapter summarizes my work related to game theory and statistical learning for network systems construed broadly. Throughout this work, my main focus was on evaluating and improving the infrastructure on which user-centric services rely. I first focus on the question of analyzing and predicting network performance, for which the statistical learning tools of causal analysis is the most natural. Then I study the problem of reducing network congestion by using incentives, in the contexts of communication networks and of energy networks. Analyzing the effect of incentives and deriving optimal incentives is naturally done using game theory to model the response of agents to incentives. I finally focus on allocation of resources in a cloud.

4.1 Practical causal analysis for network performance prediction

4.1.1 Context

Communication networks are complex systems whose operation relies on a large number of components that work together to provide services to end users. As the quality of these services depends on different parameters, understanding how each of them impacts the final performance of a service is a challenging but important problem. However, intervening on individual factors to evaluate the impact of the different parameters is often impractical due to the high cost of intervention in a network. Previously proposed approaches to study network performance fall short of allowing one to understand the role of the different parameters and to predict how a change in any of these parameters will impact performance. It is therefore desirable to adopt a formal approach that allows this.

The approach of causality pioneered by J. Pearl [Pearl, 2009] provides a powerful framework to investigate these questions, in particular using graphical causal models represented by Bayesian networks that give simple graphical criteria to infer the mathematical equations to predict the effect of interventions [Pearl, 2009, Spirtes et al., 2001]. Such an approach represents the set of parameters as a graph whose directed links mean a causal effect and it offers a concise visual presentation

of the different causal dependences, which greatly simplifies the understanding and manipulation of complex systems. It also allows performing prediction of the change of performance when intervening on one of the parameters.

In this work, we use a constraint-based method named PC algorithm [Spirtes and Glymour, 1991] to infer the causal graph, which mostly rely on testing all possible conditional independences. The choice of the independence test is then essential for inferring the correct graph. Many implementations of the PC algorithm use the Z-Fisher test, which assumes linearity and normality of the data. We observed, however, that data from communication networks are neither linear nor normal. To better take into account the nature of our data, we can use the Kernel Conditional Independence test (KCI) from [Zhang et al., 2012], which defines covariance operators in the Reproducing Kernel Hilbert Spaces (RKHS) corresponding to the parameters being tested and derives a statistic which distribution can be estimated under the hypothesis of independence. This independence test is very similar to the Hilbert Schmidt Independence Criterion (HSIC) [Gretton et al., 2007], the statistics of the two tests being exactly the same in the unconditional case. The use of kernel based independence test was proposed to handle the cases where no formal functional dependence or distribution can be assumed from the data. However, its implementation has two problems: First, the computation time to perform independence tests with this criterion is much longer than for the Z-Fisher test; in particular to test the conditional independences. Second, the implementation of the KCI test uses matrix operations (Cholesky factorization) that, although theoretically well justified, may fail due to numerical issues. In this case, the test will not return an answer.

4.1.2 Our contributions

In this context, our first contribution was to propose a bootstrap procedure to be able to apply the KCI test in practice. Indeed, both problems mentioned above can be avoided by using datasets of smaller size (or sub-datasets), but at the expense of a possible loss of accuracy. To solve these two problems, we modified the independence test by including a bootstrap method [Davison and Hinkley, 1997], which works as follows. For each independence test of $X \perp\!\!\!\perp Y \mid Z$ (X independent of Y conditionally on Z), we re-sample the dataset (with replacement) and create l new sub-datasets of N samples. The test is performed on each re-sampled sub-dataset (by comparing the p -value to an objective level of significance of the test, often denoted as α) and we accept the independence if a percentage of at least δ of sub-tests detects that the variables are independent. In [14], we show that this testing procedure KCI+bootstrap does not decrease the accuracy compared to the KCI test on the full dataset as long as N is large enough. On the other side, it offers a number of very valuable advantages. First, it reduces the computation time and offers an easy and efficient way to parallelize the computation on different machines. Second, it reduces the numerical issues due to Cholesky factorization and therefore allows to obtain a more accurate result in cases where the KCI test on the full dataset would simply fail. Also the bootstrap method offers a way to estimate the confidence in the result of the independence tests. We also show in [14] how to select the different parameters of the new KCI+bootstrap test.

Armed with the KCI+bootstrap test, we are then able to infer a causal graph from real-world datasets (about 20 features and 5000 data points) in a reasonable time. We then proposed a flexible method based on using causal rules and using copulae for estimation of the parameters distribution and dependence to do the prediction. Finally, we applied these two methods for the study of two network performance applications: TCP [14] and DNS [13].

In our study of TCP [14], we started with controlled experiments generated using the emulator

Mininet [Lantz et al., 2010]. This allowed us to validate the prediction done using our causal graph. Specifically, we generated transfers using the topology in Figure 4.1 and obtained the causal graph shown in Figure 4.2, where the parameters are explained in Table 4.1. We observe that this causal graph is in line with our domain knowledge and shows interesting new dependences that were not captured by previous models—see a complete discussion in [14]. Yet, to validate our inference method, we use predictions. Specifically, we perform a prediction of the effect of an intervention on the parameter retransmission score (representing some losses) on the throughput, using the causal graph just inferred together with our prediction method based on copulae. Then, we use the Mininet emulator to actually perform the intervention and measure the resulting throughput. The results are shown in Figure 4.3 and give us confidence in the accuracy of our methodology.

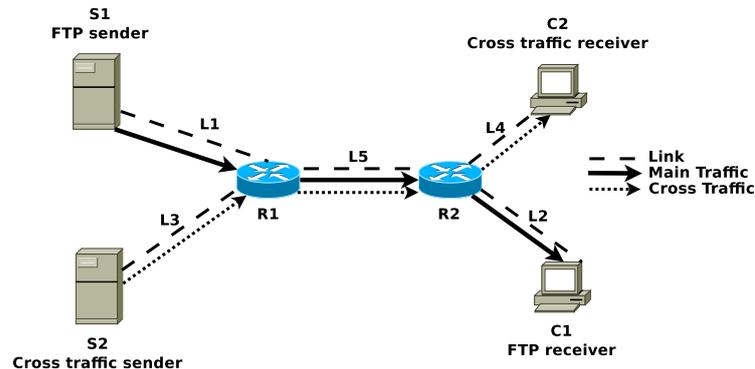


Figure 4.1: Emulated network using Mininet

Table 4.1: Summary of Mininet network emulation experiments dataset

Parameter	Definition	Min	Max	Avg	CoV
<i>bw</i>	minimum bandwidth (MBps)	1	25	7.1	0.91
<i>delay</i>	propagation delay (ms)	30	180	86	0.48
<i>queue1</i>	size of R1 buffer (pkts)	10	400	98	1.10
<i>queue2</i>	size of R2 buffer (pkts)	10	400	100	0.99
<i>nlac</i>	Narrow Link Available Capacity (kBps)	12	3070	630	5.00
<i>rwin</i>	Receiver window advertised by C1 (kB)	74	2155	288	0.65
<i>bufferingdelay</i>	part of the RTT due to queuing delay (ms)	1	6760	120	2.40
<i>rtt</i>	Round Trip Time (ms)	84	6910	310	0.99
<i>timeouts</i>	number of timeouts (units)	0	682	79	1.50
<i>retrscore</i>	fraction of retransmitted packets (no unit)	0	0.61	0.04	5.10
<i>p</i>	fraction of loss events (no unit)	0	0.64	0.04	8.40
<i>nbbytes</i>	number of bytes sent by the server (MB)	6	150	110	0.21
<i>tput</i>	throughput (kBps)	6	1100	280	0.81

To conclude our study on TCP, we collected a dataset of real TCP transfers from which we inferred the causal graph. Then, using this causal graph, we were able to perform predictions of the effect of interventions on parameters such as the RTT and loss rate. We also showed that these prediction based on causal methods differ from naive predictions based on model fitting where, to predict the effect of setting the RTT to a given value, one would simply condition on this value in the initial dataset. This shows that using causal method is a powerful tool for communication networks leading to more realistic predictions that could be useful for planning costly interventions.

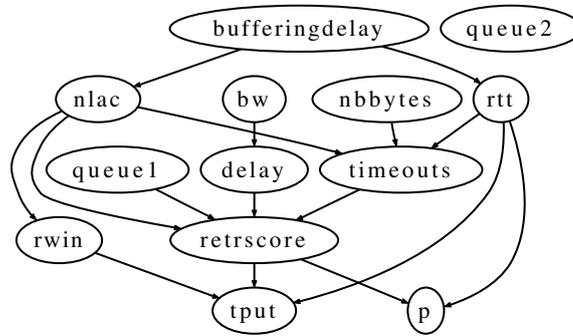


Figure 4.2: Causal model inferred by the PC algorithm, with KCI test, for the emulated network traffic

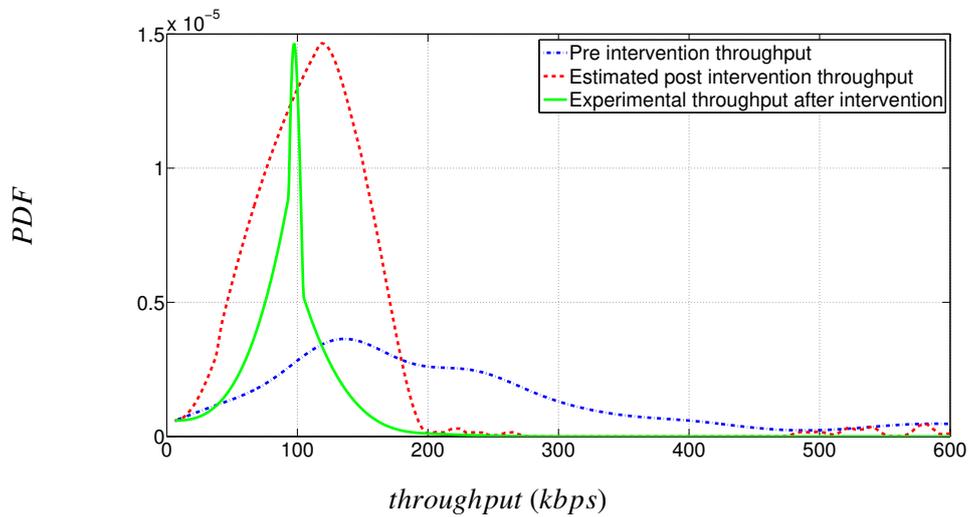


Figure 4.3: Effect of an intervention on the retransmission score, on the throughput distribution

Throughout our study of TCP, all transfers were generated using the version Reno of TCP (mainly for ease of comparison with existing analytical models such as the one of [Padhye et al., 1998]). If several TCP versions are being used, then one can account for the impact of the TCP version on the throughput simply by adding the TCP version as a (categorical) parameter in the graph. Also, in our experiments, we used transfer sizes uniformly distributed between 6 and 150 MBytes (for the Mininet experiments) and between 6 and 60 MBytes (for the real FTP transfers). While these distributions are not representative of the distribution of transfer sizes on the Internet, they were chosen to have a good statistical sample to evaluate the impact of the transfer size. Indeed, the transfer size is a parameter in our model (*nbbytes*), and we observe in the causal graphs that it does not directly affect the throughput but only directly affects the number of timeouts (the same holds for the causal graph of the real FTP transfers shown in [14]). Note also that the chosen distributions exclude small transfers (since our smallest transfer is 6 MB) hence our model will give accurate predictions only for transfers for which the transient beginning of TCP does not play a large role. In reality, file transfers on the Internet have a heavy-tailed distribution (typically well-modeled by a Pareto distribution of infinite variance [38], [21]) and this is well-known to bring statistical issues in particular for estimation of the mean and evaluation and simulation of flow-level bandwidth sharing (see [Rojas-Mora et al., 2011] which, interestingly, argues in favor of using bootstrap due to the heavy-tailed distribution of transfer sizes). It would be interesting to study how our results are changed when using a heavy-tailed distribution of transfer sizes. We suspect, however, that we would not observe much change because (i) the transfer size is part of our model so that our prediction of throughput is conditioned on the size of each transfer, and (ii) the observed impact of the transfer size is the same for Mininet experiments and for real FTP transfers which face real cross traffic that is likely heavy-tailed (probably because all important parameters impacted by the cross traffic such as retransmission score are already included in the model).

Finally, we used our causal methods to study another part of communication networks, DNS. For a user to access any resource on the Internet, it is necessary to first locate a server hosting the requested resource. The Domain Name System service (DNS) represents the first step in this process, translating a human readable name, the resource host name, into a machine readable name, an IP address. With the expansion of Content Distribution Networks (CDNs), the DNS service has seen its importance increase. In a CDN, objects are replicated on different servers geographically distributed in order to decrease the distance from the client to a server hosting the object that needs to be accessed. To attribute the demand of a user that wants to access an object hosted by a CDN, the CDN makes use of the DNS service to infer client localization. While most of the Internet Service Providers (ISPs) offer a DNS service to their customers, it is now common to see clients using a public DNS service instead. This choice may have an impact on the performance of clients retrieving objects from a given CDN.

In our work, we used our causal methods to study the impact of choosing one DNS service instead of another and we compare the performance of a large European ISP DNS service with the one of a public DNS service, Google DNS. Our approach exposed the structural dependencies of the different parameters impacted by the DNS service used and we showed how to model these dependencies with a Bayesian network. This model allowed us to explain and quantify the benefits obtained by clients using their ISP DNS service and to propose a solution to further improve their performance. Specifically, we predicted the performance a client using the local DNS service would have experienced if redirected to servers the Google DNS service would have redirected it to. In this way, we could estimate the impact of the local DNS service redirection strategy on the client performance and quantify the corresponding performance difference.

When comparing the performance of the local DNS and the Google DNS users, we observed that Google DNS users experience a throughput whose difference with the one of the local DNS service users cannot simply be explained by the redirection of Google DNS users to more distant Akamai servers. Based on the causal graph, we formulated the hypothesis that the configurations of the Akamai servers Google DNS users are redirected to allow them to eventually experience a performance similar to the one of the local DNS service users. This hypothesis is verified by our prediction consisting in giving to Akamai servers serving the local DNS users a minimum congestion window equivalent to the one of the Akamai servers serving Google DNS users. We estimated the gain in throughput corresponding to this intervention to be 32%. By comparison, the gain in terms of throughput corresponding to the better redirection of the local DNS users is estimated to 14%.

Overall, our work is amongst the first to successfully apply causal analysis in the engineering domain.

Summary of our contribution to causal analysis for communication networks performance:

We propose a variation of the KCI independence test using bootstrap that solves practical issues and allows inferring a causal graph from real-world datasets. We propose a prediction method based on copulae. We apply causal analysis to investigate how to improve the performance of TCP and DNS.

4.2 Incentives in networks with congestion

The consumer demand is steadily growing in many network areas, including communication networks and others, and numerous studies indicate that this growth will continue. The growing demand forces the providers to adopt congestion management schemes. However, it not possible to eliminate congestion in networks only using engineering techniques or smart queueing protocols. This is because, ultimately, what determines the network load at a given time is the demand, and demand is generated by humans with a natural tendency for non-uniform patterns in most networks: we observe peak times in road networks, energy networks, mobile networks, web servers, etc. Flattening the demand curve (either to reduce congestion in road or communication networks or to reduce production costs in energy networks) therefore requires acting on the demand.

4.2.1 Context

The literature has considered the use of incentives and pricing mechanisms to reduce the level of network congestion. Many pricing mechanisms have been proposed to manage quality of services (QoS) in networks, see e.g., surveys [Henderson et al., 2001, Tuffin, 2003, Sen et al., 2012, Sen et al., 2013]. For instance, in [Honig and Steiglitz, 1995], Honig and Steiglitz propose a usage-based pricing mechanism, and analyze it using a model with delay-sensitive users. Their results show how to find the price that maximizes the provider's revenue by solving a fixed-point equation. A similar model is used in [Basar and Srikant, 2002] where Bařar and Srikant analyze the many-users limit. They show that, as the number of users increases, the provider's revenue per unit of bandwidth increases and conclude that this gives providers an incentive to increase their network capacity. In a number of papers, e.g., [Mendelson and Whang, 1990, Odlyzko, 1999, Marbach, 2004], pricing mechanisms based on multiple classes of customers with different priorities are proposed and analyzed in terms of equilibrium achieved and optimal price per class. In [Shen and Bařar, 2007, Shen and Bařar, 2011], Shen and Bařar investigate the performance of non-linear pricing in a model similar to [Basar and Srikant, 2002] and find an improvement of up to 38% over

linear pricing in some cases. However, in all the aforementioned papers, the demand is assumed stationary and the price is fixed independently of the instantaneous network congestion or of the time of the day.

A few papers have proposed mechanisms with prices dependent on congestion levels. In [Paschalidis and Tsitsiklis, 2000], Paschalidis and Tsitsiklis propose a congestion-based pricing mechanism in the context of loss networks (i.e., phone). They provide a dynamic programming formulation of the revenue maximization problem and of the welfare maximization problem. Then, they show that this dynamic congestion pricing mechanism can be well approximated by a simpler static time-of-day pricing. An alternative mechanism called “Trade & Cap”, was recently proposed by Londoño, Bestavros and Laoutaris [Londoño et al., 2010]. It works in two phases. First, users engage in a trading game where they choose an amount of *reserved* bandwidth slots to buy for hard-constraints traffic. In the second phase, the remaining bandwidth is allocated to users as *fluid* bandwidth, in proportion of their remaining “buying power”. They show that this mechanism smoothes the aggregate demand to a certain level. In their model, users have a cost function that increases linearly with the total demand in a given slot. Although interesting, these schemes are complex, do not consider the elasticity of user demand do not allow user utility to be an arbitrary function of the congestion level.

Prior to our study, two recent papers had analyzed time-of-day pricing mechanisms over n time slots [Jiang et al., 2008, Joe-Wong et al., 2011]. In [Jiang et al., 2008], Jiang, Parekh and Walrand consider a model where users have unit demand. Each user chooses one time-slot in which he transmits its entire demand, to maximize his utility. The authors of [Jiang et al., 2008] obtain a bound on the price of anarchy due to users selfishness. In [Joe-Wong et al., 2011], Wong, Ha and Chiang consider a model with users transmitting *sessions* of random length. Sessions arrive as a Poisson process and each session is characterized by a *waiting function* which reflects the willingness of the user to delay his entire session for a given time, in exchange for a reward given by the provider. The authors show how to compute the optimal reward levels in order to maximize the provider profit by balancing the congestion cost due to demand exceeding capacity and the reward amount. Further analysis of this mechanism called “TUBE”, as well as implementation are provided in [Ha et al., 2012]. However, in their model, users are only sensitive to prices (the effect of congestion on the user utility is not considered) and the analysis is not game-theoretic. Their model also suffer from other limitations that we mention below in the variant for electricity networks. Finally, all models mentioned above rely on optimizing the parameters of a scheme based on estimations of the consumers response (i.e., utility).

Many recent studies have also focused on flattening the demand curve in electricity networks. There, the use of incentives is usually referred to as a Demand Response (DR hereinafter) program, and DR programs are envisioned to be a key feature of the Smart Grid paradigm [Albadi and El-Saadany, 2008] where, by means of economic incentives (discounts or penalties), users are encouraged to rearrange their consumption in response to the network state, thus mitigating the grid overload and driving wholesale prices down.

Several analytical models are available in the literature, which describe and quantify the effects of DR mechanisms. Among these contributions, the authors of [Joe-Wong et al., 2012] adapt the model of [Joe-Wong et al., 2011] to the electricity framework. They study how an energy provider should select time-dependent discounts to minimize its production costs. They assume that the percentage of users who shift their consumption from slot i to slot j is a decreasing function of the temporal distance between slots i and j and a concave and increasing function of the discount offered in slot j (R_j), *independent* from discounts in other slots. This leads to a simple convex optimization problem for the selection of optimal time-dependent discounts. However, the

assumption that the percentage of users who shift their consumption from slot i to slot j is independent from discounts in other slots is not realistic if we assume that each user chooses the time slot that gives her maximal utility (hence compares the discounts in each slot); unless the provider was able to make personalized offers and have very precise forecast of the baseline consumption of each user, which is unlikely. Several other studies consider similar models of time-dependent pricing with n time slots [Li et al., 2014, Subramanian et al., 2013, Yang et al., 2014, Song et al., 2014] and which suffer from similar limitations. Interestingly, we observe that these limitations often come from the fact that the papers start from a macroscopic description of the population that, if investigated more closely, turns out to hide assumptions about individual users that are not consistent with reasonable microscopic (user-level) models.

In summary, we observe that the literature has two main shortcomings:

- All incentive schemes rely on the optimization of parameters from the estimation of the users utility, which is sensitive to estimation errors.
- Time-dependent pricing models with n time slots are based on simplified macroscopic assumptions that are not consistent with microscopic user-level models.

4.2.2 Our contributions

In this context, we made two main contributions addressing the two shortcomings mentioned above. The first is the proposition of a “fixed budget rebate” incentive mechanism based on the ideas of lotteries and the proof that this mechanism is more robust than previous mechanism with fixed discounts to errors in the estimation of the users utilities. We present here only the main elements of our model and the main results, all details can be found in [15].

Model and incentive mechanisms

We consider a population with a large number of users represented as a continuum, sharing a resource of fixed capacity. There are two time slot: peak and off-peak. Each user has a type $\theta \in \Theta$ that parameterizes his utility. The types distribution in the population are represented by a measure μ .

In [15], we develop a complete model where each user chooses his consumption in the peak and off-peak time, with elastic demand, and then show that under very mild assumptions (that the off-peak time has very low delay due to congestion), this model is equivalent to a simplified model where each user chooses a quantity $x_\theta \in [0, d_p]$ corresponding to the reduction of demand in the peak time compare to a maximal peak time demand d_p (e.g., a cap), to maximize a utility

$$u_\theta(x_\theta, G) = \bar{u}_\theta + (d_p - x_\theta)h(G) - c_\theta(x_\theta) - p, \quad (4.1)$$

where

$$G = \int_{\Theta} x_\theta d\mu(\theta) \quad (4.2)$$

is the aggregate peak-time demand reduction, which we view here as a *public good*. In (4.1), \bar{u}_θ is a constant corresponding to the maximal user utility in the absence of delay, $h(\cdot)$ corresponds to the (unit) benefit from the public good, i.e., to the unit benefit from lower congestion at peak time (computed from the delay reduction) and $c_\theta(\cdot)$ corresponds to the cost of contribution (coming from a decrease in utility representing the fact that the user prefers peak time), and p is the price of

a monthly subscription. We assume that $h(\cdot)$ is twice differentiable, strictly concave and increasing; $c_\theta(\cdot)$ is positive, differentiable and strictly convex and increasing; and $\sup_{\theta \in \Theta} c'_\theta(d_p) < \infty$, ($\theta \in \Theta$).

We define a Nash equilibrium of the non-atomic game defined above as a function $x^{(\text{eq})} : \Theta \rightarrow [0, d_p]$ such that for all $\theta \in \Theta$, $u_\theta(x_\theta, x_{-\theta}^{(\text{eq})}) \leq u_\theta(x_\theta^{(\text{eq})}, x_{-\theta}^{(\text{eq})})$, $\forall x_\theta \in [0, d_p]$. As is typical in problems involving public goods, the Nash equilibrium of the non-atomic game defined above does not coincide with the profile x^* that maximizes the social welfare defined by

$$W = \int_{\Theta} u_\theta d\mu(\theta). \quad (4.3)$$

To align Nash equilibrium and social optimum objectives, the service provider can design mechanisms to incentivize users to reduce their peak-time demand. Here, we compare two different incentive mechanisms: a fixed-budget rebate mechanism (denoted R or FBR) and a time-of-day pricing mechanism (denoted T or TDP). Each mechanism introduces a reward based on the peak-time demand reduction x_θ below the maximum d_p . For the service provider to finance the respective reward, each mechanism also introduces an increase in the subscription price. (However, we show that each user's net utility can be improved even with this price increase.) With mechanism $j \in \{R, T\}$, the user utility becomes

$$u_\theta^j(x_\theta, G) = u_\theta(x_\theta, G) + M^j(x_\theta, G), \quad (\theta \in \Theta). \quad (4.4)$$

The fixed-budget rebate mechanism consists in giving each user a reward proportional to his fraction of the total contribution, i.e., of the functional form:

$$M^R(x_\theta, G) = R \cdot \frac{x_\theta}{G} - \Delta p_R, \quad [\text{fixed-budget rebate}] \quad (4.5)$$

where R is a parameter of the mechanism chosen *ex-anti* by the provider. In practice, this mechanism could be implemented via randomization or using other implementations. We notice here that the fixed-budget rebate mechanism introduces uncertainty in the users bill as the reward depends on the amount shifted by the other users. However, this uncertainty is only one-sided: the maximum bill is known and only the reward amount is uncertain. This asymmetry is crucial to ensure good adoption of the mechanism.

The time-of-day pricing mechanism corresponds to a fixed reward per unit of shifted demand:

$$M^T(x_\theta, G) = r \cdot x_\theta - \Delta p_T, \quad [\text{time-of-day pricing}] \quad (4.6)$$

where r is a parameter of the mechanism chosen *ex-anti* by the provider. This mechanism is a variation of a conventional time-of-day pricing mechanism, with an off-peak price subsidy. Its implementation is straightforward.

In (4.5) and (4.6), Δp_j denotes the increase in the subscription price that the service provider imposes to finance the reward mechanism. Let $G^{(\text{eq})}$ be the equilibrium level of public good (we show that the Nash equilibrium is unique for both mechanism). We assume that the price Δp_j is fixed in advance by the service provider to compensate the reward, i.e., such that $\int_{\Theta} M^j(x_\theta, G^{(\text{eq})}) d\mu(\theta) = 0$ (note that the expression of the aggregate welfare (4.3) is thus not directly modified by the mechanisms, but only through the chosen contributions x_θ). Then,

$$\Delta p_R = R \cdot \frac{d_p}{D_p} \quad \text{and} \quad \Delta p_T = r G^{(\text{eq})} \cdot \frac{d_p}{D_p}, \quad (4.7)$$

where $D_p = \int_{\Theta} d_p d\mu(\theta) = d_p \mu(\Theta)$ is the aggregate maximal peak-time demand. From (4.7), we immediately see that the service provider has to know the equilibrium to determine the price Δp_T for the time-of-day pricing mechanism. An error in the estimation of $G^{(\text{eq})}$ could have important consequences. In contrast, such knowledge is not necessary for the fixed-budget rebate mechanism where Δp_R only depends on the parameter R chosen by the service provider.

Overview of the results

Based on the model described above, we are able to establish the following results:

Social optimum: There exists a unique function x^* that maximizes the social welfare W .

Nash equilibrium: For both mechanism and for any value of the parameter R or r , there exists a unique Nash equilibrium profile $x^{(\text{eq})}$. Moreover, we have $x^{(\text{eq})} = x^*$ for well chosen values of the parameters:

$$R^* = G^* h'(G^*)(D - G^*), \quad (4.8)$$

and

$$r^* = h'(G^*)(D - G^*). \quad (4.9)$$

Variations with the parameter: For both mechanisms, the contribution of each user increases with the reward parameter value. The social welfare increases until its maximum achieved at R^*, r^* and then decreases. There is a range of parameters around R^*, r^* for which the social welfare is higher than with no mechanism.

Finally, our main result with this model is to show that the FBR mechanism is more robust than the TDP mechanism. To formalize that, we introduced perturbations of the cost-of-contribution functions (arguably the most difficult for the provider to estimate well) of the form

$$\tilde{c}_{\theta}(\cdot) = c_{\theta}(\cdot) + \epsilon \cdot p_{\theta}(\cdot), \quad (4.10)$$

where ϵ is a real number and $p_{\theta} : [0, d_p] \rightarrow \mathbb{R}$ is a continuously differentiable function satisfying

$$\sup_{\theta \in \Theta} \sup_{x \in [0, d_p]} |p'_{\theta}(x)| < \infty.$$

We define $G_R^{(\text{eq})}(\epsilon)$ and $G_T^{(\text{eq})}(\epsilon)$ the equilibrium levels of public good in the games with perturbed utilities, and similarly $W_R^{(\text{eq})}(\epsilon)$ and $W_T^{(\text{eq})}(\epsilon)$ the corresponding equilibrium welfares. Finally, let $G^*(\epsilon)$ and $W^*(\epsilon)$ be the socially optimal level of public good with perturbed utilities. We introduce the following conditions:

$$\text{(C1')} \quad |r'_R(G) - r'_{SO}(G)| < |r'_T(G) - r'_{SO}(G)|, \text{ at } G = G^*(0),$$

$$\text{(C2')} \quad |r'_R(G) - r'_{SO}(G)| > |r'_T(G) - r'_{SO}(G)|, \text{ at } G = G^*(0),$$

where r'_R, r'_T, r'_{SO} are the respective derivatives of the unit rewards

$$r_R(G) = \frac{R}{G}, \quad r_T(G) = r, \quad r_{SO}(G) = h'(G)(D - G). \quad (4.11)$$

Then we are able to state our result as follows.

Proposition 4.2.1. *There exists $\epsilon_m > 0$ such that, for any perturbation (4.10) with $\epsilon \neq 0$ and $|\epsilon| < \epsilon_m$,*

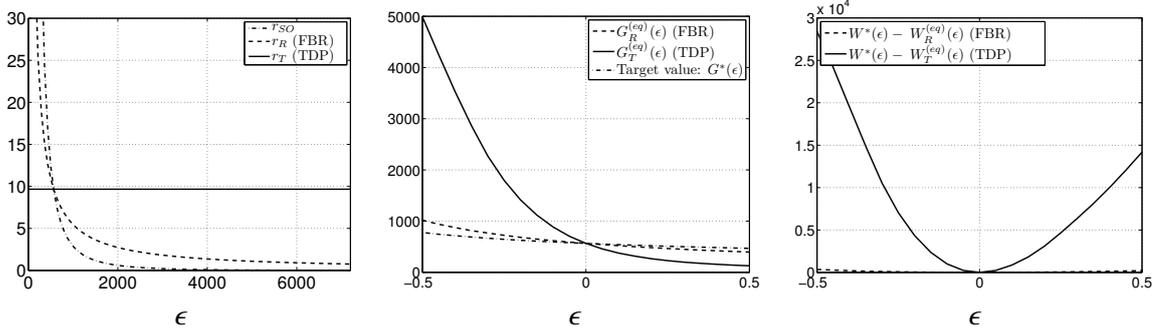


Figure 4.4: Illustration of the robustness of the FBR mechanism.

(i) if condition (C1') is satisfied, then

$$\left| G_R^{(eq)}(\epsilon) - G^*(\epsilon) \right| < \left| G_T^{(eq)}(\epsilon) - G^*(\epsilon) \right|;$$

(ii) if condition (C2') is satisfied, then

$$\left| G_R^{(eq)}(\epsilon) - G^*(\epsilon) \right| > \left| G_T^{(eq)}(\epsilon) - G^*(\epsilon) \right|.$$

The intuition behind Proposition 4.2.1 is the following: the mechanism with the unit reward closer to the optimal unit reward $r_{SO}(G)$ have an equilibrium closer to the social optimum equilibrium $G^*(\epsilon)$. Since $r_R(G)$ and $r_{SO}(G)$ are both decreasing functions, one expects $r_R(G)$ to be closer to $r_{SO}(G)$ than $r_T(G)$. It is often the case. The fact that $r_R(G)$ decreases when G increases is the *closed-loop* effect: the more users reduce their peak-time demand, the lower the incentive to reduce it is. This is the main idea behind the robustness of the FBR mechanism. However, if $r_R(G)$ decreases much faster than $r_{SO}(G)$, $r_T(G)$ can be closer to $r_{SO}(G)$. This possibility is covered by case (ii) of Proposition 4.2.1.

From Proposition 4.2.1, we deduce the our main robustness result in terms of social welfare.

Theorem 4.2.2. *There exists $\epsilon_m > 0$ such that, for any perturbation (4.10) with $\epsilon \neq 0$ and $|\epsilon| < \epsilon_m$, we have:*

(i) if condition (C1') is satisfied, then the fixed-budget rebate mechanism is more robust than the time-of-day pricing mechanism:

$$W_T^{(eq)}(\epsilon) < W_R^{(eq)}(\epsilon) < W^*(\epsilon);$$

(ii) if condition (C2') is satisfied, then the time-of-day pricing mechanism is more robust than the fixed-budget rebate mechanism:

$$W_R^{(eq)}(\epsilon) < W_T^{(eq)}(\epsilon) < W^*(\epsilon).$$

Figure 4.4 illustrates the robustness of the FBR mechanism as expressed in Proposition 4.2.1 and in Theorem 4.2.2, for a realistic scenario (see details of the utilities and parameters in [15]).

Note that in the previous exposition derived from [15], we used a non-atomic game directly. In another paper [34], we started from a model with a finite number of players and studied the convergence as the number of players tends to infinity.

The case of n time slots for electricity networks

Finally, our second contribution [23] was targeted to electricity networks, where usage patterns tend to be more complex and it makes more sense to consider a model with n time slots. In this context, finding an appropriate incentive scheme and optimizing the parameters is difficult, let alone studying the robustness. As mentioned above, most papers start from a macroscopic model of the population that leads to an easy optimization problem to compute the optimal prices or discounts. Our contribution started by showing that these macroscopic models hide assumptions and implicit requirements that are not realistic when looking at the user level; in particular that the provider can make personalized offers and has a perfect estimate of each individual's consumption in each time slot. Then, we explored four DR mechanisms with different levels of complexity:

1. the *base* mechanism corresponds to an optimization problem similar to the one considered in [Joe-Wong et al., 2012], it requires personalized offers and individual consumption forecasts; the energy production cost is optimized over the discount values, each of which is offered to a given fraction of the population,
2. the *optimized* mechanism takes full advantage of personalized offers and consumption forecasts by minimizing the cost over both the discount values and the population fractions to which the discounts are offered,
3. the *robust* mechanism relies on personalized offers, but does not need individual consumption forecasts,
4. finally the *broadcast* mechanism (analogous to that in [Yang et al., 2014]) needs neither of the two features and just broadcasts the discounts to all users.

Interestingly, contrarily to prior studies, we find that the cost-minimization problems resulting from our DR mechanisms are not convex (even for the base mechanism). Nevertheless, simple heuristics can identify (potential) minima in a reasonable amount of time in realistic scenarios. Then, our numerical results show that the simpler robust and broadcast mechanisms achieve significantly lower cost reductions than the optimized mechanism, which is difficult to implement, but that the gap reduces when the population's flexibility increases. Our results argue for more detailed models at the user level to evaluate the feasibility and benefits of incentive schemes. In future work, we plan to see if we can also make the incentive schemes more robust by introducing ideas from lotteries as we did in the simpler case of two time slots.

Summary of our contribution to incentives for network decongestion:

We propose a new incentive scheme (fixed-budget rebate) for the decongestion of the peak time in networks in a two time-slots model; and show that it is more robust than other schemes to errors in the estimation of the users utilities. For the case of n time-slots for electricity networks, we show that existing models hide crucial unrealistic assumptions, we argue for the use of user-level models and we study different possible schemes under such a model.

4.3 Cloud resources allocation

Cloud computing has now become the main paradigm for computing, storage and many other components that are crucial for user-centric services. The design of appropriate pricing schemes for cloud resources is therefore of great importance. Given the capacity of a cloud, an important objective is to increase the resource utilization so as to allow more tenants to be served. In other

words, the pricing problem cannot be separated from the question of tasks scheduling in the cloud because we need to evaluate, for a given demand, the performance that can be obtained. In this context, our initial contribution [11], [24] is in the area of tasks scheduling. We mention it here very briefly although it does not involve game theory or statistical learning, because we consider it as a pre-requisite to move towards pricing which will involve both game theory and statistical learning.

4.3.1 Context

In an effort to improve the efficiency of operation of clouds by leveraging the flexibility of user demand, there has been a significant attention on mechanisms to allow tenants to describe more precisely the characteristics of their jobs [Jalaparti et al., 2012, Ishakian et al., 2012], in particular *deadlines*. In the meantime, such technical progress also raises new algorithmic challenges on how to optimally schedule a set of malleable tasks with deadlines [Jain et al., 2011, Jain et al., 2012, Lucier et al., 2013, Menache et al., 2014, Azar et al., 2015, Bodik et al., 2014].

A fundamental model in this direction is the one of [Jain et al., 2011, Jain et al., 2012] in which a set of n malleable batch tasks has to be scheduled on C identical machines. All the jobs are available from the start and each of them is specified by a workload, a parallelism bound, a deadline and a value. Here, the number of machines assigned to a task can change during the execution and the parallelism bound decides the maximum amount of machines that can process a task simultaneously; however, the workload that is needed to complete a task will not change with the number of machines. Beyond the analysis of this basic and important model, efforts have been devoted to its online version [Lucier et al., 2013, Menache et al., 2014, Azar et al., 2015] and its extension [Bodik et al., 2014, Ferguson et al., 2012, Nagarajan et al., 2013] in which each task contains several subtasks with precedence constraints.

For the fundamental model in [Jain et al., 2011, Jain et al., 2012] under the objective of maximizing social welfare (i.e., the sum of values of tasks completed before their deadline), Jain et al. have proposed an $(1 - \frac{C}{2C-k})(1 - \epsilon)$ -approximation algorithm via deterministic rounding of linear program in [Jain et al., 2011] and a greedy algorithm GreedyRTL via dual fitting technique that achieves an approximation ratio of $\frac{C-k}{C} \cdot \frac{s-1}{s}$ in [Jain et al., 2012]. Here, k is the maximal parallelism bound of tasks, and $s (\geq 1)$ is the slackness which intuitively characterizes the resource allocation flexibility (e.g., $s = 1$ means that the maximal amount of machines have to be allocated to the task at every time slot until its deadline to ensure full completion). However, it seems difficult to improve those results using linear-programming-based techniques.

4.3.2 Our contributions

In this context, we proposed a new method to address the problem of scheduling malleable batch tasks with deadlines. Our results can be summarized as follows (see details in [24]).

Core result. Our core result is the first optimal scheduling algorithm so that C machines are optimally utilized by a set of malleable batch tasks \mathcal{S} with deadlines in terms of resource utilization. We first identify the basic constraints of malleable tasks and the optimal state in which C machines can be said to be optimally utilized by a set of tasks. Then, we propose a scheduling algorithm LDF(\mathcal{S}) that achieves such an optimal state.

Applications. This core result has applications in several new or existing algorithmic design and analysis problems for scheduling malleable tasks under different objectives. In particular, we provide:

- (1) an improved greedy algorithm GreedyRLM with an approximation ratio $\frac{s-1}{s}$ for the social welfare maximization problem with a polynomial time complexity of $O(n^2)$;
- (2) the first exact dynamic programming (DP) algorithm for the social welfare maximization problem with a pseudo-polynomial time complexity of $O(\max\{n^2, nC^L M^L\})$;
- (3) the first exact algorithm for the machine minimization problem with a polynomial time complexity of $O(n^2)$.

Here, L , D , k and M are the number of deadlines, the maximal workload, the maximal parallelism bound, and the bound of the maximal deadline of tasks. Finally, we also prove that $\frac{s-1}{s}$ is the best approximation ratio that a general greedy algorithm can achieve.

Summary of our contribution to cloud resource allocation:

We propose new algorithms for scheduling malleable tasks with deadlines on multiple machines.

CONCLUSION AND PERSPECTIVES

User-centric services based on personal data have taken a central place in our lives and economies. With great improvements to our lives, however, they also brought a set of important questions, in particular critical security and privacy problems that potentially threaten our well-being as well as the development and sustainability of the user-centric services ecosystem. This manuscript summarized our contributions to developing and studying methods that can be useful to address security, privacy and network systems performance problems.

Through our work presented in this manuscript, we showed that game theory is a key tool to tackle security, privacy and network systems performance questions because these questions involve human users whose behaviors are determined by incentives provided by the system. Statistical learning is also a key tool because it is at the core of user-centric services, both to secure the system and to exploit personal data. Even more importantly than those two tools in isolation though, Chapters 2 and 3 reveal that what is really needed to tackle security and privacy issues in user-centric services is a *combination of game theory and statistical learning* to find new learning algorithms that work well in situations where strategic human agents can alter the data. This much broader research challenge, for which the results of Chapters 2 and 3 only offer very preliminary answers, is the main perspective of my research for the years to come. I develop it briefly below. I also mention a second research perspective that investigates how learning algorithms affect humans in terms of fairness, transparency and understandability. Whereas machine learning mostly aims at automating decision making, these two research perspectives aim at bringing back the human perspective and can be summarized as the study of *'humans versus machine learning'*.

5.1 Learning from strategic data

Statistical learning exploded in the past decades and had a tremendous impact in many application domains such as computer vision or bioinformatics. As we saw, it also had a great impact in the digital area by enabling the development of user-centric services exploiting personal data.

Despite this success, most current learning algorithms face a fundamental limitation: they were developed and studied under the assumption that the data distribution is independent of the learning algorithm (the 'i.i.d. assumption'). While this assumption holds for instance when the data is generated by Nature, it fails when the data is generated by a strategic agent whose objective depends on the algorithm's outcome—we call it *strategic data*; which occurs in several important applications. In particular, as we saw, digital user-centric services face two kinds of strategic data: personal data provided by privacy-conscious users and security data generated by attackers. In both cases, using a standard learning algorithm that ignores the strategic nature of the data can lead to poor performance and there is a need to design new learning algorithms to improve privacy

and security of user-centric services.

Our overall goal is to *create and study new learning algorithms adapted to strategic data for privacy and security*. To address this challenge, we propose to build on the methodology initiated in Chapters 2 and 3, that is to build game-theoretic models involving learning agents that takes into account the objective of the agents providing or generating data (and of the learner), and to use their solutions to derive optimal learning algorithms for strategic data. We aim at working towards a general theory of learning from strategic data in the two contexts of privacy and security which lead to different types of models.

Learning from strategic personal data of privacy-conscious users. We want to create and study algorithms to learn from personal data that optimize the precision of the models learned given the objectives of privacy-conscious agents revealing the data. To this end, we plan to develop models building on the one in Chapter 3 with many agents (the users) choosing the precision of the data they reveal with an objective that combines privacy concern and interest in the learning outcome (public good). We will move towards more realistic models with incomplete information using a Bayesian framework, and extend our setting to non-linear regression and to recommendation problems. Based on the analysis of the equilibrium solutions of those models, we will then be able to tackle key learning questions such as: what is the optimal learning algorithms? what is the performance of standard learning algorithms on strategic data? how can we quantify the inefficiency due to the strategic nature of data? or how does the learning accuracy vary with the number of users? This last question opens a very broad perspective of developing results parallel to the statistical learning theory results (risk bounds, sample complexity and link to VC dimension) for the case of strategic data. We know that, in very simplistic cases [27], the rate of decrease of the variance of the mean's estimate is slower than the standard $1/n$; hence we expect that most results will be different for the strategic data case.

Learning from strategic data in security. We want to create classification algorithms from strategic data for attack detection that handle well the adversarial scenario at stake. To this end, we plan to develop models building on the one in Chapter 2 where the strategic data is a usage pattern possibly generated by an attacker whose objective is to evade detection while achieving his attack goal. While the model in Chapter 2 assumes complete information, we will model incomplete information (on both sides) in a Bayesian setting and investigate solutions of both static models (Bayesian Nash equilibrium) and dynamic models (Nash equilibrium of repeated and stochastic games) that account for different aspects of the system. Again, these solutions should allow us to derive optimal learning algorithms (incl. sequential learning) for a variety of security scenarios involving strategic data.

We note that our goal of deriving sequential learning algorithms for the security scenarios involves the difficult question of solving nonzero-sum repeated games. We plan to approach that by using simple classes of games (such as almost zero-sum) for which we hope to be able to characterize the equilibrium and by looking for heuristics that can be shown to be almost optimal. Then, our work on sequential learning could have many more applications than attack detection, in security and outside. In security for instance, it could allow to do dynamic resource allocation using a sequential version of the Blotto game, and it could also be applied to many other security scenarios not involving digital systems (e.g., defense of critical infrastructure of cyber-physical systems). Outside security, it could be useful for instance for dynamic pricing under competition.

In our work so far, we voluntarily limited ourselves to models that hit a good trade-off between tractability and realism, so that we could extract clean and generic insights from the solutions. Naturally, as we go on to more complex situations, we can expect that models will be less tractable.

This will open a challenge of a more numerical or algorithmic nature on how to compute approximately optimal algorithms to learn from personal data.

Finally, the work proposed here only scratches the surface of the extent of statistical learning methods. If our approach is successful, we hope that it can be ported to many other statistical learning problems (clustering, causal analysis, non-parametric estimation, etc.).

5.2 Human-friendly learning algorithms

With the rise of user-centric services, there is also a growing concern about how the complex learning algorithms at their core affect the users in inconspicuous ways. While users appreciate the benefits of personalization, systems often reveal very little about the specifics of how it is performed and users often have no possibility of determining whether the service that they receive is biased or discriminating in any way, or more generally how the data that they provide affects the service that they receive. In fact, this problem goes much beyond the world of online user-centric services as data-driven decision making is progressively making its way in other domains such as hiring, police, credit approvals, etc. So far, however, research in the space has mostly been limited to providing anecdotal evidence of discrimination or of outrageous outcome from learning algorithms on specific systems.

Our goal is to bring transparency and fairness to decision making systems based on learning algorithms in a principled manner. This entails several different challenges:

Defining the notions of explanation and understandability: To some extent, bringing transparency about an algorithm implies being able to *explain* to users how the algorithm works in a way that they can *understand*, but it is very unclear how to define these two notions. Here, we want to provide a definition that both complies with the intuitive notion that we have of understanding and is rigorous to allow us to tackle theoretical questions of how to compute an explanation, how to evaluate it, etc.

Building tools to bring transparency: Our second objective is to build tools and methods to bring transparency in existing systems. A key challenge here is that it is usually too costly or impossible to track all data that a given system has about a user and is using (as input to the decision making algorithm) to select the service. To tackle this challenge, we want to investigate how to bring transparency from the observation of the outputs (i.e., the service received by users) by proposing a collaborative method based on collection (on a voluntary basis) of demographic information of users. Using this idea, we are currently working on building a tool to bring transparency to targeted advertising that would infer, from the outputs, what data the ad engine has about a user and why he was targeted with a particular ad. This work entails both systems and inference challenges.

Designing fair learning algorithms: Finally, we want to tackle the question, from the perspective of the algorithm's designer, of how to design an algorithm as efficient as possible under the constraint that the outcome is 'acceptable' for humans. Here, we will start by investigating acceptability as a notion of fairness, but that could in principle cover much more general constraints that offer many perspectives.

To conclude, let us finally remark that our two broad research perspectives on 'humans versus machine learning' are not without links. In particular our ideas to build tools to bring transparency involve collecting demographic information from users, which is personal data. We plan to let users choose the precision of the data given and to use the algorithms to learn from strategic

personal data mentioned earlier. That will allow us both to improve the quality of our transparency tool and to test, in the real-world, our new algorithms to learn from strategic data.

PUBLICATIONS

PhD dissertation

- [1] **Patrick Loiseau**. *Contributions to the Analysis of Scaling Laws and Quality of Service in Networks: Experimental and Theoretical Aspects*. PhD thesis, ENS Lyon, December 2009.

Edited volumes

- [2] **Patrick Loiseau**, Aaron Roth, and Adam Wierman. The 10th Workshop on the Economics of Networks, Systems and Computation (NetEcon 2015). *ACM Performance Evaluation Review*, December 2015. (Guest editorial).
- [3] John Chuang and **Patrick Loiseau**. The joint Workshop on Pricing and Incentives in Networks and Systems (W-PIN+NetEcon 2014). *ACM Performance Evaluation Review*, 42(3):2–3, December 2014. (Guest editorial).
- [4] Costas Courcoubetis, Roch Guérin, **Patrick Loiseau**, David Parkes, Jean Walrand, and Adam Wierman. Special Issue on Pricing and Incentives in Networks and Systems: Guest Editors' Introduction. *ACM Transactions on Internet Technology*, 14(2–3):8:1–8:3, October 2014. (Guest editorial).
- [5] **Patrick Loiseau**, David Parkes, and Jean Walrand. The joint Workshop on Pricing and Incentives in Networks and Systems (W-PIN+NetEcon 2013). *ACM Performance Evaluation Review*, 41(4):2–3, March 2014. (Guest editorial).
- [6] **Patrick Loiseau** and Jean Walrand. The first Workshop on Pricing and Incentives in Networks (W-PIN 2012). *ACM Performance Evaluation Review*, 40(2):12–13, September 2012. (Guest editorial).

Preprints

- [7] Athanasios Andreou, Oana Goga, **Patrick Loiseau**, and Krishna P. Gummadi. Identity vs. attribute disclosure risks for users with multiple social profiles, 2016. (Preprint. Under Review.).
- [8] Vijay Kamble, **Patrick Loiseau**, and Jean Walrand. Regret-optimal strategies for playing repeated games with discounted losses, 2016. (Preprint. Under Review. Available as arXiv:1603.04981.).
- [9] Lemonia Dritsoula, **Patrick Loiseau**, and John Musacchio. A game-theoretic analysis of adversarial classification, 2016. (Preprint. Under Review. Available as arXiv:1610.04972.).
- [10] Michela Chessa and **Patrick Loiseau**. A cooperative game-theoretic approach to quantify the value of personal data in networks, 2016. (Preprint. Available at <http://www.eurecom.fr/~loiseau/articles/ChessaLoiseau-quantif.pdf>).
- [11] Xiaohu Wu and **Patrick Loiseau**. Algorithms for scheduling malleable cloud tasks, 2016. (Preprint. Under Review. Available as arXiv:1501.04343.).

Articles in journals

- [12] Raimo Kantola, Hammad Kabir, and **Patrick Loiseau**. Cooperation and End-to-End in the Internet. *International Journal of Communication Systems*, 2016. To appear.
- [13] Hadrien Hours, Ernst Biersack, **Patrick Loiseau**, Alessandro Finamore, and Marco Mellia. A Study of the Impact of DNS Resolvers on CDN Performance Using a Causal Approach. *Computer Networks, Special issue on "Traffic and Performance in the Big Data Era"*, 2016. To appear.
- [14] Hadrien Hours, Ernst Biersack, and **Patrick Loiseau**. A causal approach to the study of TCP performance. *ACM Transactions on Intelligent Systems and Technology, Special Issue on "Causal Discovery and Inference"* (K. Zhang, J. Li, E. Bareinboim, B. Schölkopf, and J. Pearl, editors), 7(2):25:1–25:25, January 2016.
- [15] **Patrick Loiseau**, Galina Schwartz, John Musacchio, Saurabh Amin, and S. Shankar Sastry. Incentive mechanisms for internet congestion management: Fixed-budget rebate *versus* time-of-day pricing. *IEEE/ACM Transactions on Networking*, 22(2):647–661, 2014.
- [16] **Patrick Loiseau**, Claire Médigue, Paulo Gonçalves, Najmeddine Attia, Stéphane Seuret, François Cottin, Denis Chemla, Michel Sorine, and Julien Barral. Large deviations estimates for the multiscale analysis of heart rate variability. *Physica A*, 391(22):5658–5671, November 2012.
- [17] Paulo Gonçalves, Shubhabrata Roy, Thomas Begin, and **Patrick Loiseau**. Dynamic resource management in clouds: A probabilistic approach. *IEICE Transactions on Communications, special section on Networking Technologies for Cloud Services*, E95-B(8):2522–2529, 2012. (Invited paper).
- [18] Julien Barral and **Patrick Loiseau**. Large deviations for the local fluctuations of random walks. *Stochastic Processes and their Applications*, 121(10):2272–2302, 2011.
- [19] **Patrick Loiseau**, Paulo Gonçalves, and Pascale Vicat-Blanc Primet. A long-range dependent model for network traffic with flow-scale correlations. *Stochastic Models*, 27:333–361, 2011.
- [20] Edmundo Pereira de Souza Neto, Elmer Andrés Fernández, Patrice Abry, Béatrice Cuzine, **Patrick Loiseau**, Christian Baude, Jean Frutoso, Claude Gharib, and Xavier Martin. Application of cardiac autonomous indices in the study of neurogenic erectile dysfunction. *Urologia Internationalis*, 86(3):290–297, 2011.
- [21] **Patrick Loiseau**, Paulo Gonçalves, Guillaume Dewaele, Pierre Borgnat, Patrice Abry, and Pascale Vicat-Blanc Primet. Investigating self-similarity and heavy-tailed distributions on a large scale experimental facility. *IEEE/ACM Transactions on Networking*, 18(4):1261–1274, August 2010.
- [22] Edmundo Pereira de Souza Neto, Patrice Abry, **Patrick Loiseau**, Jean-Christophe Cejka, Marc-Antoine Custaud, Jean Frutoso, Claude Gharib, and Patrick Flandrin. Empirical mode decomposition to assess cardiovascular autonomic control in rats. *Fundamental & Clinical Pharmacology*, 21(5):481–496, October 2007.

Articles in refereed conferences

- [23] Alberto Benegiamo, **Patrick Loiseau**, and Giovanni Neglia. Dissecting demand response mechanisms: the role of consumption forecasts and personalized offers. In *Proceedings of the American Control Conference (ACC)*, July 2016.
- [24] Xiaohu Wu and **Patrick Loiseau**. Algorithms for scheduling deadline-sensitive malleable tasks. In *Proceedings of the 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, September 2015.
- [25] Oana Goga, **Patrick Loiseau**, Robin Sommer, Renata Teixeira, and Krishna Gummadi. On the reliability of profile matching across large online social networks. In *Proceedings of the 21st ACM SIGKDD conference on Knowledge Discovery and Data Mining (KDD)*, August 2015.

- [26] Hadrien Hours, Ernst Biersack, **Patrick Loiseau**, Alessandro Finamore, and Marco Mellia. A study of the impact of DNS resolvers on performance using a causal approach. In *Proceedings of the 27th International Teletraffic Congress (ITC)*, September 2015. (**Selected for submission of an extended version to Computer Networks special issue on "Traffic and Performance in the Big Data Era"**).
- [27] Michela Chessa, Jens Grossklags, and **Patrick Loiseau**. A game-theoretic study on non-monetary incentives in data analytics projects with privacy implications. In *Proceedings of the 28th IEEE Computer Security Foundations Symposium (CSF)*, July 2015.
- [28] Michela Chessa and **Patrick Loiseau**. The impact of the graph structure on a public good provision game: a cooperative approach with applications to personal data release on social networks. In *SING11-GTM2015 (European meeting on game theory)*, July 2015. (1-page abstract).
- [29] Michela Chessa, Jens Grossklags, and **Patrick Loiseau**. A short paper on the incentives to share private information for population estimates. In *Proceedings of the 19th International Conference Financial Cryptography and Data Security (FC)*, January 2015. (Short paper).
- [30] Galina Schwartz, **Patrick Loiseau**, and S. Shankar Sastry. The heterogeneous colonel blotto game. In *Proceedings of the International conference on network games, control and optimization (NETG-COOP)*, October 2014.
- [31] Stratis Ioannidis and **Patrick Loiseau**. Linear regression as a non-cooperative game. In *Proceedings of the 9th conference on Web and Internet Economics (WINE)*, December 2013.
- [32] LEMONIA DRITSOULA, **Patrick Loiseau**, and John Musacchio. A game-theoretical approach for finding optimal strategies in an intruder classification game. In *Proceedings of the 51st IEEE Conference on Decision and Control (CDC)*, December 2012.
- [33] LEMONIA DRITSOULA, **Patrick Loiseau**, and John Musacchio. Computing the nash equilibria of intruder classification games. In *Proceedings of the third Conference on Decision and Game Theory for Security (GameSec)*, November 2012. (Full paper).
- [34] **Patrick Loiseau**, Galina Schwartz, John Musacchio, Saurabh Amin, and S. Shankar Sastry. Congestion pricing using a raffle-based scheme. In *Proceedings of the International conference on network games, control and optimization (NETGCOOP)*, October 2011.
- [35] Oana Goga, **Patrick Loiseau**, and Paulo Gonçalves. On the impact of the flow-size distribution's tail index on network performance with TCP connections. In *Proceedings of the 29th International Symposium on Computer Performance, Modeling, Measurements and Evaluation (IFIP Performance)*, October 2011.
- [36] **Patrick Loiseau**, Galina Schwartz, John Musacchio, and Saurabh Amin. Incentive schemes for internet congestion management: Raffles versus time-of-day pricing. In *49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, September 2011.
- [37] **Patrick Loiseau**, Paulo Gonçalves, Julien Barral, and Pascale Vicat-Blanc Primet. Modeling TCP throughput: an elaborated large-deviations-based model and its empirical validation. In *Proceedings of the 28th International Symposium on Computer Performance, Modeling, Measurements and Evaluation (IFIP Performance)*, November 2010. **Best Paper Award Runner-up**.
- [38] **Patrick Loiseau**, Paulo Gonçalves, Stéphane Girard, Florence Forbes, and Pascale Vicat-Blanc Primet. Maximum likelihood estimation of the flow size distribution tail index from sampled packet data. In *Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems (ACM SIGMETRICS / Performance)*, June 2009.
- [39] **Patrick Loiseau**, Paulo Gonçalves, Romaric Guillier, Matthieu Imbert, Yuetsu Kodama, and Pascale Vicat-Blanc Primet. Metroflux: A high performance system for analyzing flow at very fine-grain. In *Proceedings of the 5th International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom)*, April 2009.

Articles in refereed workshops

- [40] Hadrien Hours, Ernst Biersack, and **Patrick Loiseau**. Causal study of network performance. In *Proceedings of the 17ème Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications (AlgoTel)*, June 2014.
- [41] Hadrien Hours, Ernst Biersack, and **Patrick Loiseau**. A causal study of an emulated network. In *10ème Atelier en Evaluation de Performances (AEP10)*, June 2014.
- [42] **Patrick Loiseau**, Paulo Gonçalves, Romaric Guillier, Matthieu Imbert, Oana Goga, Yuetsu Kodama, and Pascale Vicat-Blanc Primet. *Metroflux*: a high performance system for very fine-grain flow analysis. In *Grid'5000 Spring School*, April 2009.
- [43] **Patrick Loiseau**, Paulo Gonçalves, and Pascale Vicat-Blanc Primet. How TCP can kill self-similarity. In *Euro-NF workshop: Traffic Engineering and Dependability in the Network of the Future*, September 2008.
- [44] **Patrick Loiseau**, Paulo Gonçalves, Yuetsu Kodama, and Pascale Vicat-Blanc Primet. Metroflux: A fully operational high speed metrology platform. In *Euro-NF workshop: New trends in modeling, quantitative methods and measurements, in cooperation with NET-COOP*, September 2008.
- [45] **Patrick Loiseau**, Paulo Gonçalves, Guillaume Dewaele, Pierre Borgnat, Patrice Abry, and Pascale Vicat-Blanc Primet. Vérification du lien entre auto-similarité et distributions à queues lourdes sur un dispositif grande échelle. In *9ème Atelier en Evaluation de Performances (AEP9)*, June 2008.
- [46] **Patrick Loiseau**, Paulo Gonçalves, and Pascale Vicat-Blanc Primet. A comparative study of different heavy tail index estimators of the flow size from sampled data. In *Proceedings of the MetroGrid Workshop, within the framework of GridNets International Conference*, October 2007.

Demonstrations

- [47] **Patrick Loiseau**, Romaric Guillier, Oana Goga, Matthieu Imbert, Paulo Gonçalves, and Pascale Vicat-Blanc Primet. Automated traffic measurements and analysis in Grid5000, June 2009. ACM SIGMETRICS / Performance demonstration contest (**Best Student Demonstration Award**).

SUMMARY OF MY PROFESSIONAL ACTIVITIES AND ADVISING EXPERIENCE AND COMPLETE CV

I joined at EURECOM as an assistant Professor in November 2011. Since then, my main professional activities have been the following:

Establishment of an independent research team on game theory and statistical learning for security, privacy and network systems. Prior to my arrival, there was no activity in game theory, network economics or statistical learning theory in the department. I have built an independent research team in the department around those topics, which is now well recognized worldwide.

My team fits well in the rest of the department's activities (my work has connections with security, data sciences, privacy and cloud computing), and is very well connected nationally and internationally. My main current collaborations are with: Oana Goga (data mining for security and privacy, MPI-SWS, Germany), Jens Grossklags (economics of security and privacy, PennState, PA, USA), Krishna Gummadi (security and privacy of social computing systems, MPI-SWS, Germany), Stratis Ioannidis (game theory and personal data, Northeastern University), John Musacchio (UC Santa Cruz, CA, USA), Giovanni Neglia (smart grid, Inria Sophia-Antipolis), Galina Schwartz (economics, UC Berkeley, CA, USA), Jean Walrand (game theory and network economics, UC Berkeley, CA, USA).

As part of building major international collaborations, I spent extended sabbaticals periods at UC Berkeley (2 months in 2012) and MPI-SWS (3 months in 2014, 6 months in 2016 under a Humboldt experienced researcher award). I have also attracted major academic personalities to visit my team at EURECOM: Jens Grossklags (PennState, stayed 3 months in sabbatical at EURECOM), Jean Walrand (UC Berkeley, visited for one week) and Galina Schwartz (UC Berkeley, visited for one week).

As part of building an independent research team, I have been hiring and advising 1 postdoc, 5 PhD students and 5 interns (see page 5-6 of my detailed CV for a list and the dedicated section below for details).

I have been disseminating my research results through publications and invited talks (see pages 6-11 of my detailed CV).

The scientific quality of my research has been recognized by various invitations of scientific nature (guest lectures in summer schools, invited talks in prestigious seminars, participation to program committees, guest editor of special issues, associate editor of ACM TOIT, etc.). See the full list of my professional service and scientific responsibilities on pages 3-5 of my detailed CV.

Research fund raising. I have received funding from ANR (Tremplin-ERC program), Institut Mines-Telecom (Futur&Ruptures program), the Labex UCN@Sophia, the France-Berkeley fund,

Symantec (faculty gift), Nokia (Cifre), SAP (Cifre) and the Data Transparency lab; for a total of 790+k€. See pages 6 of my detailed CV for a full list.

I also currently have several grant proposals under review, the main ones being an ERC Starting Grant proposal, and an ANR-DFG (France-Germany) grant proposal. They are all in areas corresponding to the perspectives mentioned in Chapter 5 and aim at continuing to support a research group working in these topics under my supervision.

Contribution to local environment and to the scientific community in general. I have created Sophia-networking, a website that allows the SophiaTech campus to share information about seminars. In addition, I created the SophiaTech networks seminar, a bi-monthly seminar with invited speakers about networks, clouds, security and other connected topics, which later became the UCN@Sophia Labex seminar. This seminar receives financial support from the Labex, which allows me to bring renowned researchers to give talks at EURECOM—see http://www.sophia-networking.org/sophiatech_networks.

In 2012, I created with Jean Walrand the W-PIN workshop on network economics. In 2013, W-PIN merged with the prestigious NetEcon workshop and I was elected chair of the steering committee. Since then, I have remained chair of the SC and I was PC chair in 2012-15. In 2012-16, I also raised financial support for W-PIN and NetEcon from the Institut Mines-Telecom.

I am serving on the steering committees of the UCN@Sophia Labex and of the “thematic network 2” of the Institut Mines-Telecom.

See the full list of my professional service activities on pages 3-5 of my detailed CV.

Teaching and advising. When arriving at EURECOM, I took over and enhanced the game theory course (short course, 21h), and I have built two new courses: network economics (a research-oriented short course, 21h) and statistical data analysis (a basic short course on the fundamentals of statistics, 21h). I teach each of these three courses as the main instructor once a year. See page 2-3 of my detailed CV for a full description of my teaching activities. In 2014-16, I was also the responsible for the networking track of the EURECOM engineering program.

Since my start at EURECOM, I have advised students at all levels (see also list in pages 5-6 of my detailed CV):

Postdocs

- Michela Chessa (Sept. 2013–Aug. 2015, now assistant Professor at University of Nice Sophia-Antipolis)
Topic: Game theory and economics of personal data
Output: [10], [27–29] (+ 1 paper ready to submit to journal within 1 month).

PhD students

- Quan Vu (Starting Dec. 2016, co-advised with Alonso Silva from Nokia Bell Labs)
Topic: Learning in Blotto Games and Applications to Modeling Attention in Social Networks
- Amine Lahouel (Since Sept. 2016, co-advised with Cédric Hébert from SAP labs)
Topic: Data anonymity / utility tradeoff in big data applications
- Athanasios Andreou (Since Oct. 2015, co-advised with Oana Goga from MPI-SWS)
Topic: Bringing transparency to personalized services through statistical inference
Output: [7]

- Xiaohu Wu (Nov. 2012–Feb. 2016, now postdoc at Aalto University)
Topic: Techniques for Scheduling and Pricing in Cloud Computing
Ph.D. Telecom ParisTech, thesis successfully defended on Feb. 16, 2016
Output: [11], [24] (+ 1 paper submitted in double-blind)
- Hadrien Hours (Nov. 2011–Sep. 2015, co-advised with Ernst Biersack from EURECOM, postdoc at ENS Lyon, now data scientist at Booking.com)
Topic: A causal approach to the study of telecommunication networks
Ph.D. Telecom ParisTech, thesis successfully defended on Sept 16, 2016
Output: [13, 14], [26], [40, 41]
- Alberto Benegiamo (Nov. 2013–Aug. 2015, co-advised with Giovanni Neglia from Inria)
Topic: Mathematical tools for smart grids
PhD interrupted after 18 months by mutual agreement
Output: [23] (+ an extended version ready to submit to journal within 1 month)

Interns

- Yannick Terme (Eng. EURECOM/Telecom ParisTech and ENSAE, intern at MPI-SWS in July-August 2016)
- Nina Grgić-Hlača (M.A. University of Zagreb, intern at EURECOM in Feb.-July 2016 with an ERASMUS+ grant)
- Vijay Kamble (Ph.D. UC Berkeley, intern at EURECOM in April-May 2015, now postdoc at Stanford)
Output: [8]
- Athanasios Andreou (M.Sc. EURECOM, intern at MPI-SWS in Feb.-Sept. 2015, co-advised with Oana Goga and Krishna Gummadi, now Ph.D. student at EURECOM)
Output: initial work on [7], won a grant “thèse d’excellence” from institut Mines-Telecom to do a PhD at EURECOM
- Yifan Pi (B.Sc. Tsinghua University, intern at EURECOM during summer 2013, now software engineer at Google)

Student projects

- Supervision of 15+ semester projects and master projects since 2012 (2-3 per semester) in the areas of game theory, social networks, security and privacy

Patrick Loiseau

EURECOM
Campus SophiaTech
450 Route des Chappes
06410 Biot
France

Phone (Office): +33 4 93 00 81 47
Phone (Mobile): +33 6 17 38 15 19
Email: patrick.loiseau@eurecom.fr
Homepage: <http://www.eurecom.fr/~loiseau/>

Research Interests

My research lies in the areas of game theory and statistics and their application to security, personal data and networks economics (designing better online systems and applications through good incentives). I like to do interdisciplinary research and to work on problems that have both fundamental theoretical aspects and important practical applications. My main current research interests include adversarial and game-theoretic statistical learning (i.e., statistical learning with strategic agents), multi-armed bandit games, game theory for security, cyber insurance, economics of personal data, statistical inference in online social systems, resource allocation and pricing in clouds and smart grids, and causal analysis.

Education

Université Pierre et Marie Curie (Paris 6) / École Polytechnique, Paris, France July 2010

M.Sc. Degree in Mathematics – *Probability and random models*

First class honors: “mention très bien”

Thesis: *Large deviations for mixing processes*

École Normale Supérieure de Lyon, Lyon, France Dec. 2009

Ph.D. in Computer Science, prepared at LIP lab., within the Inria RESO team

Thesis: *Contributions to the Analysis of Scaling Laws and Quality of Service in Networks: Experimental and Theoretical Aspects*

Advisors: Paulo Gonçalves, Pascale Vicat-Blanc Primet

Committee: Christophe Diot, Daniel Kofman, Jean-Yves Le Boudec (reviewer), Rudolf Riedi (reviewer),

Philippe Robert (reviewer)

École Normale Supérieure de Lyon, Lyon, France 2002 – 2006

“Elève normalien”: undergraduate and graduate studies at the physics department

M.Sc. Degree in Physics – *Non linear and statistical physics* (July 2006)

Thesis: *Complex wavelets for the analysis of scaling phenomena*

Degree of Professeur-Agrégé in physics (July 2005)

B.Sc. Degree in physics (July 2003)

Lycée Marcelin Berthelot, Saint-Maur des Fossés, France 2000 – 2002

French preparatory classes with physics and chemistry majors (classes préparatoires PCSI and PC*)

Admission to École Normale Supérieure de Lyon

Employment

EURECOM, Sophia-Antipolis, France

Nov. 2011 – present

Assistant Professor in the Data Science department (*first class* since July 2015)

(Previously in the former Networking and Security department until Jan. 2016)

- University of California**, Santa Cruz, CA, USA Dec. 2010 – Oct. 2011
 Post-doctoral scholar in Basking Engineering school, working with Prof. John Musacchio
 Research topic: *Game theory and application to network economics*
- Inria Paris-Rocquencourt**, Le Chesnay, France Jan. 2010 – Nov. 2010
 Post-doctoral fellow in Sisyphe team, working with Julien Barral and Michel Sorine
 Research topic: *Multi-scale analysis of heart-rate variability: estimation and control-theoretic modeling*
- École Normale Supérieure de Lyon**, Lyon, France Sept. 2006 – Dec. 2009
 Doctoral fellow at LIP lab. in Inria RESO team, supervised by Paulo Gonçalves and Pascale Vicat-Blanc
 Research topic: *Analysis and modeling of network traffic and performance: from theory to practice*
- École Normale Supérieure de Lyon**, Lyon, France Sept. 2002 – Aug. 2006
 “Elève fonctionnaire stagiaire”

Visiting positions

- Max Planck Institute for Software Systems**, Saarbrücken, Germany April 2016 – present
 Visiting researcher, hosted by Prof. Krishna Gummadi
Funded by a Humboldt Research Fellowship for experienced researchers
- Max Planck Institute for Software Systems**, Saarbrücken, Germany July 2014 – Sept. 2014
 Visiting researcher, hosted by Prof. Krishna Gummadi
- University of California**, Berkeley, CA, USA July 2012 – Aug. 2012
 Visiting researcher in the EECS department, hosted by Prof. Jean Walrand
- University of California**, Berkeley, CA, USA Dec. 2010 – Oct. 2011
 Visiting member of the Network Economics Group, hosted by Prof. Jean Walrand
- University of Waterloo**, Waterloo, ON, Canada Oct. 2010
 Visiting researcher in the ECE department, hosted by Prof. Ravi Mazumdar

Internships

- École Normale Supérieure de Lyon**, Lyon, France Apr. 2006 – July 2006
 Research Intern in the Physics lab., supervised by Patrice Abry, Pierre Borgnat and Paulo Gonçalves
 Research topic: *Complex wavelets for the analysis of scaling phenomena*
- École Normale Supérieure de Lyon**, Lyon, France May 2004 – July 2004
 Research Intern in the Physics lab., supervised by P. Abry, P. Flandrin and E. Pereira de Souza Neto
 Research topic: *Application of the Empirical Mode Decomposition to the study of the heart beat rate*
- École Normale Supérieure de Lyon**, Lyon, France June 2003 – July 2003
 Research Intern in the Chemistry lab., supervised by Vincent Krakoviak
 Research topic: *Numerical study of the pressure in a random porous matrix via Monte-Carlo simulations*

Teaching experience

- EURECOM**, Sophia-Antipolis, France
Statistical data analysis every Fall since 2013
 Instructor, graduate course (short)

<i>Game Theory</i> Instructor, graduate course (short)	every Fall since 2013
<i>Network Economics</i> Instructor, graduate course (short)	every Fall since 2012
<i>Performance Evaluation of Computer Systems</i> Instructor, graduate course (long)	Spring 2012 and 2013

University of California, Santa Cruz, CA, USA

ISM207: <i>Random Process Models in Engineering</i> Guest lecturer, graduate course (instructor: Prof. Musacchio), TIM program	Spring 2011
---	-------------

École Normale Supérieure de Lyon, Lyon, France

<i>Network traffic models</i> Guest lecturer, M2 graduate course (instructors: C. Touati and P. Gonçalves), CS department	Spring 2010
--	-------------

Université de Versailles Saint-Quentin-en-Yvelines, Versailles, France

<i>Introduction to probability</i> Part-time teacher, L2 undergraduate level (instructor: A. Rouault), Mathematics department	Fall 2010
--	-----------

École Normale Supérieure de Lyon, Lyon, France

Teaching assistant (“moniteur”) in the physics and CS departments 2006 – 2009

<i>Electromagnetic waves and telecommunications</i> Tutorials, graduate level (preparatory class to “agrégation” in physics)
<i>Introduction to signal processing</i> Lab. sessions, bachelor level (L3), physics program
<i>Principles of hydrodynamics, linear acoustics and shock waves</i> Tutorials, graduate level (preparatory class to “agrégation” in physics)
<i>Computer architecture, systems and networks</i> Tutorials, bachelor level (L3), fundamental CS program

MediPlus Lyon, Lyon, France

Part-time teaching for first year medicine and pharmacy students 2006 – 2009

<i>Basics of physics and biophysics</i> Lectures and tutorials, medicine program
<i>Basics of general mathematics and statistics</i> Lectures and tutorials, pharmacy program

Professional service

Teaching and internal responsibilities

- Responsible for the networking track of the engineering studies at EURECOM (since 2014)
- Member of the restricted committee for strategic reflection at EURECOM (2014-16)

Steering committees

- Chair of the steering committee of NetEcon (since 2013)

Chair of the steering committee of Sophia-networking (sophia-networking.org) (since 2013)
Member of the scientific council of the Labex UCN@Sophia (since 2015)
Member of the steering committee of the “thematic network 2” of Institut Mines-Telecom (since 2014)

Conference organization

Registration chair of ACM SIGMETRICS 2016
Co-organizer, seminar on modeling, optimization and control in wireless networks, Paris 2015
PC co-chair of NetEcon 2015 (with Aaron Roth and Adam Wierman)
PC co-chair of W-PIN+NetEcon 2014 (with John Chuang)
PC co-chair of W-PIN+NetEcon 2013 (with David Parkes and Jean Walrand)
Registration chair of ACM SIGMETRICS 2013
PC co-chair of W-PIN 2012 (with Jean Walrand)

Editorial activities

Guest editor, ACM TOIT special issue on economics of security and privacy (2016)
Associate editor, ACM TOIT (since 2015)
Guest editor, ACM TOIT special issue on pricing and incentives in networks and systems (2013)

TPCs

WiOpt 2017
NetGCooP 2016
WWW 2016 (demo track)
SDP 2016
FC 2016
ITC 2016
WWW 2015 (demo track)
SDP 2015
NetGCooP 2014
ACM SIGMETRICS 2014
SDP 2014
ICQT 2013
W-PIN+NetEcon 2013
SDP 2013
GameSec 2012
W-PIN 2012
CFIP 2009 (shadow)

Invited referee for journals and conferences (each listed only once)

ACM Transactions on Privacy and Security, IEEE Transactions on Information Forensics and Security, ACM Transactions on the Web, IEEE Networks, IEEE Transactions on Dependable and Secure Computing, ACM Transactions on Information and System Security, Operation Research, IEEE/ACM Transactions on Networking, International Journal of Information Security, ISAAC 2015, IEEE INFOCOM 2013, Computer Communications journal, IEEE Transactions on Communications, Computer Networks Journal, IEEE Communication Letters, Stochastic Models, ACM SigComm CCR, 20th ITC Specialist Seminar on Network Virtualization - Concepts and Performance, CFIP 2009

Evaluation of PhD dissertations

Committee member for the PhD of Antoine Rault (Inria Rennes, 2016)

Reviewer and committee member for the PhD of Áron Lászka (Budapest University of Technology and Economics, 2014)

Member of the mid-term evaluation committee for 7 PhD students (since 2013)

Panel member for grant proposal selection

Expert for the F.R.S.-FNRS, Belgium (2016)

Member of the selection committee for the “Future & Ruptures” program from IMT (2014)

External reviewer for the Informatics and Mathematics Panel of the Academic Research Council, Ministry of Education, Singapore (2014)

Professional membership

IEEE, ACM, Data Transparency lab, GDR Multifractal, GDR Jeux

Honors and awards

Data Transparency lab research grant (top 11% of the projects)	2016
Humboldt Research Fellowship for experienced researchers (Alexander von Humboldt Foundation)	2016
Symantec research faculty gift	2015
Data Transparency lab travel grant (top 30% of the projects)	2015
Best Paper Award Runner-up at IFIP Performance (6 papers selected)	2010
ERCIM Alain Bensoussan European Post-doctoral fellowship (declined)	2010
Best Student Demonstration Award at ACM SIGMETRICS/Performance	2009
PhD fellowship and teaching assistanship from École Normale Supérieure de Lyon	2006

Advising experience

Postdocs

Michela Chessa (Sept. 2013–Aug. 2015, now assist. Prof at University of Nice Sophia-Antipolis)

PhD students

Quan Dong Vu (Start expected Dec. 2016, co-advised with Alonso Silva from Nokia Bell Labs)

Amine Lahouel (Since Sept. 2016, co-advised with Cédric Hebert from SAP)

Athanasios Andreou (Since Oct. 2015, co-advised with Oana Goga from MPI-SWS)

Xiaohu Wu (Nov. 2012–Feb. 2016, Ph.D. Telecom ParisTech, now postdoc at Aalto University)

Hadrien Hours (Nov. 2011–Sep. 2015, Ph.D. Telecom ParisTech, co-advised with Ernst Biersack, postdoc at ENS Lyon, now data scientist at Booking.com)

Interns

Yannick Terme (Eng. EURECOM/Telecom ParisTech, intern at MPI-SWS in July-August 2016, now student at ENSAE)

Nina Grgić-Hlača (M.A. University of Zagreb, intern at EURECOM in Feb.-July 2016 with an ERASMUS+ grant, now PhD student at MPI-SWS)

Vijay Kamble (Ph.D. UC Berkeley, intern at EURECOM in April-May 2015, now postdoc at Stanford)

Athanasios Andreou (M.Sc. EURECOM, intern at MPI-SWS in Feb.-Sept. 2015, co-advised with Oana Goga and Krishna Gummadi, now Ph.D. student at EURECOM)

Yifan Pi (B.Sc. Tsinghua University, intern at EURECOM during summer 2013, now software engineer at Google)

Student projects

Supervision of 15+ semester projects and master projects since 2012 (2-3 per semester)

Research funding

- | | |
|--|-------------|
| ANR Tremplin-ERC | 2017 – 2018 |
| <i>CONNECTED: Towards secure and private personal-data-based online services in the networked world</i> (€ 150k) | |
| Patrick Loiseau (PI) | |
| Cifre contract with Nokia Bell Labs | 2016 – 2019 |
| <i>Learning in Blotto Games and Applications to Modeling Attention in Social Networks</i> (€ 45k) | |
| Patrick Loiseau (co-PI), Alonso Silva (co-PI at Nokia Bell Labs) | |
| Cifre contract with SAP Research | 2016 – 2019 |
| <i>Approche de l'anonymisation des données en fonction du niveau de risque associé</i> (€ 45k) | |
| Patrick Loiseau (co-PI), Cédric Hebert (co-PI at SAP Research) | |
| Data Transparency lab research grant | 2016 – 2017 |
| <i>TranspAd: A Collaborative Tool to Bring Transparency to Targeted Advertising</i> (€ 50k) | |
| Patrick Loiseau (co-PI), Oana Goga (co-PI) | |
| Institut Mines-Telecom Futur&Ruptures program, doctoral support grant | 2015 – 2018 |
| <i>TRANSPA: Bringing transparency to personalized services through statistical inference</i> (€ 108k) | |
| Patrick Loiseau (PI) | |
| France-Berkeley fund | 2014 – 2016 |
| <i>Multi-armed bandit games and applications</i> (\$ 10k) | |
| Patrick Loiseau (PI), Jean Walrand (PI) | |
| Symantec research faculty gift | 2015 |
| <i>Cyber insurance</i> (\$ 30k) | |
| Patrick Loiseau (PI) | |

- Institut Mines-Telecom Futur&Ruptures program**, post-doctoral support grant 2015
MONET: MONETization of personal data in social networks: A game-theoretic approach (€ 30k)
 Patrick Loiseau (PI)
- Labex UCN@Sophia**, post-doctoral support grant 2013 – 2015
PRIMO: PRivate data MONetization: a public good approach using cooperative game theory (€ 90k)
 Patrick Loiseau (PI)
- Labex UCN@Sophia**, doctoral support grant 2013 – 2016
Mathematical tools for the smart grid (€ 105k)
 Patrick Loiseau (co-PI), Giovanni Neglia (co-PI)
- Institut Mines-Telecom Futur&Ruptures program**, doctoral support grant 2012 – 2015
Robust pricing of cloud resources through mean-field games (€ 121k)
 Patrick Loiseau (PI)

Publications

PhD dissertation

- [1] **Patrick Loiseau**. *Contributions to the Analysis of Scaling Laws and Quality of Service in Networks: Experimental and Theoretical Aspects*. PhD thesis, ENS Lyon, December 2009.

Edited volumes

- [2] **Patrick Loiseau**, Aaron Roth, and Adam Wierman. The 10th Workshop on the Economics of Networks, Systems and Computation (NetEcon 2015). *ACM Performance Evaluation Review*, December 2015. (Guest editorial).
- [3] John Chuang and **Patrick Loiseau**. The joint Workshop on Pricing and Incentives in Networks and Systems (W-PIN+NetEcon 2014). *ACM Performance Evaluation Review*, 42(3):2–3, December 2014. (Guest editorial).
- [4] Costas Courcoubetis, Roch Guérin, **Patrick Loiseau**, David Parkes, Jean Walrand, and Adam Wierman. Special Issue on Pricing and Incentives in Networks and Systems: Guest Editors' Introduction. *ACM Transactions on Internet Technology*, 14(2–3):8:1–8:3, October 2014. (Guest editorial).
- [5] **Patrick Loiseau**, David Parkes, and Jean Walrand. The joint Workshop on Pricing and Incentives in Networks and Systems (W-PIN+NetEcon 2013). *ACM Performance Evaluation Review*, 41(4):2–3, March 2014. (Guest editorial).
- [6] **Patrick Loiseau** and Jean Walrand. The first Workshop on Pricing and Incentives in Networks (W-PIN 2012). *ACM Performance Evaluation Review*, 40(2):12–13, September 2012. (Guest editorial).

Preprints

- [7] Athanasios Andreou, Oana Goga, **Patrick Loiseau**, and Krishna P. Gummadi. Identity vs. attribute disclosure risks for users with multiple social profiles, 2016. (Preprint. Under Review.).
- [8] Vijay Kamble, **Patrick Loiseau**, and Jean Walrand. Regret-optimal strategies for playing repeated games with discounted losses, 2016. (Preprint. Under Review. Available as arXiv:1603.04981.).
- [9] Lemonnia Dritsoula, **Patrick Loiseau**, and John Musacchio. A game-theoretic analysis of adversarial classification, 2016. (Preprint. Under Review. Available as arXiv:1610.04972.).
- [10] Michela Chessa and **Patrick Loiseau**. A cooperative game-theoretic approach to quantify the value of personal data in networks, 2016. (Preprint. Available at <http://www.eurecom.fr/~loiseau/articles/ChessaLoiseau-quantif.pdf>).

- [11] Xiaohu Wu and **Patrick Loiseau**. Algorithms for scheduling malleable cloud tasks, 2016. (Preprint. Under Review. Available as arXiv:1501.04343.).

Articles in journals

- [12] Raimo Kantola, Hammad Kabir, and **Patrick Loiseau**. Cooperation and End-to-End in the Internet. *International Journal of Communication Systems*, 2016. To appear.
- [13] Hadrien Hours, Ernst Biersack, **Patrick Loiseau**, Alessandro Finamore, and Marco Mellia. A Study of the Impact of DNS Resolvers on CDN Performance Using a Causal Approach. *Computer Networks, Special issue on "Traffic and Performance in the Big Data Era"*, 2016. To appear.
- [14] Hadrien Hours, Ernst Biersack, and **Patrick Loiseau**. A causal approach to the study of TCP performance. *ACM Transactions on Intelligent Systems and Technology, Special Issue on "Causal Discovery and Inference"* (K. Zhang, J. Li, E. Bareinboim, B. Schölkopf, and J. Pearl, editors), 7(2):25:1–25:25, January 2016.
- [15] **Patrick Loiseau**, Galina Schwartz, John Musacchio, Saurabh Amin, and S. Shankar Sastry. Incentive mechanisms for internet congestion management: Fixed-budget rebate *versus* time-of-day pricing. *IEEE/ACM Transactions on Networking*, 22(2):647–661, 2014.
- [16] **Patrick Loiseau**, Claire Médigue, Paulo Gonçalves, Najmeddine Attia, Stéphane Seuret, François Cottin, Denis Chemla, Michel Sorine, and Julien Barral. Large deviations estimates for the multiscale analysis of heart rate variability. *Physica A*, 391(22):5658–5671, November 2012.
- [17] Paulo Gonçalves, Shubhabrata Roy, Thomas Begin, and **Patrick Loiseau**. Dynamic resource management in clouds: A probabilistic approach. *IEICE Transactions on Communications, special section on Networking Technologies for Cloud Services*, E95-B(8):2522–2529, 2012. (Invited paper).
- [18] Julien Barral and **Patrick Loiseau**. Large deviations for the local fluctuations of random walks. *Stochastic Processes and their Applications*, 121(10):2272–2302, 2011.
- [19] **Patrick Loiseau**, Paulo Gonçalves, and Pascale Vicat-Blanc Primet. A long-range dependent model for network traffic with flow-scale correlations. *Stochastic Models*, 27:333–361, 2011.
- [20] Edmundo Pereira de Souza Neto, Elmer Andrés Fernández, Patrice Abry, Béatrice Cuzine, **Patrick Loiseau**, Christian Baude, Jean Frutoso, Claude Gharib, and Xavier Martin. Application of cardiac autonomous indices in the study of neurogenic erectile dysfunction. *Urologia Internationalis*, 86(3):290–297, 2011.
- [21] **Patrick Loiseau**, Paulo Gonçalves, Guillaume Dewaele, Pierre Borgnat, Patrice Abry, and Pascale Vicat-Blanc Primet. Investigating self-similarity and heavy-tailed distributions on a large scale experimental facility. *IEEE/ACM Transactions on Networking*, 18(4):1261–1274, August 2010.
- [22] Edmundo Pereira de Souza Neto, Patrice Abry, **Patrick Loiseau**, Jean-Christophe Cejka, Marc-Antoine Custaud, Jean Frutoso, Claude Gharib, and Patrick Flandrin. Empirical mode decomposition to assess cardiovascular autonomic control in rats. *Fundamental & Clinical Pharmacology*, 21(5):481–496, October 2007.

Articles in refereed conferences

- [23] Alberto Benegiamo, **Patrick Loiseau**, and Giovanni Neglia. Dissecting demand response mechanisms: the role of consumption forecasts and personalized offers. In *Proceedings of the American Control Conference (ACC)*, July 2016.
- [24] Xiaohu Wu and **Patrick Loiseau**. Algorithms for scheduling deadline-sensitive malleable tasks. In *Proceedings of the 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, September 2015.

- [25] Oana Goga, **Patrick Loiseau**, Robin Sommer, Renata Teixeira, and Krishna Gummadi. On the reliability of profile matching across large online social networks. In *Proceedings of the 21st ACM SIGKDD conference on Knowledge Discovery and Data Mining (KDD)*, August 2015.
- [26] Hadrien Hours, Ernst Biersack, **Patrick Loiseau**, Alessandro Finamore, and Marco Mellia. A study of the impact of DNS resolvers on performance using a causal approach. In *Proceedings of the 27th International Teletraffic Congress (ITC)*, September 2015. (**Selected for submission of an extended version** to *Computer Networks special issue on "Traffic and Performance in the Big Data Era"*).
- [27] Michela Chessa, Jens Grossklags, and **Patrick Loiseau**. A game-theoretic study on non-monetary incentives in data analytics projects with privacy implications. In *Proceedings of the 28th IEEE Computer Security Foundations Symposium (CSF)*, July 2015.
- [28] Michela Chessa and **Patrick Loiseau**. The impact of the graph structure on a public good provision game: a cooperative approach with applications to personal data release on social networks. In *SING11-GTM2015 (European meeting on game theory)*, July 2015. (1-page abstract).
- [29] Michela Chessa, Jens Grossklags, and **Patrick Loiseau**. A short paper on the incentives to share private information for population estimates. In *Proceedings of the 19th International Conference Financial Cryptography and Data Security (FC)*, January 2015. (Short paper).
- [30] Galina Schwartz, **Patrick Loiseau**, and S. Shankar Sastry. The heterogeneous colonel blotto game. In *Proceedings of the International conference on network games, control and optimization (NETG-COOP)*, October 2014.
- [31] Stratis Ioannidis and **Patrick Loiseau**. Linear regression as a non-cooperative game. In *Proceedings of the 9th conference on Web and Internet Economics (WINE)*, December 2013.
- [32] Lemonia Dritsoula, **Patrick Loiseau**, and John Musacchio. A game-theoretical approach for finding optimal strategies in an intruder classification game. In *Proceedings of the 51st IEEE Conference on Decision and Control (CDC)*, December 2012.
- [33] Lemonia Dritsoula, **Patrick Loiseau**, and John Musacchio. Computing the nash equilibria of intruder classification games. In *Proceedings of the third Conference on Decision and Game Theory for Security (GameSec)*, November 2012. (Full paper).
- [34] **Patrick Loiseau**, Galina Schwartz, John Musacchio, Saurabh Amin, and S. Shankar Sastry. Congestion pricing using a raffle-based scheme. In *Proceedings of the International conference on network games, control and optimization (NETGCOOP)*, October 2011.
- [35] Oana Goga, **Patrick Loiseau**, and Paulo Gonçalves. On the impact of the flow-size distribution's tail index on network performance with TCP connections. In *Proceedings of the 29th International Symposium on Computer Performance, Modeling, Measurements and Evaluation (IFIP Performance)*, October 2011.
- [36] **Patrick Loiseau**, Galina Schwartz, John Musacchio, and Saurabh Amin. Incentive schemes for internet congestion management: Raffles versus time-of-day pricing. In *49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, September 2011.
- [37] **Patrick Loiseau**, Paulo Gonçalves, Julien Barral, and Pascale Vicat-Blanc Primet. Modeling TCP throughput: an elaborated large-deviations-based model and its empirical validation. In *Proceedings of the 28th International Symposium on Computer Performance, Modeling, Measurements and Evaluation (IFIP Performance)*, November 2010. **Best Paper Award Runner-up**.
- [38] **Patrick Loiseau**, Paulo Gonçalves, Stéphane Girard, Florence Forbes, and Pascale Vicat-Blanc Primet. Maximum likelihood estimation of the flow size distribution tail index from sampled packet data. In *Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems (ACM SIGMETRICS / Performance)*, June 2009.

- [39] **Patrick Loiseau**, Paulo Gonçalves, Romaric Guillier, Matthieu Imbert, Yuetsu Kodama, and Pascale Vicat-Blanc Primet. Metroflux: A high performance system for analyzing flow at very fine-grain. In *Proceedings of the 5th International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom)*, April 2009.

Articles in refereed workshops

- [40] Hadrien Hours, Ernst Biersack, and **Patrick Loiseau**. Causal study of network performance. In *Proceedings of the 17ème Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications (AlgoTel)*, June 2014.
- [41] Hadrien Hours, Ernst Biersack, and **Patrick Loiseau**. A causal study of an emulated network. In *10ème Atelier en Evaluation de Performances (AEP10)*, June 2014.
- [42] **Patrick Loiseau**, Paulo Gonçalves, Romaric Guillier, Matthieu Imbert, Oana Goga, Yuetsu Kodama, and Pascale Vicat-Blanc Primet. *Metroflux*: a high performance system for very fine-grain flow analysis. In *Grid'5000 Spring School*, April 2009.
- [43] **Patrick Loiseau**, Paulo Gonçalves, and Pascale Vicat-Blanc Primet. How TCP can kill self-similarity. In *Euro-NF workshop: Traffic Engineering and Dependability in the Network of the Future*, September 2008.
- [44] **Patrick Loiseau**, Paulo Gonçalves, Yuetsu Kodama, and Pascale Vicat-Blanc Primet. Metroflux: A fully operational high speed metrology platform. In *Euro-NF workshop: New trends in modeling, quantitative methods and measurements, in cooperation with NET-COOP*, September 2008.
- [45] **Patrick Loiseau**, Paulo Gonçalves, Guillaume Dewaele, Pierre Borgnat, Patrice Abry, and Pascale Vicat-Blanc Primet. Vérification du lien entre auto-similarité et distributions à queues lourdes sur un dispositif grande échelle. In *9ème Atelier en Evaluation de Performances (AEP9)*, June 2008.
- [46] **Patrick Loiseau**, Paulo Gonçalves, and Pascale Vicat-Blanc Primet. A comparative study of different heavy tail index estimators of the flow size from sampled data. In *Proceedings of the MetroGrid Workshop, within the framework of GridNets International Conference*, October 2007.

Software Demonstrations

- [47] **Patrick Loiseau**, Romaric Guillier, Oana Goga, Matthieu Imbert, Paulo Gonçalves, and Pascale Vicat-Blanc Primet. Automated traffic measurements and analysis in Grid5000, June 2009. ACM SIGMETRICS / Performance demonstration contest (**Best Student Demonstration Award**).

Invited talks

- | | |
|---|---------------|
| ENS Lyon (SIESTE seminar), Lyon, France
<i>Learning from strategic data: a game-theoretic perspective</i> | October 2016 |
| MPI-SWS, Saarbrücken, Germany
<i>Classification from strategic data: a game-theoretic perspective</i> | April 2016 |
| 11ème Atelier en Evaluation de Performances (keynote), Toulouse, France
<i>Strategic resource allocation in adversarial environments</i> | March 2016 |
| Harvard University (EconCS seminar), Cambridge, MA, USA
<i>Classification from strategic data: a game-theoretic perspective</i> | November 2015 |
| MIT (Special Henry L. Pierce laboratory seminar), Cambridge, MA, USA
<i>Classification from strategic data: a game-theoretic perspective</i> | November 2015 |

- Northeastern University (ECE department seminar), Boston, MA, USA November 2015
Classification from strategic data: a game-theoretic perspective
- MIT Media Lab (Data Transparency Lab conference), Cambridge, MA, USA November 2015
Bringing Transparency to Targeted Advertising
- LRI, Université Paris-Sud (Séminaire d'algorithmique et de complexité du plateau de Saclay), Saclay, France October 2015
Learning to classify from strategic data
- UCLA, IPAM Graduate Summer School: Games and Contracts for Cyber-Physical Security (invited lecturer), Los Angeles, CA, USA July 2015
Learning with Strategic Agents: From Adversarial Learning to Game-Theoretic Statistics
- Inria Grenoble (In'tech seminar), Grenoble, France June 2015
On the impact of game theory in security
- ACM SIGMETRICS (invited tutorial), Portland, OR, USA June 2015
Learning with Strategic Agents: From Adversarial Learning to Game-Theoretic Statistics
- LINCS (LINCS seminar), Paris, France March 2015
Game-theoretic statistics: Learning from data generated by strategic agents
- Institut Henri Poincaré (Paris game theory seminar), Paris, France March 2015
Game-theoretic statistics: Learning from data generated by strategic agents
- Data transparency lab (DTL) kickoff workshop, Telefonica, Barcelona, Spain November 2014
Game theory and statistics for data transparency: 3 directions
- AlgoGT, Saint Nizier du Moucherotte, France July 2013
Classification games
- Campus SophiaTech (SophiaTech networks seminar), Sophia-Antipolis, France April 2013
A Robust Incentive Mechanism for Congestion Management
- Mines ParisTech (Séminaire du CMA), Sophia-Antipolis, France March 2013
Incentive Mechanisms for Decongestion: Fixed-Budget Rebate versus Time-of-Day Pricing
- UC Berkeley (TRUST seminar), Berkeley, USA August 2012
Incentive mechanisms for congestion management
- RESCOM summer school (guest lecture), Vittel, France June 2012
Game theory for network security and privacy
- Supélec, Gif-sur-Yvette, France February 2012
Large games for Internet congestion management
- INRIA Paris-Rocquencourt (RAP seminar), Le Chesnay, France February 2012
Large games for Internet congestion management
- UCLA (EE department), Los Angeles, CA, USA October 2011
Raffle-based Incentive Schemes for Congestion Management
- Caltech (RSRG Seminars), Pasadena, CA, USA October 2011
Raffle-based Incentive Schemes for Congestion Management
- Orange Labs (France Telecom), Sophia-Antipolis, France March 2011
TCP traffic modeling using an almost-sure large-deviations result

- University of Nice, Laboratoire J.A. Dieudonné (Séminaire de Probabilités et Statistiques), Nice, France
Principe de grandes déviations presque-sur et applications March 2011
- Alcatel-Lucent Bell Laboratories (Mathematics of Networks and Communications Research Department),
Murray Hill, NJ, USA March 2011
Almost-sure large deviations and application to TCP traffic
- University of Waterloo (Department of Electrical and Computer Engineering invited seminar), Waterloo,
Canada
Large deviations and application to fine TCP modeling October 2010
- UC Berkeley (Networking, Communications and DSP seminars), Berkeley, CA, USA September 2010
Large deviations and application to fine TCP modeling
- Caltech (RSRG Seminars), Pasadena, CA, USA September 2010
Large deviations and application to fine TCP modeling
- Politecnico di Torino (Telecommunication Network Group), Torino, Italy May 2010
Heavy-tails and correlations in network traffic
- INRIA Paris-Rocquencourt (RAP seminar), Le Chesnay, France November 2009
Large deviations and application to TCP performance

REFERENCES

- [Abernethy et al., 2011] Abernethy, J., Bartlett, P. L., and Hazan, E. (2011). Blackwell approachability and no-regret learning are equivalent. In *Proceedings of COLT*.
- [Abernethy et al., 2015] Abernethy, J., Chen, Y., Ho, C.-J., and Waggoner, B. (2015). Low-cost learning via active data procurement. In *Proceedings of EC*, pages 619–636.
- [Abernethy et al., 2008] Abernethy, J., Warmuth, M. K., and Yellin, J. (2008). Optimal strategies from random walks. In *Proceedings of COLT*.
- [Acquisti et al., 2011] Acquisti, A., Gross, R., and Stutzman, F. (2011). Faces of facebook: Privacy in the age of augmented reality. In *BlackHat*.
- [Acquisti and Grossklags, 2012] Acquisti, A. and Grossklags, J. (2012). An online survey experiment on ambiguity and privacy. *Communications & Strategies*, 88(4):19–39.
- [Agrawal and Srikant, 2000] Agrawal, R. and Srikant, R. (2000). Privacy-preserving data mining. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 439–450.
- [Aitken, 1935] Aitken, A. C. (1935). On least squares and linear combinations of observations. *Proceedings of the Royal Society of Edinburgh*.
- [Albadi and El-Saadany, 2008] Albadi, M. and El-Saadany, E. (2008). A summary of demand response in electricity markets. *Electric Power Systems Research*, 78(11):1989 – 1996.
- [Alpcan and Başar, 2003] Alpcan, T. and Başar, T. (2003). A game theoretic approach to decision and analysis in network intrusion detection. In *Proceedings of IEEE CDC*, pages 2595–2600.
- [Alpcan and Başar, 2010] Alpcan, T. and Başar, T. (2010). *Network Security: A Decision and Game-Theoretic Approach*. Cambridge University Press.
- [Arce et al., 2012] Arce, D. G., Kovenock, D., and Roberson, B. (2012). Weakest-link attacker-defender games with multiple attack technologies. *Naval Research Logistics (NRL)*, 59(6):457–469.
- [Atallah et al., 1999] Atallah, M., Bertino, E., Elmagarmid, A., Ibrahim, M., and Verykios, V. (1999). Disclosure limitation of sensitive rules. In *Proceedings of Workshop on Knowledge and Data Engineering Exchange (KDEX'99)*, pages 45–52.
- [Auer et al., 2002] Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2002). The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77.
- [Aumann and Maschler, 1995] Aumann, R. J. and Maschler, M. (1995). *Repeated Games with Incomplete Information*. MIT Press.
- [Azar et al., 2015] Azar, Y., Kalp-Shaltiel, I., Lucier, B., Menache, I., Naor, J. S., and Yaniv, J. (2015). Truthful online scheduling with commitments. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation (EC)*, pages 715–732.
- [Backes et al., 2016] Backes, M., Berrang, P., Goga, O., Gummadi, K., and Manoharan, P. (2016). On profile linkability despite anonymity in social media systems. In *WPES*.

- [Backstrom et al., 2010] Backstrom, L., Sun, E., and Marlow, C. (2010). Find me if you can: improving geographical prediction with social and spatial proximity. In *WWW*.
- [Barreno et al., 2010] Barreno, M., Nelson, B., Joseph, A. D., and Tygar, J. D. (2010). The security of machine learning. *Machine Learning*, 81(2):121–148.
- [Bartlett et al., 2015] Bartlett, P. L., Koolen, W. M., Malek, A., Takimoto, E., and Warmuth, M. K. (2015). Minimax fixed-design linear regression. In *Proceedings of The 28th Annual Conference on Learning Theory (COLT)*, pages 226–239.
- [Basar and Srikant, 2002] Basar, T. and Srikant, R. (2002). Revenue-maximizing pricing and capacity expansion in a many-users regime. In *In proceedings of IEEE INFOCOM*, pages 294–301.
- [Baye et al., 1996] Baye, M. R., Kovenock, D., and De Vries, C. G. (1996). The all-pay auction with complete information. *Economic Theory*, 8(2):291–305.
- [Biega et al., 2016] Biega, J. A., Gummadi, K. P., Mele, I., Milchevski, D., Tryfonopoulos, C., and Weikum, G. (2016). R-susceptibility: An ir-centric approach to assessing privacy risks for users in online communities. In *SIGIR*.
- [Blackwell, 1956a] Blackwell, D. (1956a). An analog of the minimax theorem for vector payoffs. *Pacific J. Math.*, 6(1):1–8.
- [Blackwell, 1956b] Blackwell, D. (1956b). Controlled random walks. In De Groot, J. and Gerretsen, J., editors, *Proceedings of the International Congress of Mathematicians 1954*, volume 3, pages 336–338.
- [Blum and Mansour, 2005] Blum, A. and Mansour, Y. (2005). From external to internal regret. In *Proceedings of COLT*, pages 621–636.
- [Bodik et al., 2014] Bodik, P., Menache, I., Naor, J. S., and Yaniv, J. (2014). Brief announcement: Deadline-aware scheduling of big-data processing jobs. In *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 211–213.
- [Borel, 1921] Borel, E. (1921). La théorie du jeu et les équations intégrales à noyau symétrique. *Comptes Rendus de l'Académie des Sciences*, 173(1304-1308):58.
- [Borel and Ville, 1938] Borel, E. and Ville, J. (1938). *Applications de la théorie des probabilités aux jeux de hasard*. J. Gabay.
- [Brückner et al., 2012] Brückner, M., Kanzow, C., and Scheffer, T. (2012). Static prediction games for adversarial learning problems. *Journal of Machine Learning Research*, 13:2617–2654.
- [Brückner and Scheffer, 2011] Brückner, M. and Scheffer, T. (2011). Stackelberg games for adversarial prediction problems. In *Proceedings of ACM SIGKDD*, pages 547–555.
- [Bubeck and Cesa-Bianchi, 2012] Bubeck, S. and Cesa-Bianchi, N. (2012). Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122.
- [Cai et al., 2015] Cai, Y., Daskalakis, C., and Papadimitriou, C. H. (2015). Optimum statistical estimation with strategic data sources. *JMLR W&CP (Proceedings of COLT 2015)*, 40:40.1–40.40.
- [Caragiannis et al., 2016] Caragiannis, I., Procaccia, A. D., and Shah, N. (2016). Truthful univariate estimators. In *Proceedings of the 33rd International Conference on Machine Learning (ICML '16)*.
- [Caruana and Li, 2012] Caruana, G. and Li, M. (2012). A survey of emerging approaches to spam filtering. *ACM Computing Surveys*, 44(2):9:1–9:27.
- [Cesa-Bianchi et al., 1997] Cesa-Bianchi, N., Freund, Y., Haussler, D., Helmbold, D. P., Schapire, R. E., and Warmuth, M. K. (1997). How to use expert advice. *J. ACM*, 44(3):427–485.
- [Cesa-Bianchi and Lugosi, 2003] Cesa-Bianchi, N. and Lugosi, G. (2003). Potential-based algorithms in on-line prediction and game theory. *Machine Learning*, 51(3):239–261.

- [Chambers and Lambert, 2014] Chambers, C. P. and Lambert, N. S. (2014). Dynamically eliciting unobservable information. In *Proceedings of EC*, pages 987–988.
- [Chang et al., 2010] Chang, J., Rosenn, I., Backstrom, L., and Marlow, C. (2010). epluribus: Ethnicity on social networks. In *ICWSM*.
- [Chen and Leneutre, 2009] Chen, L. and Leneutre, J. (2009). A game theoretical framework on intrusion detection in heterogeneous networks. *IEEE Transactions on Information Forensics and Security*, 4(2):165–178.
- [Chen et al., 2012] Chen, T., Kaafar, M. A., Friedman, A., and Boreli, R. (2012). Is more always merrier?: A deep dive into online social footprints. In *WOSN*.
- [Chernov and Zhdanov, 2010] Chernov, A. and Zhdanov, F. (2010). Prediction with expert advice under discounted loss. In *Algorithmic Learning Theory*, pages 255–269. Springer.
- [Chorppath and Alpcan, 2013] Chorppath, A. K. and Alpcan, T. (2013). Trading privacy with incentives in mobile commerce: A game theoretic approach. *Pervasive and Mobile Computing*, 9(4):598–612.
- [Cover, 1966] Cover, T. M. (1966). Behavior of sequential predictors of binary sequences. Technical report, DTIC Document.
- [Cummings et al., 2015] Cummings, R., Ioannidis, S., and Ligett, K. (2015). Truthful linear regression. In *Proceedings of the 28th Annual Conference on Learning Theory (COLT 2015)*, volume 40, pages 1–36.
- [Dalvi et al., 2004] Dalvi, N., Domingos, P., Mausam, Sanghai, S., and Verma, D. (2004). Adversarial classification. In *Proceedings of ACM KDD*, pages 99–108.
- [Dandekar et al., 2012] Dandekar, P., Fawaz, N., and Ioannidis, S. (2012). Privacy auctions for recommender systems. In *Proceedings of WINE*.
- [Dasgupta and Ghosh, 2013] Dasgupta, A. and Ghosh, A. (2013). Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd International Conference on World Wide Web (WWW '13)*, pages 319–330.
- [Davison and Hinkley, 1997] Davison, A. C. and Hinkley, D. (1997). *Bootstrap Methods and their Application*. Cambridge University Press.
- [Dekel et al., 2008] Dekel, O., Fischer, F., and Procaccia, A. D. (2008). Incentive compatible regression learning. In *Proceedings of SODA*, pages 884–893.
- [Domingo-Ferrer, 2008] Domingo-Ferrer, J. (2008). A survey of inference control methods for privacy-preserving data mining. In *Privacy-preserving data mining*, pages 53–80. Springer.
- [Duchi et al., 2013] Duchi, J., Jordan, M., and Wainwright, M. (2013). Local privacy and statistical minimax rates. In *Proceedings of the 54th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 429–438.
- [Duncan and Mukherjee, 2000] Duncan, G. T. and Mukherjee, S. (2000). Optimal disclosure limitation strategy in statistical databases: Detering tracker attacks through additive noise. *Journal of the American Statistical Association*, 95(451):720–729.
- [Dwork, 2006] Dwork, C. (2006). Differential privacy. In *Proceedings of ICALP*, pages 1–12.
- [Dwork and Roth, 2014] Dwork, C. and Roth, A. (2014). The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3–4):211–407.
- [Ferguson et al., 2012] Ferguson, A. D., Bodik, P., Kandula, S., Boutin, E., and Fonseca, R. (2012). Jockey: Guaranteed job latency in data parallel clusters. In *Proceedings of the 7th ACM European Conference on Computer Systems (EuroSys)*, pages 99–112.
- [Foster and Vohra, 1997] Foster, D. P. and Vohra, R. V. (1997). Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1–2):40–55.

- [Freund and Schapire, 1999] Freund, Y. and Schapire, R. E. (1999). Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1–2):79–103.
- [Frongillo et al., 2015] Frongillo, R. M., Chen, Y., and Kash, I. A. (2015). Elicitation for aggregation. In *Proceedings of AAAI*.
- [Fudenberg and Tirole, 1991] Fudenberg, D. and Tirole, J. (1991). *Game Theory*. MIT Press.
- [Ghosh and Roth, 2011] Ghosh, A. and Roth, A. (2011). Selling privacy at auction. In *Proceedings of ACM EC*, pages 199–208.
- [Globerson and Roweis, 2006] Globerson, A. and Roweis, S. (2006). Nightmare at test time: Robust learning by feature deletion. In *Proceedings of ICML*.
- [Gneiting and Raftery, 2007] Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378.
- [Gravin et al., 2014] Gravin, N., Peres, Y., and Sivan, B. (2014). Towards optimal algorithms for prediction with expert advice. *arXiv preprint arXiv:1409.3040*.
- [Gretton et al., 2007] Gretton, A., Fukumizu, K., Teo, C. H., Song, L., Schölkopf, B., and Smola, A. J. (2007). A kernel statistical test of independence. In *NIPS*.
- [Gross and Wagner, 1950] Gross, O. and Wagner, R. (1950). A continuous Colonel Blotto game. *Rand*.
- [Guzella and Caminhas, 2009] Guzella, T. S. and Caminhas, W. M. (2009). A review of machine learning approaches to spam filtering. *Expert Systems with Applications*, 36(7):10206–10222.
- [Ha et al., 2012] Ha, S., Sen, S., Joe-Wong, C., Im, Y., and Chiang, M. (2012). TUBE: Time Dependent Pricing for Mobile Data. In *In proceedings of ACM SIGCOMM*.
- [Hannan, 1957] Hannan, J. (1957). Approximation to Bayes risk in repeated plays. In Drescher, M., Tucker, A. W., and Wolfe, P., editors, *Contributions to the Theory of Games*, volume 3, pages 97–139. Princeton University Press.
- [Hastie et al., 2009] Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer, second edition.
- [Henderson et al., 2001] Henderson, T., Crowcroft, J., and Bhatti, S. (2001). Congestion pricing: paying your way in communication networks. *IEEE Internet Computing*, 5(5):85–89.
- [Hillman and Riley, 1989] Hillman, A. L. and Riley, J. G. (1989). Politically contestable rents and transfers. *Economics & Politics*, 1(1):17–39.
- [Honig and Steiglitz, 1995] Honig, M. L. and Steiglitz, K. (1995). Usage-based pricing of packet data generated by a heterogeneous user population. In *Proc. of INFOCOM '95*, pages 867–874.
- [Hortala-Vallve and Llorente-Saguer, 2012] Hortala-Vallve, R. and Llorente-Saguer, A. (2012). Pure strategy nash equilibria in non-zero sum colonel blotto games. *International Journal of Game Theory*, 41(2):331–343.
- [Huang et al., 2011] Huang, L., Joseph, A. D., Nelson, B., Rubinstein, B. I., and Tygar, J. D. (2011). Adversarial machine learning. In *Proceedings of ACM AISec*, pages 43–58.
- [Huberman et al., 2005] Huberman, B. A., Adar, E., and Fine, L. R. (2005). Valuating privacy. *IEEE Security & Privacy*, 3(5):22–25.
- [Ishakian et al., 2012] Ishakian, V., Sweha, R., Bestavros, A., and Appavoo, J. (2012). Cloudpack* exploiting workload flexibility through rational pricing. In *Proceedings of the 13th International Middleware Conference (Middleware)*, pages 374–393.
- [Jain et al., 2011] Jain, N., Menache, I., Naor, J., and Yaniv, J. (2011). A truthful mechanism for value-based scheduling in cloud computing. In *Proceedings of the 4th International Conference on Algorithmic Game Theory (SAGT)*, pages 178–189.

- [Jain et al., 2012] Jain, N., Menache, I., Naor, J., and Yaniv, J. (2012). Near-optimal scheduling mechanisms for deadline-sensitive jobs in large computing clusters. In *Proceedings of the Twenty-fourth Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 255–266.
- [Jalapati et al., 2012] Jalapati, V., Ballani, H., Costa, P., Karagiannis, T., and Rowstron, A. (2012). Bridging the tenant-provider gap in cloud services. In *Proceedings of the Third ACM Symposium on Cloud Computing (SoCC)*, pages 10:1–10:14.
- [Jiang et al., 2008] Jiang, L., Parekh, S., and Walrand, J. (2008). Time-dependent network pricing and bandwidth trading. In *Proc. of IEEE NOMS International Workshop on Bandwidth on Demand*.
- [Joe-Wong et al., 2011] Joe-Wong, C., Ha, S., and Chiang, M. (2011). Time-dependent broadband pricing: Feasibility and benefits. In *Proceedings the 31st International Conference on Distributed Computing Systems (ICDCS)*, pages 288–298.
- [Joe-Wong et al., 2012] Joe-Wong, C., Sen, S., Ha, S., and Chiang, M. (2012). Optimized day-ahead pricing for smart grids with device-specific scheduling flexibility. *IEEE Journal on Selected Areas in Communications*, 30(6):1075–1085.
- [Kairouz et al., 2016] Kairouz, P., Oh, S., and Viswanath, P. (2016). Extremal mechanisms for local differential privacy. *Journal of Machine Learning Research*, 17(17):1–51.
- [Kantarcioglu et al., 2011] Kantarcioglu, M., Xi, B., and Clifton, C. (2011). Classifier evaluation and attribute selection against active adversaries. *Data Mining and Knowledge Discovery*, 22(1):291–335.
- [Kifer et al., 2012] Kifer, D., Smith, A., and Thakurta, A. (2012). Private convex empirical risk minimization and high-dimensional regression. *JMLR W&CP (Proceedings of COLT 2012)*, 23:25.1–25.40.
- [Kleinberg et al., 2001] Kleinberg, J., Papadimitriou, C. H., and Raghavan, P. (2001). On the value of private information. In *Proceedings of TARK*, pages 249–257.
- [Koolen et al., 2014] Koolen, W. M., Malek, A., and Bartlett, P. L. (2014). Efficient minimax strategies for square loss games. In *Advances in Neural Information Processing Systems*, pages 3230–3238.
- [Koolen et al., 2015] Koolen, W. M., Malek, A., Bartlett, P. L., and Abbasi, Y. (2015). Minimax time series prediction. In *Advances in Neural Information Processing Systems*, pages 2548–2556.
- [Kovenock et al., 2010] Kovenock, D., Mauboussin, M. J., and Roberson, B. (2010). Asymmetric conflicts with endogenous dimensionality. *Korean Economic Review*, 26:287–305.
- [Kovenock and Roberson, 2012a] Kovenock, D. and Roberson, B. (2012a). Coalitional Colonel Blotto games with application to the economics of alliances. *Journal of Public Economic Theory*, 14(4):653–676.
- [Kovenock and Roberson, 2012b] Kovenock, D. and Roberson, B. (2012b). Coalitional Colonel Blotto games with application to the economics of alliances. *Journal of Public Economic Theory*, 14(4):653–676.
- [Kvasov, 2007] Kvasov, D. (2007). Contests with limited resources. *Journal of Economic Theory*, 136(1):738–748.
- [Labitzke et al., 2011] Labitzke, S., Taranu, I., and Hartenstein, H. (2011). What your friends tell others about you: Low cost linkability of social network profiles. In *SNA-KDD*.
- [Lantz et al., 2010] Lantz, B., Heller, B., and McKeown, N. (2010). A network in a laptop: Rapid prototyping for software-defined networks. In *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks, Hotnets-IX*, pages 19:1–19:6, New York, NY, USA. ACM.
- [Laskov and Lippmann, 2010] Laskov, P. and Lippmann, R. (2010). Machine learning in adversarial environments. *Machine Learning*, 81(2):115–119.
- [Lehrer, 2003] Lehrer, E. (2003). Approachability in infinite dimensional spaces. *International Journal of Game Theory*, 31(2):253–268.

- [Li and Vorobeychik, 2015] Li, B. and Vorobeychik, Y. (2015). Scalable optimization of randomized operational decisions in adversarial classification settings. In *Proceedings of AISTATS*.
- [Li et al., 2014] Li, P., Guan, X., Wu, J., and Wang, D. (2014). Pricing strategy for device-level demand response in a microgrid. In *Control Conference (CCC), 2014 33rd Chinese*, pages 7579–7584.
- [Ligett and Roth, 2012] Ligett, K. and Roth, A. (2012). Take it or Leave it: Running a Survey when Privacy Comes at a Cost. In *WINE*, pages 378–391.
- [Littlestone and Warmuth, 1994] Littlestone, N. and Warmuth, M. (1994). The weighted majority algorithm. *Information and Computation*, 108(2):212–261.
- [Liu et al., 2013] Liu, J., Zhang, F., Song, X., Song, Y.-I., Lin, C.-Y., and Hon, H.-W. (2013). What’s in a name?: An unsupervised approach to link users across communities. In *WSDM*.
- [Londoño et al., 2010] Londoño, J., Bestavros, A., and Laoutaris, N. (2010). Trade & cap: a customer-managed, market-based system for trading bandwidth allowances at a shared link. In *Proc. of ACM NetEcon Workshop*, pages 7:1–7:7.
- [Lowd and Meek, 2005] Lowd, D. and Meek, C. (2005). Adversarial learning. In *Proceedings of ACM KDD*, pages 641–647.
- [Lucier et al., 2013] Lucier, B., Menache, I., Naor, J. S., and Yaniv, J. (2013). Efficient online scheduling for deadline-sensitive jobs: Extended abstract. In *Proceedings of the Twenty-fifth Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 305–314.
- [Luo and Schapire, 2013] Luo, H. and Schapire, R. E. (2013). Towards minimax online learning with unknown time horizon. *arXiv preprint arXiv:1307.8187*.
- [Luo et al., 2015] Luo, Y., Shah, N. B., Huang, J., and Walrand, J. (2015). Parametric prediction from parametric agents. In *Proceedings of the 10th Workshop on the Economics of Networks, Systems and Computation (NetEcon '15)*, pages 57–57.
- [Malhotra et al., 2012] Malhotra, A., Totti, L., Meira, W., Kumaraguru, P., and Almeida, V. (2012). Studying user footprints in different online social networks. In *CSOSN*.
- [Manshaei et al., 2013] Manshaei, M., Zhu, Q., Alpcan, T., Basar, T., and Hubaux, J.-P. (2013). Game theory meets network security and privacy. *ACM Computing Surveys*, 45(2).
- [Marbach, 2004] Marbach, P. (2004). Analysis of a static pricing scheme for priority services. *IEEE/ACM Transactions on Networking*, 12(2):312–325.
- [Masucci and Silva, 2014] Masucci, A. M. and Silva, A. (2014). Strategic resource allocation for competitive influence in social networks. In *52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 951–958.
- [Meir et al., 2012] Meir, R., Procaccia, A. D., and Rosenschein, J. S. (2012). Algorithms for strategyproof classification. *Artificial Intelligence*, 186:123–156.
- [Menache et al., 2014] Menache, I., Shamir, O., and Jain, N. (2014). On-demand, spot, or both: Dynamic resource allocation for executing batch jobs in the cloud. In *11th International Conference on Autonomic Computing (ICAC)*, pages 177–187.
- [Mendelson and Whang, 1990] Mendelson, H. and Whang, S. (1990). Optimal incentive-compatible priority pricing for the M/M/1 queue. *Operations Research*, 38:870–883.
- [Mislove et al., 2010] Mislove, A., Viswanath, B., Gummadi, K. P., and Druschel, P. (2010). You are who you know: inferring user profiles in online social networks. In *WSDM*.
- [Motoyama and Varghese, 2009] Motoyama, M. and Varghese, G. (2009). I seek you: searching and matching individuals in social networks. In *WIDM*.

- [Nagarajan et al., 2013] Nagarajan, V., Wolf, J., Balmin, A., and Hildrum, K. (2013). Flowflex: Malleable scheduling for flows of mapreduce jobs. In Eyers, D. and Schwan, K., editors, *Proceedings of the 14th International Middleware Conference (Middleware)*, pages 103–122.
- [Nelson et al., 2009] Nelson, B., Barreno, M., Chi, F. J., Joseph, A. D., Rubinstein, B. I. P., Saini, U., Sutton, C., Tygar, J. D., and Xia, K. (2009). Misleading learners: Co-opting your spam filter. In Yu, P. S. and Tsai, J. J. P., editors, *Machine Learning in Cyber Trust: Security, Privacy, and Reliability*. Springer.
- [Nelson et al., 2010] Nelson, B., Rubinstein, B. I. P., Huang, L., Joseph, A. D., Lau, S., Lee, S., Rao, S., Tran, A., and Tygar, J. D. (2010). Near optimal evasion of convex-inducing classifiers. In *Proceedings of AISTATS*.
- [Nix and Kantarciouglu, 2012] Nix, R. and Kantarciouglu, M. (2012). Incentive compatible privacy-preserving distributed classification. *IEEE Transactions on Dependable and Secure Computing*, 9(4):451–462.
- [Northern and Nelson, 2011] Northern, C. T. and Nelson, M. L. (2011). An unsupervised approach to discovering and disambiguating social media profiles. In *MDSW*.
- [Odlyzko, 1999] Odlyzko, A. (1999). Paris metro pricing for the internet. In *Proc. of EC '99*, pages 140–147.
- [Oliveira and Zaiane, 2003] Oliveira, S. R. and Zaiane, O. R. (2003). Privacy preserving clustering by data transformation. In *Proceedings of SBBD*, pages 304–318.
- [Padhye et al., 1998] Padhye, J., Firoiu, V., Towsley, D., and Kurose, J. (1998). Modeling tcp throughput: a simple model and its empirical validation. *SIGCOMM Comput. Commun. Rev.*, 28(4):303–314.
- [Paridhi Jain and Joshi, 2013] Paridhi Jain, P. K. and Joshi, A. (2013). @i seek 'fb.me': Identifying users across multiple online social networks. In *WoLE*.
- [Paschalidis and Tsitsiklis, 2000] Paschalidis, I. C. and Tsitsiklis, J. N. (2000). Congestion-dependent pricing of network services. *IEEE/ACM Transactions on Networking*, 8(2):171–184.
- [Pearl, 2009] Pearl, J. (2009). *Causality: Models, Reasoning and Inference*. Cambridge University Press, New York, NY, USA.
- [Peled et al., 2013] Peled, O., Fire, M., Rokach, L., and Elovici, Y. (2013). Entity matching in online social networks. In *SocialCom*.
- [Perchet, 2014] Perchet, V. (2014). Approachability, regret and calibration: Implications and equivalences. *Journal of Dynamics and Games*, 1(2):181–254.
- [Perito et al., 2011] Perito, D., Castelluccia, C., Ali Kâafar, M., and Manils, P. (2011). How unique and traceable are usernames? In *PETS*.
- [Perote and Perote-Pena, 2004] Perote, J. and Perote-Pena, J. (2004). Strategy-proof estimators for simple regression. *Mathematical Social Sciences*, 47(2):153–176.
- [Powell, 2009] Powell, R. (2009). Sequential, nonzero-sum Blotto: Allocating defensive resources prior to attack. *Games and Economic Behavior*, 67(2):611 – 615.
- [Raad et al., 2010] Raad, E., Chbeir, R., and Dipanda, A. (2010). User profile matching in social networks. In *NBiS*.
- [Ricci et al., 2011] Ricci, F., Rokach, L., Shapira, B., and Kantor, P. B. (2011). *Recommender Systems Handbook*. Springer.
- [Riederer et al., 2011] Riederer, C., Erramilli, V., Chaintreau, A., Krishnamurthy, B., and Rodriguez, P. (2011). For sale : your data: by : you. In *Proceedings of HotNets*, pages 13:1–13:6.
- [Rinott et al., 2012] Rinott, Y., Scarsini, M., and Yu, Y. (2012). A colonel Blotto gladiator game. *Math. Oper. Res.*, 37(4):574–590.

- [Roberson, 2006] Roberson, B. (2006). The Colonel Blotto game. *Economic Theory*, 29(1):1–24.
- [Roberson, 2010] Roberson, B. (2010). Allocation games. In Cochran, J. J., Cox, L. A., Keskinocak, P., Kharoufeh, J. P., and Smith, J. C., editors, *Wiley Encyclopedia of Operations Research and Management Science*. John Wiley and Sons, Inc.
- [Rojas-Mora et al., 2011] Rojas-Mora, J., Jiménez, T., and Altman, E. (2011). Simulating flow level bandwidth sharing with pareto distributed file sizes. In *Proceedings of the 5th International ICST Conference on Performance Evaluation Methodologies and Tools (VALUETOOLS)*, pages 265–273.
- [Rosenthal, 1974] Rosenthal, R. W. (1974). Correlated Equilibria in Some Classes of Two-Person Games. In *International Journal of Game Theory*, 3(3):119–128.
- [Roth and Schoenebeck, 2012] Roth, A. and Schoenebeck, G. (2012). Conducting truthful surveys, cheaply. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC '12)*, pages 826–843.
- [Sen et al., 2012] Sen, S., Joe-Wong, C., Ha, S., and Chiang, M. (2012). Incentivizing time-shifting of data: a survey of time-dependent pricing for internet access. *IEEE Communications Magazine*, 50:91–99.
- [Sen et al., 2013] Sen, S., Joe-Wong, C., Ha, S., and Chiang, M. (2013). A survey of smart data pricing: Past proposals, current plans, and future trends. *ACM Comput. Surv.*, 46(2):15:1–15:37.
- [Shen and Başar, 2007] Shen, H. and Başar, T. (2007). Optimal nonlinear pricing for a monopolistic network service provider with complete and incomplete information. *IEEE Journal on Selected Areas in Communications (JSAC), Special Issue: Non-Cooperative Behavior in Networking*, 25(6):1216–1223.
- [Shen and Başar, 2011] Shen, H. and Başar, T. (2011). Pricing under information asymmetry for a large population of users. *Telecommunication Systems*, 47(1-2):123–136.
- [Sommer and Paxson, 2010] Sommer, R. and Paxson, V. (2010). Outside the Closed World: On Using Machine Learning For Network Intrusion Detection. In *Proceedings of IEEE S&P*.
- [Song et al., 2014] Song, L., Xiao, Y., and van der Schaar, M. (2014). Demand side management in smart grids using a repeated game framework. *IEEE Journal on Selected Areas in Communications*, 32(7):1412–1424.
- [Song et al., 2010] Song, Y., Locasto, M. E., Stavrou, A., Keromytis, A. D., and Stolfo, S. J. (2010). On the infeasibility of modeling polymorphic shellcode. *Machine Learning*, 81(2):179–205.
- [Sorin, 2002] Sorin, S. (2002). *A First Course on Zero Sum Repeated Games Paperback*. Springer.
- [Spirtes and Glymour, 1991] Spirtes, P. and Glymour, C. (1991). An Algorithm for Fast Recovery of Sparse Causal Graphs. *Social Science Computer Review*, 9:62–72.
- [Spirtes et al., 2001] Spirtes, P., Glymour, C., and Scheines, R. (2001). *Causation, Prediction, and Search*. The MIT Press, Cambridge, MA, USA, second edition.
- [Stoltz and Lugosi, 2005] Stoltz, G. and Lugosi, G. (2005). Internal regret in on-line portfolio selection. *Machine Learning*, 59(1–2):125–159.
- [Stringhini et al., 2012] Stringhini, G., Egele, M., Zarras, A., Holz, T., Kruegel, C., and Vigna, G. (2012). B@bel: leveraging email delivery for spam mitigation. In *Proceedings of USENIX Security*.
- [Subramanian et al., 2013] Subramanian, S., Ghosh, S., Hosking, J. R. M., Natarajan, R., and Zhang, X. (2013). Dynamic price optimization models for managing time-of-day electricity usage. In *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*, pages 163–168.
- [Tambe, 2011] Tambe, M. (2011). *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.
- [Taylor et al., 2007] Taylor, B., Fingal, D., and Aberdeen, D. (2007). The war against spam: A report from the front line. In *Proceedings of the NIPS Workshop on Machine Learning in Adversarial Environments for Computer Security*.

- [Thomas et al., 2013] Thomas, K., McCoy, D., Grier, C., Kolcz, A., and Paxson, V. (2013). Trafficking fraudulent accounts: The role of the underground market in twitter spam and abuse. In *Proceedings of USENIX Security*, pages 195–210.
- [Traub et al., 1984] Traub, J. F., Yemini, Y., and Woźniakowski, H. (1984). The statistical security of a statistical database. *ACM Transactions on Database Systems (TODS)*, 9(4):672–679.
- [Tsai and Yu, 2009] Tsai, J. J. P. and Yu, P. S., editors (2009). *Machine Learning in Cyber Trust: Security, Privacy, and Reliability*. Springer.
- [Tuffin, 2003] Tuffin, B. (2003). Charging the internet without bandwidth reservation: an overview and bibliography of mathematical approaches. *Journal of Information Science and Engineering*, 19(5):765–786.
- [Vaidya et al., 2006] Vaidya, J., Clifton, C. W., and Zhu, Y. M. (2006). *Privacy Preserving Data Mining*. Springer.
- [Vieille, 1992] Vieille, N. (1992). Weak approachability. *Mathematics of Operations Research*, 17(4):pp. 781–791.
- [Vosecky et al., 2009] Vosecky, J., Hong, D., and Shen, V. (2009). User identification across multiple social networks. In *NDT*.
- [Vovk, 1990] Vovk, V. G. (1990). Aggregating strategies. In *Proceedings of COLT*, pages 371–386.
- [Wang et al., 2014] Wang, G., Wang, T., Zheng, H., and Zhao, B. Y. (2014). Man vs. machine: Practical adversarial detection of malicious crowdsourcing workers. In *Proceedings of USENIX Security*, pages 239–254.
- [Warner, 1965] Warner, S. L. (1965). Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69.
- [Yang et al., 2014] Yang, W., Yu, R., and Nambiar, M. (2014). Quantifying the benefits to consumers for demand response with a statistical elasticity model. *IET Generation, Transmission Distribution*, 8(3):503–515.
- [You et al., 2011] You, G.-w., Hwang, S.-w., Nie, Z., and Wen, J.-R. (2011). Socialsearch: enhancing entity search with social network matching. In *EDBT/ICDT*.
- [Zafarani and Liu, 2009] Zafarani, R. and Liu, H. (2009). Connecting corresponding identities across communities. In *ICWSM*.
- [Zafarani and Liu, 2013] Zafarani, R. and Liu, H. (2013). Connecting users across social media sites: A behavioral-modeling approach. In *KDD*.
- [Zamir, 1992] Zamir, S. (1992). Chapter 5 repeated games of incomplete information: Zero-sum. In Aumann, R. and Hart, S., editors, *Handbook of Game Theory with Economic Applications*, volume 1, pages 109–154. Elsevier.
- [Zhang et al., 2012] Zhang, K., Peters, J., Janzing, D., and Schölkopf, B. (2012). Kernel-based conditional independence test and application in causal discovery. *CoRR*, abs/1202.3775.
- [Zheleva and Getoor, 2009] Zheleva, E. and Getoor, L. (2009). To join or not to join: The illusion of privacy in social networks with mixed public and private user profiles. In *WWW*.
- [Zhou and Kantarcioglu, 2014] Zhou, Y. and Kantarcioglu, M. (2014). Adversarial learning with bayesian hierarchical mixtures of experts. In *Proceedings of SIAM SDM*, pages 929–937.
- [Zhou et al., 2012] Zhou, Y., Kantarcioglu, M., Thuraisingham, B., and Xi, B. (2012). Adversarial support vector machine learning. In *Proceedings of KDD*, pages 1059–1067.

