

Multi-Agent Systems for Computational Economics and Finance

Michael Kampouridis, Panagiotis Kanellopoulos, Maria Kyropoulou, Themistoklis Melissourgos and Alexandros A. Voudouris

School of Computer Science and Electronic Engineering, University of Essex, United Kingdom

E-mails: mkampo@essex.ac.uk, panagiotis.kanellopoulos@essex.ac.uk, maria.kyropoulou@essex.ac.uk, themistoklis.melissourgos@essex.ac.uk, alexandros.voudouris@essex.ac.uk

Abstract. In this article we survey the main research topics of our group at the University of Essex. Our research interests lie at the intersection of theoretical computer science, artificial intelligence, and economic theory. In particular, we focus on the design and analysis of mechanisms for systems involving multiple strategic agents, both from a theoretical and an applied perspective. We present an overview of our group's activities, as well as its members, and then discuss in detail past, present, and future work in multi-agent systems.

Keywords: Algorithmic game theory, Mechanism design, Computational social choice, Machine learning

1. The Group

The School of Computer Science and Electronic Engineering at the University of Essex (UoE) is the home of many researchers actively working on the broad field of Artificial Intelligence. We are members of the Centre for Computational Finance and Economic Agents (CCFEA) in UoE, and our research spans various core areas of multi-agent systems (MAS), focusing primarily at the intersection of theoretical computer science, artificial intelligence, finance and economics. Specifically, our expertise is on algorithmic game theory, computational social choice, machine learning and financial engineering. We have a strong publication track record that demonstrates our close working relationship, as well as our national and international collaborations. Our relevant research is published in highly respected international peer-reviewed journals and conferences, such as AAAI, IJCAI, AAMAS, EC, ICALP, STOC, FOCS, Artificial Intelligence Journal, Journal of Artificial Intelligence Research, Games and Economic Behavior, and more. We also participate as Program Committee members in renowned international conferences and have participated in the organisation of several international conferences and workshops, including the 15th International Symposium on Algorithmic Game Theory (SAGT 2022) and the 2022 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics (CIFER).

We remark that our work is intertwined with industrial applications and we have been involved in four Knowledge Transfer Partnerships funded by Innovate UK (total value exceeding £1M), as well as an ESRC-funded project with Essex County Council. For example, in collaboration with a large healthcare provider, we are studying more efficient ways to elicit preferences, match healthcare workers and clients, and increase the number of shifts that are covered. Our approach relies heavily on game theory, mechanism design, and machine learning. In another case involving a leading law firm, we worked towards developing a pricing tool that predicts the cost of a requested piece of legal work and allocates the most suitable lawyers to it, based on historical data. This is a fascinating and innovative research direction linking once again machine learning and algorithmic game theory.

2. Core Areas of our Research

The study of multi-agent systems (MAS) is a core area of Artificial Intelligence and many different perspectives and methodologies have been adopted, each contributing to the development of MAS. One key area where our group has focused on is the *existence, computation and quality of equilibria* in environments where multiple intelligent agents act selfishly aiming to optimize their personal agenda, which may collide with the interests of others or the system designer's objectives. Another class of problems we have extensively studied falls within the area of *computational social choice*, where agents express preferences over different outcomes and our goal is to make social decisions based on these preferences that satisfy desired properties, such as *utilitarian efficiency* or *fairness criteria*. We have studied game-theoretic and social choice settings in which the main objective is to provide the agents with the necessary incentives to avoid strategic manipulations, an area known as *mechanism design*. Finally, we have considered *financial* applications both through aforementioned frameworks as well as using Machine Learning. In the rest of this section we provide background on these topics, discuss our contribution, and identify interesting directions for future research.

2.1. Existence, Computation, and Quality of Equilibria

Strategic games emerge naturally in environments where rational and intelligent individuals participate. Games model situations where selfish entities aim to maximize their own payoff without caring about that of their competitors. Informally, a game consists of a set of *agents*, each having a set of available *actions*. In such an MAS, a combination of the agents' actions yields an *outcome* which is experienced by all agents; each agent has her own *preference* over the outcomes, often expressed through a valuation or utility function. A probability distribution over an agent's actions is called a *strategy*, while a combination of agent strategies in which no agent would get a better (in expectation) outcome by unilaterally deviating to a different strategy is called a *Nash equilibrium* [1]. Such notions of stability often serve as predictions on the outcomes of the corresponding game, thus, have received a great deal of attention from the community. Various other equilibrium notions have also been considered as solution concepts, depending on the game at hand.

There are three main questions regarding equilibria in games: *Is an equilibrium guaranteed to exist?*, *how fast can we compute one (if it exists)?*, and *how good for the society are the equilibria reached compared to a centrally enforced optimal state?* The last question gives rise to the notions of *price of anarchy* and *price of stability* that quantify how bad the worst (best, respectively) Nash equilibrium is when compared to the state maximizing the aggregate welfare of the society (optimal *social welfare*). We have studied these questions in many different and challenging types of games. In what follows, we present some of our relevant work grouped by similarity of model.

Auctions. In a typical auction environment, there is a set of items for sale and a set of potential buyers. The possible actions of the buyers consist of the potential bids they can place in the auction, and the game is defined by a specific resolution method of the bids, which, given the bids, determines the outcome. Consider, for example, the generalized second price (GSP) auction, which is the primary auction used for monetizing the use of the Internet. In a quite influential paper (see [2]) we studied the space of equilibria under GSP and quantified the efficiency loss that can arise under a wide range of sources of uncertainty. We also focused on the revenue of the GSP auction, and studied a Bayesian setting wherein the valuations of advertisers are drawn independently from a common regular probability distribution [3].

In a combinatorial auction, an outcome consists of an allocation of the items to the agents and a price for each agent's bundle. One of the most prominent equilibrium notions for auctions, where we require per-item prices, is that of a *Walrasian equilibrium*, which occurs when no agent prefers any other bundle for the given set of prices and the market clears. Walrasian equilibria have the additional desirable property of maximizing the social welfare, however they are not known to exist even for very restricted valuations, as we showed in [4]. In the same paper, we presented some special cases for which existence of an equilibrium is guaranteed and provided polynomial time algorithms to compute it.

Furthermore, we have analyzed the efficiency of auctions in settings where the bidders have *budget constraints* that limit the amount of payments they can afford. More specifically, we have considered auctions for the allocation

1 of divisible resources with bidders that have concave valuation functions for fractions of the resources [5], position
2 auctions (including the GSP auction mentioned above) [6], and simultaneous combinatorial auctions with payment
3 rules that are convex combinations of the bids [7].

4 We have also studied a series of questions related to the existence, efficiency and computational complexity of
5 equilibria in two-stage games modelling settings where multiple vendors compete with each other by offering similar
6 products to unit-demand buyers who in turn have different preferences over the vendors [8].

7 *Congestion games.* In a *congestion game* there is a set of resources and each agent can use a subset of them,
8 possibly under some restrictions. Each resource incurs a cost that is a function of the number of agents that use it.
9 Congestion games have been studied extensively due to the fact that they capture many real-life applications and
10 since they admit *pure Nash equilibria*.

11 In [9] we completely characterized the impact of selfishness and greediness in singleton congestion games, where
12 each strategy consists of a single resource. In particular, we presented tight bounds on the price of anarchy, as well as
13 on the competitiveness of the greedy algorithm for the online setting, when the objective is to minimize the total cost
14 of all agents, as well as a tight upper bound on the price of stability of congestion games with linear cost functions.
15 We have also introduced a setting where agents are partially altruistic and partially selfish and we have determined
16 the impact of this behavior on the overall system performance [10]. In an attempt to mitigate the impact of selfish
17 behavior, we studied the question of how much the performance can be improved if agents are forced to pay taxes
18 for using resources, for the setting of atomic congestion games with linear cost functions [11, 12].

19 A well known special case of congestion games is that of task scheduling. In [13] we considered the problem of
20 scheduling tasks on strategic unrelated machines when no payments are allowed, under the objective of minimizing
21 the makespan. We adopted the monitoring paradigm where a machine is bound by her declarations and provided
22 a (non-truthful) randomized algorithm whose pure price of anarchy is considerably better than the approximation
23 ratio that can be achieved by any truthful mechanism.

24 Scheduling tasks to machines is also the focus of [14]. Each task has a particular latency requirement at each
25 machine and may choose either to be assigned to some machine in order to get serviced provided that her latency
26 requirement is met, or not to participate in the assignment at all. From a global perspective, one would aim to
27 maximize the number of tasks that participate in the assignment. However, tasks may behave selfishly aiming to
28 get serviced by some machine where her latency requirement is met with no regard to overall system performance.
29 We modelled this selfish behavior as a strategic game, showed how to compute pure Nash equilibria efficiently, and
30 assessed the impact of selfishness on system performance.

31 *Games on graphs.* *Segregation* has been the subject of multidisciplinary research for more than half a century,
32 with most work focusing on Thomas Schelling's model [15, 16] involving two different types of agents. Each agent
33 occupies a node on a graph and is happy if a least some fraction of her neighbors are of the same type; an unhappy
34 agent either jumps to an unoccupied node or swaps locations with another unhappy agent. In recent work, we
35 considered Schelling's model through the lens of game theory, by assuming that agents choose their next location
36 to maximize the fraction of same-type neighbors. We have studied questions related to the existence, computation,
37 and quality of equilibria for many variants of Schelling games [17–20]. Besides this, we have also studied the
38 complexity of computing welfare-maximizing allocations [21].

39 Regarding the classical *stable matching* problem, we considered its generalization that allows for cardinal prefer-
40 ences (as opposed to ordinal) and fractional matchings (as opposed to integral) [22]. In this cardinal setting, stable
41 fractional matchings can have much larger social welfare than stable integral ones. Our goal was to understand the
42 computational complexity of finding an optimal (i.e., welfare-maximizing) stable fractional matching. We consid-
43 ered both exact and approximate stability notions, and provided simple approximation algorithms with weak welfare
44 guarantees; somewhat surprisingly, achieving better approximations is computationally hard.

45 *Opinion formation* is another important topic, where the objective is to study how the opinion of an individual
46 is affected by her social acquaintances. In this context, we have studied the existence and quality of equilibria in
47 games consisting of agents, each of whom has a personal belief that can be represented by a real number, but may
48 express a different opinion aiming minimize the maximum distance of her opinion from her belief and the most
49 distant opinion expressed by any of her friends [23].

50 In graph-based *hedonic games*, a group of utility maximizing agents have hedonic preferences over the agents'
51 set, and wish to be partitioned into clusters so that they are grouped together with agents they prefer. The agents cor-

1 respond to nodes in a connected graph and their preferences are defined so that shorter graph distance implies higher
2 preference. In [24] we focused on fractional hedonic games and social distance games and we proved improved
3 bounds on the price of stability.

4 *Security games.* To decide the optimal strategy to commit to, the leader in a Stackelberg game can use a variety
5 of learning algorithms that operate by querying the best responses or the payoffs of the follower. Consequently, the
6 follower can potentially deceive the leader by responding as if her payoffs are much different than what they actually
7 are. In [25], we proved that it is always possible for the follower to compute payoffs, and thus deceive the leader, that
8 yield near-optimal utility, for various scenarios about the learning interaction between leader and follower. In [26],
9 we considered a two-stage district-based voting manipulation scenario which leads to a Stackelberg game. First, an
10 attacker attempts to manipulate the election outcome in favor of a preferred candidate by changing the vote counts
11 in some districts. Afterwards, a defender has the ability to demand a recount in a subset of the manipulated districts,
12 thus restoring the vote counts therein to the original, true values.

13 An attacker-defender setting was also studied in [27] where multiple attackers try to inflict damage to nodes on a
14 network, and a single defender tries to protect the nodes by probabilistically patrolling induced connected subgraphs
15 of the network. In that work we showed a general LP-based algorithm for computing Nash equilibria and then
16 focused on special cases of graphs. First we characterized equilibria, and analyzed some of their properties, like
17 the *defense ratio*; that is, the ratio of all attackers over the expected number of the ones caught by the defender in
18 equilibrium. We also computed the (parameterized) value of the *price of defense*, i.e., the worst defense ratio over
19 all possible networks.

20 Investigating a significantly challenging aspect of cryptocurrency, in [28] we focused on the strategic considera-
21 tions of miners participating in bitcoin’s protocol. We formulated and studied the stochastic game underlying such
22 strategic considerations and showed that, when each miner’s computational power is relatively small, their best re-
23 sponse matches the expected behavior of the bitcoin designer. However, when a miner’s computational power is
24 large, she deviates from the expected behavior, and other Nash equilibria arise. This work has served as a starting
25 point for a long line of works on blockchain protocols.

26 *Tractability of equilibria.* As a side quest of their main subject, almost all of the works in this section have studied
27 the question of tractability of equilibria. In other words, to categorize the equilibrium computation problems accord-
28 ing to how fast algorithms they (can) admit under the Turing machine model. [29] solely studied this question on
29 evolutionarily stable strategies, and showed that deciding existence of such strategies is co-NP-hard for a wide range
30 of games in nature.

31 On the bright side, we have managed to show (quasi) polynomial approximation schemes for a significantly broad
32 family of problems in algorithmic game theory and beyond. In [30], we presented a master algorithm which finds
33 approximate solutions to non-linear programs with non-linear constraints. Our algorithm is simple and, in particular,
34 of the brute force kind. Given that the domain is succinctly representable – as is the case for the strategy set in a
35 great number of games – the algorithm is efficient. Importantly, it is able to handle classical optimization problems
36 such as multi-objective non-linear optimization, and for some families of these problems performs at least as good
37 as the currently best algorithms. Nonetheless, it can be considered a unification of older optimization methods under
38 a single approximation scheme with remarkable simplicity.

39 *Future directions.* Even though there have been many results on the computational complexity of the aforemen-
40 tioned problems by now there is still a vast unexplored area for research. The focus of the computer science commu-
41 nity has been almost entirely on the worst-case complexity and only little is known about the average case or similar.
42 This raises the following very important open question: *What is the average or smoothed complexity of algorithms*
43 *for these problems?* Furthermore, as blockchain technology conquers more and more fields, and more protocols are
44 being designed based on, it is important to analyze the theoretical properties of such protocols and the incentives
45 of the miners therein. Finally, a recent series of papers have studied questions about the manipulability of machine
46 learning algorithms for regression or classification, and have shown bounds on the price of anarchy and stability of
47 the induced games. The limits of this research direction have not yet been reached and there are many interesting
48 settings to be considered.

2.2. Computational Social Choice

Another area of MAS research that we have studied in the recent years is that of *computational social choice* [31]. One of the most prominent models within this area is that of *utilitarian voting*, where we have a set of *agents* and a set of *alternatives*. Every agent has a private valuation function, which maps the alternatives to numerical values, indicating how much the agent likes each alternative. Due to cognitive limitations, it is typical for each agent to express their preferences using only *limited* information, typically by ranking the alternatives from the one she considers the best to the one she considers the worst. A *mechanism* (or *voting rule*) takes as input the information provided by the agents about their preferences and outputs a subset of alternatives. The goal is to design mechanisms that make decisions to (approximately) maximize some objective function of the values of the agents for the selected alternatives, such as the *social welfare* (the total value of the agents) or the *egalitarian welfare* (the minimum value over all agents). As mechanisms only have limited information about the valuation functions of the agents the choices they can make are inevitably suboptimal. This loss of efficiency is measured by the *distortion*, that is the worst-case ratio between the maximum possible objective value over all feasible subsets of alternatives and the objective value of the subset chosen by the mechanism.

Our work in utilitarian voting has focused both on improving the state of the art (see our recent survey for an overview [32]), as well as introducing new agendas and open avenues for future research. In single-winner voting (where the goal is to choose a single alternative), we initiated the study of *query-enhanced mechanisms* that take as input the rankings over alternatives provided by the agents, and also have the capability of making a small number of *queries* to (approximately) learn parts of the valuations, which they can then use to make better-informed decisions. For this class of mechanisms, we showed tradeoffs between the necessary amount of information (number of queries) to achieve low distortion [33]. In follow-up work [34, 35], we showed how our methodology can be adapted and extended to capture several other fundamental social choice models beyond utilitarian voting, including the well-known *one-sided matching* problem (in which there is a set of agents that have preferences over a same-sided set of items and the objective is to match each agent to a single distinct item so as to maximize the total value of the agents) and several generalizations of it, such as *constrained resource allocation* (where each agent can be matched to more than one item), *two-sided matching* (in which each agent has preferences over all other agents and the goal is to create pairs of agents with maximum value for one another), and *graph matching* (where the preferences of the agents for other agents are partial and defined according to a directed graph).

In another line of work, we initiated the study of *distributed* voting settings [36–38]. In contrast to centralized voting where there is a single pool of agents and we obtain direct access to the information provided by them, in distributed voting the agents are partitioned into *disjoint groups* and the mechanisms work in two steps: We first decide a *representative alternative* for each group by holding a local election for the agents therein, and then choose one of these representative alternatives as the overall winner. This model captures important applications such as district-based presidential elections, and many other scenarios in which the agents might be grouped together because of logistic or geographic reasons. Recently, we also initiated the study of multiwinner voting (where the goal is to choose some fixed-size subset of alternatives) in settings where the agents have metric preferences over the alternatives (that is, agents and alternatives are represented as points in some — possibly multidimensional — metric space, i.e., the distances between agents and alternatives that define the valuations satisfy the triangle inequality) [39].

Future directions. There are plenty of interesting directions for future work on utilitarian voting. One possible direction would be to combine the models discussed above and investigate the interplay between information and distortion in several single-winner and multiwinner distributed voting, as well as other important settings beyond these, such as *participatory budgeting* (where each alternative is associated with a cost and the goal is to choose a subset of alternatives of total cost within a given budget). While the distortion of ordinal voting rules is well-understood in most important settings by now, the known impossibilities do not rule out that machine learning methods could help in improving the efficiency in an instance-by-instance case. In fact, recent work has shown that it is possible to teach machine learning methods to approximately mimic some voting rules such as scoring rules [40]. While this by itself is not groundbreaking, it indicates that it might be possible to learn voting rules that cannot be described through a simple formula or well-defined instructions. So, machine learning has the potential to allow us to learn voting rules that approximately maximize the social welfare.

2.3. Fair Division

An important class of MAS problems in computer science is that of *fair division*, or in other words, fair resource allocation. While such problems may also fall within the general umbrella of computational social choice, the amount of attention they have received in the last century, and particularly in the last couple of decades, justifies discussing them in isolation. In a fair division problem, there is a set of agents that typically have cardinal preferences over a set of items, where these items can correspond to divisible or indivisible goods or chores. The goal is to compute an allocation of the items to the agents such that every agent considers the allocation to be fair. What fairness means can have multiple interpretations, which is why many different fairness notions have been proposed and studied in the literature. The most compelling and widely adopted ones are those of *envy-freeness*, which demands that no agent envies what another agent has been allocated, and *proportionality*, which demands that every agent obtains at least her proportional share of all items.

Motivated by work on price of anarchy and stability (see Section 2.1), we initiated the study of the *price of fairness* which aims to quantify the impact that the requirement for fairness has on the efficiency of allocations with respect to the social welfare. Formally, it is defined as the worst-case ratio between the maximum social welfare that can be achieved by any allocation and the maximum social welfare over allocations that satisfy particular fairness criteria. In our work, we focused on envy-freeness, proportionality, and another well-known criterion known as equitability, and showed bounds on the price of fairness for divisible and indivisible goods and chores [41]. Our work has received a lot of attention and the notion of price of fairness has been considered in many subsequent papers.

When the items are divisible resources, a setting known as *cake-cutting*, it is known from previous work that there always exist envy-free and proportional allocations for any number of agents. Such allocations are quite easy to compute when the number of agents is rather small and becomes harder when more agents come into play. Recently, we took an experimental approach aiming to understand how well-known cake-cutting algorithms for computing fair allocations perform in practice [42]. We ran many different experiments, where human participants were asked to engage with each other or against automata, and investigated possible manipulations of the algorithms, a question that had only been considered in a small number of previous related papers. Our results indicate that the extent of envy that can arise from strategic behavior is lower in (theoretically) envy-free procedures than proportional and not envy-free ones.

Another problem that recently rose to prominence in fair division of divisible resources is *consensus halving*. In this problem we ask to divide a one-dimensional resource among multiple agents by cutting it into many pieces, and then creating two parts of the pieces such that each agent believes the two parts are of equal value. In [43] we showed that the problem is PPA-complete even for computing a relaxed solution where the agents believe that the difference in value of the two parts is at most a large constant value. Our work has important complexity implications to other major problems in fair division of divisible and indivisible resources, namely discrete ham-sandwich, necklace splitting, pizza sharing, and more. Prior to this, in [44] we studied the problem of finding an exact solution of consensus halving and showed that its complexity differs tremendously from its approximate version. In particular, the problem becomes FIXP-hard and belongs BU, classes that are not even known to be inside FNP.

In contrast to the divisible case, in settings with indivisible items, it is not hard to see that there exist simple instances that do not admit envy-free or proportional allocations (e.g., when there are two agents and one good), and in fact it is generally hard to even decide when such allocations exist, with the exception of a few special cases such as when each agent is to be assigned only a single item [45]. To circumvent such impossibilities, various relaxations of envy-freeness and proportionality have been proposed, including *envy-freeness up to one item* (EF1), which demands that every agent does not envy any agent after the hypothetical removal of any item (from the ones given to them), or the stronger *envy-freeness up to any item* (EFX), which demands that the envy is eliminated by the removal of any item. EF1 allocations always exist and are easy to compute, but the existence of EFX still remains a mystery and has been shown only for a few special cases, such as when there are at most two possible values for the goods [46]. In our recent survey [47], we discussed important results that have been obtained for these relaxations and other important ones in the discrete fair division literature during the last 10 years.

We have also studied other interesting settings with different assumptions on the relationships between agents or on the nature of the possible allocations. For example, we have studied the existence and computation of EF1 and

EFX allocations for settings where the agents are partitioned into groups and the agents of a group can extract value from the items that are allocated to their group [48]. We have also considered the case of random allocations and, in particular, the notion of *interim envy-freeness*, which aims to provide a tradeoff between the very restrictive notion of (deterministic) envy-freeness and the, perhaps too weak, notion of ex-ante envy-freeness [49].

Future directions. As mentioned above, there are only a few cases for which EFX allocations are known to exist and designing an algorithm that always computes an EFX allocation (even in exponential time) has turned out to be a highly non-trivial task. Experiments on randomly generated data indicate that at least one EFX allocation should always exist, but these do not exclude the possibility of instances that do not admit EFX allocations and have such a structure that coming up with them at random or by hand is quite unlikely. As in utilitarian voting, one possible idea to bypass this bottleneck might be to train a machine learning model to try to identify EFX allocations. Such a model has the potential to lead to new insights, for example it may identify particular properties shared by all EFX allocations that could be used in the design of an algorithm, or lead to an impossibility result. We plan to investigate this direction in the near future to tackle this enigmatic question in fair division.

2.4. Approximate Mechanism Design

As discussed in Section 2.1, in many MAS the participating agents usually have incentives to act selfishly aiming to optimize their personal payoff, thus inducing strategic games. However, in many applications, it is important that we manage to control the behavior of the agents as this can give us a better estimation about the expected system usage and efficiency. This can be done by anticipating possible manipulations by the agents, and thus appropriately specifying the rules of the games they will end up playing. In particular, in *mechanism design* the goal is to define the rules such that it is a (weakly) dominant strategy for each agent to act truthfully, that is, no agent has incentive to deviate and report a false input. In some situations this can be coupled together with full efficiency, but usually the truthfulness requirement leads to some loss of societal efficiency, and thus our goal is to design truthful mechanisms that are approximately optimal. In some environments this can be done by using money transfers (*mechanism design with money*), whereas in others the use of money is prohibited (*mechanism design without money*).

In settings where monetary transfers are allowed and there are no budget restrictions, it is always possible to simultaneously achieve truthfulness and full efficiency by deploying the well-known Vickrey-Clarke-Groves (VCG) mechanism. Implementing VCG can be hard in certain applications as it demands to identify a welfare-maximizing outcome and then pair it together with particular payments. In some applications, however, it turns out that VCG can be implemented efficiently. One such example is the problem of scheduling tasks to selfish machines in order to minimize the total execution time. In this context, VCG has a very simple interpretation: greedily and myopically allocate each task to a machine that minimizes its processing time (note that this can lead to a sub-optimal completion time). This approach, however, is clearly suboptimal under the more common objective of minimizing the completion time of the schedule, that we have studied in [50]. We compute the approximation ratio of VCG when the execution times required for each task are independent random variables, identical across machines. Another simple example is that of single-item auctions, where the VCG mechanism reduces to the simple second-price auction that is truthful and maximizes the social welfare (the agent with the highest value for the item obtains it). When the goal is to optimize the revenue in single item auctions rather than welfare, in [51] we showed that no deterministic truthful mechanism can approximate the optimal revenue within a factor better than $63/64 = 98.5\%$ when the participating bidders have correlated valuations; relatively simple optimal mechanisms exist for the case of independent valuations.

In most related literature in mechanism design (and in algorithmic game theory more generally) it is implicitly assumed that the participating agents are perfectly rational in the sense that they always seek to optimize their personal payoffs. But there are also settings with imperfectly rational agents for whom it is not immediate what type of incentives we should provide. *Obviously strategyproof* (OSP) mechanisms have been proposed for such cases to require that at each point during the execution where an agent is queried to communicate information, it should be “obvious” for the agent what strategy to adopt in order to maximize her utility: roughly speaking, the worst possible outcome after selecting her true type is at least as good as the best possible outcome after misreporting her type. In [52] we considered single-minded combinatorial auctions where each bidder is interested in receiving a specific

(known or unknown) bundle of items or a superset of it. We designed novel OSP mechanisms and provided OSP implementations of known truthful mechanisms that are approximately optimal in terms of the social welfare.

There are also many settings where payments are not allowed, but we still want to design truthful mechanisms. The most prominent such setting is that of *facility location*. In the simplest model, there is a set of agents with private positions on the line of real numbers and our goal is to decide where to locate a facility such that the total or the maximum distance of the agents from the facility is minimized, while at the same time truthfully eliciting the agent positions. When there are two facilities or more, the agents might also have different *preferences* over them leading to the so-called *heterogeneous* facility location models. In our work, we have studied a constrained version of the problem with agents that either like or dislike the facilities and the facilities must be located only within a given feasible region of the Euclidean plane [53]. We showed a characterization of the feasible regions for which the optimal solution can be truthfully implemented when the agents can only lie about their preferences or about their locations. In another work, we focused on models with agents that have approval preferences over the facilities (that is, an agent either likes a facility or is indifferent about it) and there are different types of constraints, such as that the agents occupy the nodes of a discrete line graph and the facilities cannot be located at the same node [54], or that we have a limited budget which allows just a subset of the facilities to be built in which case we have to decide which of them to locate and where [55].

In the context of *machine scheduling* without payments, one way to incentivize the agents to act truthfully is that of *monitoring*. We have studied a version of the problem where the machines can be strategic but are bound by their declarations in the sense that if a machine reports that it needs a particular amount of time to process some jobs assigned to it, then it has to spend at least this much processing time [56]. We showed that by using a novel and particularly harsh notion of monitoring, we can design OSP mechanisms that achieve the same approximation guarantee as the best possible truthful mechanism, and that such a monitoring framework is essential. In the context of *fair division*, we have also considered truthful allocations for the fundamental case of two agents and two items [57]. By characterizing the possible truthful mechanisms, we were able to show an impossibility: deterministic truthful allocations do not minimize the envy between agents.

Finally, we have also designed strategyproof mechanisms for *hybrid social choice* settings, where the objective is to choose a welfare-maximizing outcome when there are both agents with non-monetary preferences (corresponding to voters) over the possible outcomes as well as agents with monetary preferences (corresponding to bidders) [58]. Since this problem is a combination of mechanism design with and without monetary transfers, classical solutions like VCG cannot be applied. We focused on the simple but fundamental scenario of one voter and two agents, and provided tight approximation guarantees of the optimal social welfare. We distinguished between mechanisms that use ordinal and cardinal information, as well as between mechanisms that base their decisions on one of the two sides (either the voter or the agents) or both.

Future directions. The field of approximate mechanism design is broad and contains a number of interesting unexplored directions. For example, a number of questions are left open by our works: how well can we approximate the optimal social welfare in a single-minded combinatorial auction with an OSP mechanism? In what other contexts where payments are not allowed can paradigms like monitoring help incentivize truthful behavior? Beyond the problems we have already considered, it would also be interesting to explore the design of strategyproof machine learning algorithms for settings where the samples might be provided by strategic agents.

2.5. Multi-Agent Systems in Finance

The study of financial applications through the lens of algorithmic game theory is a relatively novel research direction that has attracted our interest. *Machine learning* research on finance is more extensive, as it is commonly used in trading strategies and for designing markets, among others.

1 *Financial Network Games.* In [59] we studied financial networks of firms that can strategize over their payments. 1
2 In such networks, nodes correspond to firms, and directed labeled edges correspond to debt contracts between them. 2
3 The existence of cycles in the network indicates that a payment of a firm to one of its lenders might result to some 3
4 incoming payment. So, if a firm cannot fully repay its debt, then the exact (partial) payments it makes to each of its 4
5 creditors can affect the cash inflow back to itself. Note that firms are interested in their financial well-being (utility) 5
6 which is aligned with the amount of incoming payments they receive from the network. We considered a natural set 6
7 of payment strategies called priority-proportional payments, and we investigated the existence and (in)efficiency of 7
8 equilibrium strategies. 8

9 We have also studied a setting where a financial authority offers bailout money to some firm(s) or forgives the 9
10 debts of others in order to maximize liquidity, and examine efficient ways to achieve this goal, together with game- 10
11 theoretic questions on the behavior of the firms [60, 61]. We investigated the approximation ratio provided by 11
12 the greedy bailout policy compared to the optimal one, and we studied the computational hardness of finding the 12
13 optimal debt removal and budget-constrained optimal bailout policy, respectively. From a game-theoretic standpoint, 13
14 we observed that the removal of some incoming debt might be in the best interest of a firm, if that helps one of 14
15 its borrowers remain solvent and avoid costs related to default. Therefore, firms aim to maximize their utility by 15
16 strategically giving up some incoming payments; we studied the existence and quality of pure Nash equilibria, as 16
17 well as the computational complexity of finding them. 17

18 *Exploiting Machine Learning.* MAS allows the simulation of markets which consist of heterogeneous agents, 18
19 with differing risk attitudes and differing expectations to future outcomes, in contrast to traditional assumptions 19
20 of homogeneity and rationality. MAS attempts to explain market behavior, replicate documented features of real- 20
21 world markets, and allows us to gain insight into the likely outcomes of different regulatory policies. The essence of 21
22 MAS in financial markets lies in the notion of autonomous agents whose behavior evolves endogenously leading to 22
23 complex, emergent, system dynamics, which are not predictable from the properties of the individual agents. 23

24 In developing a multi-agent system, a key question is: *How do agents learn and adapt their strategies over time?* 24
25 These ‘learning mechanisms’ can be implemented in many ways including use of different types of machine learning 25
26 algorithms, such as neural networks or evolutionary algorithms, where each agent’s behavior or strategy adapts over 26
27 time in response to environmental feedback, as we have discussed in [62]. 27

28 Machine learning has been very popular in financial applications; we have applied it in areas such as trading [63, 28
29 64] and derivatives [65, 66]. Past work that has particularly focused on MAS includes studying the *microstructure* 29
30 *dynamics* of financial markets. In [67, 68], we investigated the market fraction hypothesis (MFH), which states that 30
31 the fraction of the different types of trading strategies that exist in a financial market changes (swings) over time. 31
32 Our framework used a genetic programming algorithm to evolve the multi-agent trading strategies and a time-variant 32
33 self-organizing map to cluster the different types of agents. Our findings led us to reject the hypothesis, as we found 33
34 that dominant types of trading strategies can remain popular over long time periods with very little swinging present. 34

35 Furthermore, in [69, 70], we studied the plausibility of the dinosaur hypothesis (DH), which states that because 35
36 market’s behavior constantly changes, traders need to continuously co-evolve their strategies with the market in 36
37 order to avoid becoming obsolete or dinosaurs. We again used a multi-agent system that combined evolutionary 37
38 algorithms and neural networks and provided empirical evidence confirming the DH, thus stressing the need of 38
39 agents continuously adapting their strategies to the market environment. 39

40 Another area that we have worked on is applying machine learning to automated negotiation problems, where we 40
41 used a genetic programming algorithm to co-evolve negotiation strategies of agents that have different preference 41
42 criteria, namely price and negotiation speed [71]. In cloud/grid computing environments, any delay in acquiring re- 42
43 sources is considered an overhead, hence the negotiation agents need to adopt strategies that enable them to not only 43
44 optimize resource price, but also to reach early agreements. Our work was the first to apply genetic programming to 44
45 this type of problem and outperformed previous results in the area. 45

46 *Future directions.* Applying algorithmic game theory techniques to financial networks only recently received 46
47 attention from the community, hence the field remains to a big extent unexplored. Existing models can be extended 47
48 by considering more financial instruments (e.g. possibility of claim transfers) and made even more realistic. Com- 48
49 putational results could provide insights, e.g. on the inefficiency of equilibria in a random network. Machine learn- 49
50 ing techniques are very broad and can be applied in this and many other contexts; we propose its use in order to 50
51 overcome intractability or incomplete information issues also in the previous sections. 51

3. Conclusions

To conclude, our group combines research expertise that ranges from theoretical computer science to machine learning. Our diverse expertise puts us in the advantageous position of being able to tackle problems in a broad domain of Algorithmic Game Theory, Computational Economics and Finance, both from a theoretical and an applied perspective. This is evident from the extensive (and diverse) list of research problems our group has contributed to, as well as from the numerous joint publications among the group members. We aim to continue advancing research on multi-agent systems for computational economics and finance, and explore the vast landscape in this fascinating area.

References

- [1] J. Nash, Non-Cooperative Games, *Annals of Mathematics* **54**(2) (1951), 286–295.
- [2] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R.P. Leme and É. Tardos, Bounding the inefficiency of outcomes in generalized second price auctions, *Journal of Economic Theory* **156** (2015), 343–388.
- [3] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos and M. Kyropoulou, Revenue Guarantees in the Generalized Second Price Auction, *ACM Transactions on Internet Technology* **14**(2–3) (2014), 17:1–17:19.
- [4] A. Deligkas, T. Melissourgos and P.G. Spirakis, Walrasian Equilibria in Markets with Small Demands, in: *AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems*, 2021, pp. 413–419.
- [5] I. Caragiannis and A.A. Voudouris, The Efficiency of Resource Allocation Mechanisms for Budget-Constrained Users, *Mathematics of Operations Research* **46**(2) (2021), 503–523.
- [6] A.A. Voudouris, A note on the efficiency of position mechanisms with budget constraints, *Information Processing Letters* **143** (2019), 28–33.
- [7] A.A. Voudouris, Simple combinatorial auctions with budget constraints, *Theoretical Computer Science* **842** (2020), 6–17.
- [8] I. Caragiannis, X. Chatzigeorgiou, P. Kanellopoulos, G.A. Krimpas, N. Protopapas and A.A. Voudouris, Efficiency and complexity of price competition among single-product vendors, *Artificial Intelligence* **248** (2017), 9–25.
- [9] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos and L. Moscardelli, Tight Bounds for Selfish and Greedy Load Balancing, *Algorithmica* **61**(3) (2011), 606–637.
- [10] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou and E. Papaioannou, The Impact of Altruism on the Efficiency of Atomic Congestion Games, in: *Proceedings of the 5th International Symposium on Trustworthy Global Computing (TGC)*, Lecture Notes in Computer Science, Vol. 6084, Springer, 2010, pp. 172–188.
- [11] I. Caragiannis, C. Kaklamanis and P. Kanellopoulos, Taxes for linear atomic congestion games, *ACM Transactions on Algorithms* **7**(1) (2010), 13:1–13:31.
- [12] I. Caragiannis, C. Kaklamanis and P. Kanellopoulos, Improving the Efficiency of Load Balancing Games through Taxes, in: *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE)*, Lecture Notes in Computer Science, Vol. 5385, Springer, 2008, pp. 374–385.
- [13] Y. Giannakopoulos, E. Koutsoupias and M. Kyropoulou, The anarchy of scheduling without money, *Theoretical Computer Science* **778** (2019), 19–32.
- [14] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos and E. Papaioannou, Scheduling to maximize participation, *Theoretical Computer Science* **402**(2–3) (2008), 142–155.
- [15] T.C. Schelling, Models of segregation, *American Economic Review* **59**(2) (1969), 488–493.
- [16] T.C. Schelling, Dynamic models of segregation, *Journal of Mathematical Sociology* **1**(2) (1971), 143–186.
- [17] E. Elkind, J. Gan, A. Igarashi, W. Suksompong and A.A. Voudouris, Schelling Games on Graphs, in: *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019, pp. 266–272.
- [18] P. Kanellopoulos, M. Kyropoulou and A.A. Voudouris, Modified Schelling games, *Theoretical Computer Science* **880** (2021), 1–19.
- [19] P. Kanellopoulos, M. Kyropoulou and A.A. Voudouris, Not all Strangers are the Same: The Impact of Tolerance in Schelling Games, *CoRR abs/2105.02699* (2021).
- [20] A. Agarwal, E. Elkind, J. Gan and A.A. Voudouris, Swap Stability in Schelling Games on Graphs, in: *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020, pp. 1758–1765.
- [21] M. Bullinger, W. Suksompong and A.A. Voudouris, Welfare Guarantees in Schelling Segregation, *Journal of Artificial Intelligence Research* **71** (2021), 143–174.
- [22] I. Caragiannis, A. Filos-Ratsikas, P. Kanellopoulos and R. Vaish, Stable fractional matchings, *Artificial Intelligence* **295** (2021), 103416.
- [23] I. Caragiannis, P. Kanellopoulos and A.A. Voudouris, Bounding the Inefficiency of Compromise in Opinion Formation, *Algorithmica* **84**(1) (2022), 234–271.
- [24] C. Kaklamanis, P. Kanellopoulos, K. Papaioannou and D. Patouchas, On the price of stability of some simple graph-based hedonic games, *Theoretical Computer Science* **855** (2021), 1–15.

- 1 [25] G. Birmpas, J. Gan, A. Hollender, F.J.M. Cossío, N. Rajgopal and A.A. Voudouris, Optimally Deceiving a Learning Leader in Stackelberg Games, *Journal of Artificial Intelligence Research* **72** (2021), 507–531. 1
- 2 [26] E. Elkind, J. Gan, S. Obratzsova, Z. Rabinovich and A.A. Voudouris, Protecting elections by recounting ballots, *Artificial Intelligence* **290** 2
- 3 (2021), 103401. 3
- 4 [27] E.C. Akrida, A. Deligkas, T. Melissourgos and P.G. Spirakis, Connected Subgraph Defense Games, *Algorithmica* **83**(11) (2021), 3403– 4
- 5 3431. 5
- 6 [28] A. Kiayias, E. Koutsoupias, M. Kyropoulou and Y. Tselekounis, Blockchain Mining Games, in: *Proceedings of the 2016 ACM Conference* 6
- 7 *on Economics and Computation, EC '16*, ACM, 2016, pp. 365–382. 7
- 8 [29] T. Melissourgos and P.G. Spirakis, Existence of Evolutionarily Stable Strategies Remains Hard to Decide for a Wide Range of Payoff 8
- 9 Values, in: *Algorithms and Complexity - 10th International Conference, CIAC, Lecture Notes in Computer Science*, Vol. 10236, 2017, 9
- 10 pp. 418–429. 10
- 11 [30] A. Deligkas, J. Fearnley, T. Melissourgos and P.G. Spirakis, Approximating the existential theory of the reals, *Journal of Computer and* 11
- 12 *System Sciences* **125** (2022), 106–128. 11
- 13 [31] F. Brandt, V. Conitzer, U. Endriss, J. Lang and A.D. Procaccia (eds), *Handbook of Computational Social Choice*, Cambridge University 12
- 14 Press, 2016. 13
- 15 [32] E. Anshelevich, A. Filos-Ratsikas, N. Shah and A.A. Voudouris, Distortion in Social Choice Problems: The First 15 Years and Beyond, in: 14
- 16 *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 2021, pp. 4294–4301. 14
- 17 [33] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas and A.A. Voudouris, Peeking behind the ordinal curtain: Improving distortion via cardinal 15
- 18 queries, *Artificial Intelligence* **296** (2021), 103488. 16
- 19 [34] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas and A.A. Voudouris, A Few Queries Go a Long Way: Information-Distortion Tradeoffs in 17
- 20 Matching, *Journal of Artificial Intelligence Research* (2022). 17
- 21 [35] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas and A.A. Voudouris, Don't Roll the Dice, Ask Twice: The Two-Query Distortion of Matching 18
- 22 Problems and Beyond, *CoRR abs/2203.01872* (2022). 19
- 23 [36] A. Filos-Ratsikas, E. Micha and A.A. Voudouris, The distortion of distributed voting, *Artificial Intelligence* **286** (2020), 103343. 20
- 24 [37] A. Filos-Ratsikas and A.A. Voudouris, Approximate Mechanism Design for Distributed Facility Location, in: *Proceedings of the 14th* 21
- 25 *International Symposium on Algorithmic Game Theory (SAGT)*, 2021, pp. 49–63. 22
- 26 [38] E. Anshelevich, A. Filos-Ratsikas and A.A. Voudouris, The distortion of distributed metric social choice, *Artificial Intelligence* (2022), 23
- 27 103713. 24
- 28 [39] I. Caragiannis, N. Shah and A.A. Voudouris, The Metric Distortion of Multiwinner Voting, in: *Proceedings of the 36th AAAI Conference* 24
- 29 *on Artificial Intelligence (AAAI)*, 2022. 25
- 30 [40] C. Anil and X. Bao, Learning to Elect, *CoRR abs/2108.02768* (2021). 26
- 31 [41] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos and M. Kyropoulou, The Efficiency of Fair Division, *Theory of Computing Systems* **50**(4) 27
- 32 (2012), 589–610. 28
- 33 [42] M. Kyropoulou, J. Ortega and E. Segal-Halevi, Fair cake-cutting in practice, *Games and Economic Behavior* **133** (2022), 28–49. 29
- 34 [43] A. Deligkas, J. Fearnley, A. Hollender and T. Melissourgos, Constant Inapproximability for PPA, *CoRR abs/2201.10011* (2022). 30
- 35 [44] A. Deligkas, J. Fearnley, T. Melissourgos and P.G. Spirakis, Computing exact solutions of consensus halving and the Borsuk-Ulam theorem, 31
- 36 *Journal of Computer and System Sciences* **117** (2021), 75–98. 31
- 37 [45] J. Gan, W. Suksompong and A.A. Voudouris, Envy-freeness in house allocation problems, *Mathematical Social Sciences* **101** (2019), 32
- 38 104–106. 33
- 39 [46] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas, A. Hollender and A.A. Voudouris, Maximum Nash welfare and other stories about EFX, 34
- 40 *Theoretical Computer Science* **863** (2021), 69–85. 34
- 41 [47] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas and A.A. Voudouris, Fair Division of Indivisible Goods: A Survey, in: *Proceedings of the* 35
- 42 *31st International Joint Conference on Artificial Intelligence (IJCAI)*, 2022. 36
- 43 [48] M. Kyropoulou, W. Suksompong and A.A. Voudouris, Almost envy-freeness in group resource allocation, *Theoretical Computer Science* 37
- 44 **841** (2020), 110–123. 38
- 45 [49] I. Caragiannis, P. Kanellopoulos and M. Kyropoulou, On Interim Envy-Free Allocation Lotteries, in: *Proceedings of the 22nd ACM Con-* 39
- 46 *ference on Economics and Computation (EC)*, ACM, 2021, pp. 264–284. 39
- 47 [50] Y. Giannakopoulos and M. Kyropoulou, The VCG Mechanism for Bayesian Scheduling, *ACM Trans. Economics and Comput.* **5**(4) (2017), 40
- 48 19:1–19:16. 41
- 49 [51] I. Caragiannis, C. Kaklamanis and M. Kyropoulou, Limitations of Deterministic Auction Design for Correlated Bidders, *ACM Trans.* 42
- 50 *Comput. Theory* **8**(4) (2016), 13:1–13:18. 42
- 51 [52] B. de Keijzer, M. Kyropoulou and C. Ventre, Obviously Strategyproof Single-Minded Combinatorial Auctions, in: *47th International* 43
- Colloquium on Automata, Languages, and Programming, *ICALP 2020*, LIPIcs, Vol. 168, 2020, pp. 71:1–71:17. 44
- [53] M. Kyropoulou, C. Ventre and X. Zhang, Mechanism Design for Constrained Heterogeneous Facility Location, in: *Proceedings of the 12th* 45
- International Symposium Algorithmic Game Theory (SAGT)*, 2019, pp. 63–76. 46
- [54] P. Kanellopoulos, A.A. Voudouris and R. Zhang, On Discrete Truthful Heterogeneous Two-Facility Location, in: *Proceedings of the 31st* 47
- International Joint Conference on Artificial Intelligence (IJCAI)*, 2022. 48
- [55] A. Deligkas, A. Filos-Ratsikas and A.A. Voudouris, Heterogeneous facility location with limited resources, in: *Proceedings of the 36th* 49
- AAAI Conference on Artificial Intelligence (AAAI)*, 2022. 49
- [56] M. Kyropoulou and C. Ventre, Obviously Strategyproof Mechanisms without Money for Scheduling, in: *Proceedings of the 18th Interna-* 50
- tional Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19*, pp. 1574–1581. 51

- [57] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos and M. Kyropoulou, On Low-Envy Truthful Allocations, in: *Proceedings of the First International Conference on Algorithmic Decision Theory (ADT)*, Lecture Notes in Computer Science, Vol. 5783, Springer, 2009, pp. 111–119.
- [58] I. Caragiannis, A. Filos-Ratsikas, S. Nath and A.A. Voudouris, Truthful mechanisms for ownership transfer with expert advice, *CoRR abs/1802.01308* (2018).
- [59] P. Kanellopoulos, M. Kyropoulou and H. Zhou, Financial Network Games, in: *Proceedings of the Second ACM International Conference on AI in Finance (ICAIF)*, 2021.
- [60] P. Kanellopoulos, M. Kyropoulou and H. Zhou, Forgiving Debt in Financial Network Games (Extended Abstract), in: *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2022.
- [61] P. Kanellopoulos, M. Kyropoulou and H. Zhou, Forgiving Debt in Financial Network Games, in: *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 2022.
- [62] A. Brabazon, M. Kampouridis and M. O’Neill, Applications of genetic programming to finance and economics: past, present, future, *Genetic Programming and Evolvable Machines* **21**(1) (2020), 33–53.
- [63] A. Adegbeye and M. Kampouridis, Machine learning classification and regression models for predicting directional changes trend reversal in FX markets, *Expert Systems with Applications* **173** (2021), 114645.
- [64] M. Kampouridis and F.E. Otero, Evolving trading strategies using directional changes, *Expert Systems with Applications* **73** (2017), 145–160.
- [65] S. Cramer, M. Kampouridis and A.A. Freitas, Decomposition genetic programming: An extensive evaluation on rainfall prediction in the context of weather derivatives, *Applied Soft Computing* **70** (2018), 208–224.
- [66] S. Cramer, M. Kampouridis, A.A. Freitas and A. Alexandridis, Stochastic model genetic programming: Deriving pricing equations for rainfall weather derivatives, *Swarm and Evolutionary Computation* **46** (2019), 184–200.
- [67] M. Kampouridis, S.-H. Chen and E. Tsang, The market fraction hypothesis under different genetic programming algorithms, in: *Information Systems for Global Financial Markets: Emerging Developments and Effects*, IGI Global, 2012, pp. 37–54.
- [68] M. Kampouridis, S.-H. Chen and E. Tsang, Market fraction hypothesis: A proposed test, *International Review of Financial Analysis* **23** (2012), 41–54, Complexity and Non-Linearities in Financial Markets: Perspectives from Econophysics.
- [69] S.-H. Chen, M. Kampouridis and E. Tsang, Microstructure dynamics and agent-based financial markets, in: *International workshop on multi-agent systems and agent-based simulation*, Lecture Notes in Artificial Intelligence, Springer, 2010, pp. 121–135.
- [70] M. Kampouridis, S.-H. Chen and E. Tsang, Microstructure dynamics and agent-based financial markets: can dinosaurs return?, *Advances in Complex Systems* **15**(supp02) (2012), 1250060.
- [71] M. Kampouridis and K.M. Sim, A GP approach for price-speed optimizing negotiation, in: *2013 IEEE Congress on Evolutionary Computation*, IEEE, 2013, pp. 1170–1177.