ERC Consolidator Grant 2019
Research Proposal Part B2 : Robust Algorithmic Game Theory

Section a. State of The Art and Objectives

Background and State of The Art

Algorithmic Game Theory (AGT) is a fascinating interdisciplinary field of research that emerged in the beginning of the millennium to address the challenges posed by the Internet. It lies at the crossroads of computer science, game theory, and economics, and is driven by real-life problems, such as keyword auctions in search engines, electronic marketplaces and social networks. This field is full of intellectual excitement internally, and has already significantly influenced its three parent disciplines. It integrates the economic theory of mechanism design with the computer science theory of algorithms and computations, and puts forward a framework for designing computational mechanisms with proper incentives for cooperations in complex markets. For the last two decades AGT studied with great success the interaction between algorithms and incentives, and developed a remarkable theory that facilitates rigorous analysis and groundbreaking results (see [122, 138, 137, 134] for influential surveys and textbooks on the area). This theory has generated several impressive applications, ranging from huge advertising auctions in search engines and web pages [53, 159] to FCC spectrum auctions for wireless bandwidth [113], all involving transactions worth hundreds of billions of dollars. Beyond combinatorial auctions, AGT has studied a wide range of applications, such as matching markets, bilateral trade, scheduling, pricing schemes, diffusion in social networks, voting methods, and more. AGT has become the arena of productive cross-fertilization between game theory, economics and algorithms.

The field of AGT has received tremendous recognition within the academic community, including the Gödel Prize in 2012 [109, 121, 143] and the Nevanlinna Prize in 2018. This recognition reflects the field's excitement, breadth, accomplishment, and promise. Indeed, AGT has developed novel approaches for addressing the unique challenges posed by the Internet, injected new concepts that benefit the very core of research in economics, and provided theory that studied the economic, algorithmic and computational considerations in a unified way.

During its first two decades, AGT invented itself from scratch, combining two presumably different and distant disciplines into a coherent framework. This process entailed overcoming numerous challenges and obstacles. As a newly emerging field, it was natural and necessary for AGT to employ simplified models to facilitate this non-trivial synergy, and to abstract away many real-world considerations that complicate the models. However, for all its theoretical appeal, AGT theory suffers from some overly strong assumptions that may render the theory unsuitable for addressing some of the practical problems AGT seeks to tackle. Use of these non-realistic assumptions impedes our ability to analyze existing mechanisms and to design new ones for real applications. In light of the far-reaching implications of computerized economic markets, developing a more robust theory of AGT is crucial and timely.

Objectives

The goal of this proposed project is to develop theory of robust AGT that will make it more applicable to practical settings and help it better inform the design of real systems. To achieve this goal, I will highlight limitations in the traditional AGT paradigm, by identifying the unrealistic assumptions employed in AGT theory, and considering aspects that exist in real-world applications that are often abstracted away in theory. The unrealistic assumptions must be made as flexible and nuanced as possible. There are large gaps to be filled before the complete theory of AGT finds full applicability in real Internet applications. Nonetheless, at the same time, the field of AGT is becoming increasingly influential and relevant to an increasing number of applications, and this proposed project exhibits a crucial step in developing a theoretical framework that will guide the design of real applications. Despite recent progress on this matter in recent years, many problems remain open. More than a decade ago, Papadimitriou wrote [122]: "*A fascinating fusion of ideas from both fields game theory and algorithms came into being and was used productively in the effort to illuminate the mysteries of the Internet. It has come to be called algorithmic game theory*." It is my vision that a fascinating fusion of ideas will emerge from this project, that will come to be called *robust algorithmic game theory"*.

Research Approach: Wilson's doctrine as an inspiration In [161] Wilson wrote: "*I foresee the progress of game theory as depending on successive reductions in the base of common knowledge required to conduct useful analyses of practical problems. Only by repeated weakening of common knowledge assumptions will the theory approximate reality*." In a similar spirit, Bergemann and Morris wrote in [19]: "*The mechanism design literature assumes too much common knowledge of the environment among the players and planner ... The modelling strategy must be to first make explicit the implicit common knowledge assumptions and then weaken them*." This approach, which has come to be known as *Wilson's doctrine*, animates this current project. It posits that simple, detail-free mechanisms should be preferred in order to alleviate the risks of applying excessively strong assumptions. If invalid or ineffective assumptions are applied, the resulting theory can lose some of its predictive validity for real-world settings. This proposed research is inspired by Wilson's doctrine, and seeks to develop theory that is robust to model misspecification. Of course, because some of these original, more rigid assumptions have some validity and usefulness, an optimal balance must be found that modifies these older assumptions in such way that still leads to tractable models amenable to rigorous mathematical analysis.

Canonical Example: Mechanism Design and Combinatorial Markets Before enumerating the robustness aspects that will be the core of this project, I give a high-level description of the canonical model in AGT, namely combinatorial markets/auctions and algorithmic mechanism design (see surveys [123, 27]). A detailed description of the model is given in Section b. Combinatorial auctions are held in every sphere of activity, from government spectrum auctions in the United States and Europe for assigning wireless spectrum licenses [113], through sponsored search auctions in virtually all search engines [53, 159], to sales of millions of items on eBay every day. The scope of the current research proposal is by no means restricted to combinatorial auctions. Yet, this application will serve as the primary model for demonstrating our ideas, and is given here for concreteness.

In a combinatorial market problem, a seller is allocating m heterogeneous and indivisible items to n agents with heterogeneous and private preferences over bundles of items. The preference of an agent i is represented by a valuation function v_i that assigns to each bundle $S \subseteq \{1, \ldots, m\}$ a non-negative real value that specifies this agent's value for receiving this bundle. The utility a player derives from a bundle $S \subseteq [m]$ which she values at $v_i(S)$ and pays p for is $v_i(S) - p$. A (centralized) *mechanism* receives a vector of agent valuations as input, and chooses an outcome. An outcome has two components: the allocation, which specifies which bundle is allocated to each agent, and the payments, which specify how much (if at all) each agent pays.

This proposal targets the following areas of application related to the combinatorial markets model: (1) Social welfare maximization: the sum of agent valuations for their allocations. (2) Revenue maximization: the expected revenue the seller extracts from the buyers. (3) Market equilibrium: the steady state allocation reached among agents, which consists of market-clearing prices.

Robustness Aspects In what follows, we specify the different robustness aspects we seek to study to achieve the ultimate goal of developing a robust AGT theory. For each one of the aspects, we identify the overly strong assumption being employed, and explain why we believe it is unsuitable for real-world scenarios. In Section b we expand on each of these aspects. In particular: (1) we consider ways to weaken the assumption in a way that leads to a more realistic model while being amenable to rigorous analysis, (2) we suggest both high-level research directions and concrete open problems that we find particularly interesting for the associated aspect, and (3) we provide preliminary results from our research. The concrete open problems often range from important extensions of existing results (low risk) to more ambitious and open-ended problems related to mostly unexplored areas (high risk). The different aspects are highly interleaved and complementary. Together, they build a theory of robust AGT and portray the limits of possibilities within the regime of economics and computation.

(1) Informational and outcome externalities. Traditional AGT assumes that every agent fully knows his or her value for each outcome, and values are independent and private. Moreover, the value an agent has for some outcome depends only on her own allocation. In practice, an agent's valuation may depend on information owned by others and may be highly correlated with other agent valuations. In addition, an agent's valuation for some allocation vector may depend on other agent allocations. We seek mechanisms that are robust to such externalities.

(2) Unknown priors. Traditional AGT assumes that the seller and buyers have perfect distributional knowledge of the preferences of all the buyers. In practice, neither the seller not the buyers have perfect distributional information about the preferences of the participating agents. We seek mechanisms that are robust to uncertainty regarding the underlying priors.

(3) Behavioral considerations. Traditional AGT assumes that agents are fully rational, acting as utility maximizing agents. At the same time, it is well known by now that in practice, agent actions are influenced by various behavioral considerations and cognitive biases. We aim to devise models that capture various behavioral considerations and cognitive biases, and study their effects on the stability and efficiency of markets.

(4) Divergence from precise structural properties. Classes of set functions, distributions, and other central objects of study are usually characterized by some convenient properties that make them useful in optimization, learning, etc. In practice, the input might only approximately adhere to the given property. We seek to study the robustness of the AGT theory to small divergence from these properties; namely, whether and to what extent small deviations form these properties can be tolerated.

(5) Beyond worst case analysis. In traditional AGT, a mechanism's performance is measured with respect to its worst-case inputs. While worst-case analysis is often more tractable analytically, and gives strong guarantees (when it succeeds to do so), it bases its analysis on worst case scenarios, that do not necessarily resemble real-world instances. It is of great interest to develop input models that resemble real-world instances, and develop analysis methods to study them.

The Ramifications of Robust AGT Theory Unrealistic assumptions may lead to too optimistic results. For example, it is commonly assumed in AGT literature that agents have private and independent values. This assumption leads to the sweeping result showing that a welfare-maximizing outcome can be achieved, in great generality, via the celebrated Vickrey-Clarke-Groves (VCG) mechanism [160, 40, 85] (disregarding computational constraints). In VCG mechanisms, it is in the best interest of every agent to reveal his or her true preferences, regardless of the preferences and strategies of others. In reality, however, values may be highly correlated. An elegant model that captures this reality is the *interdependent values* model of Milgrom and Weber [116]. In this more realistic model, such a powerful result is unattainable even for some of the simplest imaginable settings. Thus, while a common conviction in AGT is that welfare maximization is well understood, when we turn to the more realistic interdependent values model, welfare maximization is extremely challenging and far from being understood.

In some cases, however, more realistic assumptions may be harnessed to obtain better results than those obtained under the oversimplified assumptions. For example, it is commonly assumed in AGT literature that agents act to maximize their utility. Under this assumption, a *market equilibrium* — a highly desirable state, composing of a set of item prices that clear the market — rarely exists [86]. However, it has been recently shown by [13] that if agents are cognitively biased by the *endowment effect* [154], then a market equilibrium exists much more broadly. Thus, in this example, relaxing the overly strong assumption of full rationality leads to more stable markets.

The proposed project will identify and study both types of scenarios. On the one hand, we seek to identify and leverage more realistic assumptions that can be harnessed for more stable and efficient results. On the other hand, we will study cases where more realistic assumptions lead to unavoidable deterioration of results and may require newly designed mechanisms.

Section b. Methodology

Canonical Example: Mechanism Design and Combinatorial Markets. Our starting point is the following model of combinatorial markets/auctions (see the surveys [123, 27]): A seller is allocating m heterogeneous and indivisible items to n agents with heterogeneous and private preferences over bundles of items. The preference of an agent i is represented by a valuation function v_i that assigns to each bundle $S \subseteq \{1, \ldots, m\}$ a non-negative real value that specifies this agent's value for receiving this bundle of items. We assume throughout free disposal, i.e., that for $S \subseteq T$ we have that $v_i(S) \leq v_i(T)$, and normalization, i.e., that $v_i(\emptyset) = 0$. The setting is quasi-linear, meaning that the utility a player derives from a bundle $S \subseteq [m]$ which she values at $v_i(S)$ and pays p for is $v_i(S) - p$. It may be assumed that agent values are drawn from a joint distribution D.

A (centralized) *mechanism* receives one valuation function from each agent as input, and chooses an outcome. An outcome has two components: the allocation, which specifies which bundle is allocated to every agent, and the payments, which specify how much (if at all) each agent pays. An allocation is denoted by a vector $x = (x_1, \ldots, x_n)$, where x_i is the bundle allocated to agent i. The payments are denoted by $p_1, \ldots, p_n$.

We distinguish between *truthful* and *non-truthful* mechanisms. A truthful mechanism is a mechanism that creates incentives for the agents to reveal their private vales truthfully. To analyze the performance of non-truthful mechanisms, we must make behavioral assumptions. Traditionally, the assumption has been that selfish agents will play Nash equilibrium strategies, and a mechanism's performance is measured by the *price of anarchy* (PoA) — the ratio of the performance in the worst Nash equilibrium and the optimal performance [109, 125].

The Approximation Paradigm. A remark regarding the approximation paradigm is in order. An enormous amount of research in CS over the past decades uses the approximation lens to reason about fundamental questions that cannot be optimally solved due to bounded computation, bounded communication, partial information about future events and other limitations. The approximation paradigm has not been historically part of the economic theory literature, which concentrated instead on identifying often strong assumptions under which optimal solutions can be attained. There is a tremendous opportunity to use this methodology to extend the scope of research of traditional economic theory and establish near optimality guarantees under much weaker conditions, thereby significantly expanding the scope of design in which rigorous analysis can be made. Indeed, the approximation approach has begun penetrating into the theoretical economic literature in recent years [8].

In Robust AGT theory, where many traditional assumptions are successfully challenged, approximation will play an even more important role. Often, relaxing simplifying assumptions leads to scenarios where optimal solutions are unattainable. Approximation seems to be the right paradigm for reasoning about what can and cannot be done in such scenarios. Although the approximation results we expect to obtain are quantitative in nature, they should not be taken too literally. Rather, they serve mainly to provide qualitative insights, to help identify features that are potentially beneficial in practice, and to distinguish between tolerable and intolerable scenarios.

Objective 1: Informational and Outcome Externalities

The traditional assumption in the combinatorial auction model is that agents have *independent private values* (IPV). This means that every agent is fully aware of her own value, and there is no correlation between agent values. Over the last two decades, beautiful theory was developed that reconciles the computational and economic viewpoints within the IPV model (see, e.g., [121, 137, 120] for influential papers and surveys).

For all its theoretical appeal, the IPV model is not rich enough to capture common informational structures exhibited by many auction settings. Indeed, agent preferences in practice are typically neither independent nor private. Work in AGT recently started to look into the realm of correlated valuations [126, 50, 141, 32] but only scratched the surface.

To capture informational externalities, we propose to use the model of *interdependent values* (IDV), introduced by Milgrom and Weber [116]. In this mode, every agent $i \in [n]$ has a signal s_i, which captures her *private* information about the good (e.g., data about a particular impression for online advertising); the signals $s_1, \ldots, s_n$ are drawn from a joint (possibly correlated) distribution. Every agent i also has a public valuation function v_i, which is a monotone function of all n private signals; i.e., $v_i = v_i(s_1, \ldots, s_n)$.

A broad research goal is to apply the algorithmic mechanism design lens to the study of mechanisms for the general interdependent model (IDV). Notably, as we turn from the IPV model to the more realistic IDV model, even the social welfare objective, which is considered to be well-understood, become extremely challenging. This is true even for the simplest setting of a single item, where positive results are known only under the stringent condition known as *single-crossing* (SC) [114, 11, 141, 32, 44, 93, 44].

The IDV model is very different in nature than its IPV counterpart. In studying this model, we need to define what constitutes a robust mechanism and what captures a robust truthfulness notion. Two relevant notions are (a) *ex-post incentive compatible* (EPIC), which requires that truth telling is an ex-post equilibrium, meaning that truth telling is a best response to truth telling by others for *every* signal profile, and (b) *Bayesian IC* (BIC), where truth telling is a best response to truth telling by others only in expectation over signals. Note that the BIC notion implicitly assumes that agents know precisely the probability distribution of the signals of other agents, while the EPIC notion requires no such knowledge. Adhering to Wilson's doctrine, EPIC mechanisms should be preferred over BIC mechanisms.

As it turns out, *single-crossing* (SC) is a crucial property that enables positive results in the realm of interdependent values. SC means that the impact of agent i's private signal, s_i, on her own valuation is greater than the impact of s_i on the valuation of any other agent. Under the SC condition, a generalization of the VCG mechanism maximizes social welfare [114, 11, 141, 32], but in the absence of SC, an efficient outcome may not be implementable [44, 93]. However, in many practical scenarios, single crossing is violated [44]. Moreover, while SC is a celebrated property assumed in virtually all works about interdependent values, we recently found that even in the simplest imaginable scenarios beyond a single item (e.g., 2 items and 2 agents with unit-demand valuations), SC does not seem to capture efficiency.

Major research goal: Develop a coherent and broad theory that studies the computational and economic aspects of the design and analysis of mechanisms for combinatorial auctions in IDV settings. In particular, we aim to: (1) identify key properties (such as submodular over signals [56], and signal affiliations [141]) that enable the analysis and design of mechanisms for the IDV settings. (2) identify appropriate notions of truthfulness for this setting (e.g., DSIC, EPIC, BIC), and study the tradeoffs related to these notions. (3) derive welfare and revenue guarantees for truthful poly-time mechanisms for various families of valuations. (4) study the tradeoffs between prior-dependent and prior-independent mechanisms. (5) quantify inefficiencies that arise in equilibrium of simple auctions, using the *price of anarchy* measure. (6) extend the above settings to scenarios where the valuations are unknown.

Preliminary results on informational externalities. In [56] we invoke for the first time a computational and economic study of auctions in IDV settings, without assuming single crossing. We introduce a parameterized version of single-crossing, c-single-crossing, and obtain a host of welfare and revenue guarantees for single-item settings as a function of the parameter c.

In a more recent working paper [57], we study IDV settings with valuations that are "submodular over signals" (SM). This condition captures scenarios with signal substitutability (informally, the effect of a signal increase on the value is more significant when other signals are lower). With this property (and a parameterized version, d-SM), we study the problem in much greater generality, namely for settings that go beyond single-parameter. In particular, we show that for an interesting class of combinatorial auctions with valuations satisfying d-SM, there exists a truthful mechanism that gives at least a constant fraction (that depends on d) of the optimal welfare. This result opens up a wide terrain of unexplored problems.

Up until now we considered informational externalities but no outcome externalities. The general model of social choice and mechanism design handles outcome externalities in great generality. In particular, every agent has some value for every possible outcome from a given set of outcomes. This general model gives rise to the celebrated VCG mechanism [160, 40, 85], showing that the welfare maximizing outcome can always be implemented, but also to famous impossibility results, such as Gibbard-Satterthwaite [80, 148] (showing essentially that in the absence of payments, every "reasonable" voting rule is susceptible to manipulation) and Robert's [133] (showing that every social choice function that is implementable must be weighted VCG). The combinatorial auctions model described above is much more structures: an outcome is an allocation vector and the value of every agent depends only on her own allocation. What happens if this assumption is relaxed?

Preliminary results on outcome externalities. In [54] we initiate a preliminary study that considers positive outcome externalities — goods that grant value not only to their owners but also to the owners' surroundings — and their effect on the design of posted price mechanisms for revenue maximization. We introduce a Bayesian model for this scenario, and devise nearly-optimal pricing schemes for various types of externalities, both for simultaneous sales and for sequential sales.

Open problem 1. *Derive tight welfare and revenue guarantees for single item settings, under c-SC and d-SM valuations. In particular, close the gaps that remain from [56, 57]. Moreover, extend these results to more general single-parameter settings (e.g., allocation of k identical items).*

following are several more ambitious problems. The following problems apply to various economic objectives (welfare, revenue), various truthfulness notions (EPIC,BIC), various families of valuations (unit-demand, submodular, subadditive, etc.), and deterministic vs. random mechanisms. For each of these problems, we are interested in both existential (do good mechanisms exist?) and computational (can good mechanisms be computed?) results.

Open problem 2. *Extend the IDV study of [56] to multi-parameter settings, and to general combinatorial auctions.*

Open problem 3. *In aligning with Wilson's doctrine, it is more appropriate to consider scenarios where the valuation functions are private as well (in addition to signals). What are the implications of private valuations? This question is interesting already for a single item setting.*

Open problem 4. *What are the implications of considering* outcome externalities *in addition to informational externalities? (for a model that considers both types of externalities, see [93]).*

Open problem 5. *What is the price of anarchy (PoA) of simple auctions, such as simultaneous first-price and second-price auctions [39, 24, 140, 153, 69] in IDV settings? To what extent the smoothness framework [140] applies to IDV settings? What are the computational implications of IDV settings in this respect?*

Objective 2: Unknown Priors

In the classic economic and AGT literature, in many canonical models and questions, the prior partial knowledge of a decision maker regarding the state of the world is stylistically modeled as a Bayesian prior. This is the case in the revenue maximization problem, where the model stylistically assumes that the valuations of the buyers are drawn from some underlying distribution that is known to the seller, and the seller's goal is to maximize her revenue in expectation over this distribution.

In adherence to Wilson's doctrine, a simple, detail-free mechanism, which does not depend on the priors, should be preferred over complex prior-dependent mechanisms. While for welfare guarantees prior-independent mechanisms are often sufficient, to obtain revenue guarantees complex, prior-dependent mechanisms are required [119]. Two approaches have been taken with respect to this problem, namely *competition complexity* (**a term coined in our recent paper [58]**), and *sample complexity* [42].

Approach 1: competition complexity Perhaps with more competition in the market, e.g., with additional buyers, the optimal revenue can be extracted by simple prior-independent mechanisms. This key observation led to the beautiful result of Bulow and Klemperer [30], showing that for a single item sale with any number of i.i.d. bidders, one additional bidder suffices to guarantee the optimal revenue in the prior-independent VCG mechanism. This result is particularly interesting in light of the observation that simple mechanisms tend to attract more participants [115]. Thus, it is natural to assume that more bidders will be present if a simple mechanism is chosen. The following natural question arises: how much more attractive (in terms of the required additional bidders) prior-independent mechanisms should be to ensure sufficient revenue?

Preliminary results. In [58] we coined the term *competition complexity* to quantify this tradeoff, and provided the first bounds on the competition complexity in multi-parameter settings (specifically, additive valuations subject to downward-closed constraints). This notion is essentially the economic equivalence of the *resource augmentation* approach in computer science [151, 97, 130]. Variants of the competition complexity notions were later studied in [68, 112]. Interestingly, this approach was also considered in the problem of identifying the optimal seeds in diffusion processes in social networks [3].

Major research goals: First, we seek to study the competition complexity with respect to revenue maximization in more complex markets. Specifically, our goal is to identify market scenarios that admit low competition complexity. This applies to: (1) various types of preferences (e.g., unit-demand/submodular/subaddtive valuations, Fedex-type valuations [75]), (2) various truthfulness notions (DSIC, BIC), (3) various prior-independent mechanisms (e.g., VCG or other), and (4) deterministic vs. randomized mechanisms. Our more ambitious goal is to apply the competition complexity paradigm to additional areas of applications, such as (1) gains from trade in bilateral trade scenarios [12], (2) combinatorial public projects [29], and (3) scheduling. We believe that this direction bears a huge potential.

Approach 2: sample complexity A second natural approach for addressing the problem of unknown priors is to *learn* them. Over the past few years, work in the interface of AGT and learning theory has turned to study a more demanding and realistic model where the assumption of complete information of the prior distribution is relaxed to the more realistic assumption of having access only to samples from this fixed, yet unknown, distribution. The goal is to learn with high probability, from a moderate number of samples, a solution whose quality approximately matches that of the best solution that could be designed had one gained access to complete information of the prior. This approach was initiated by Cole and Roughgarden [42] with their study of revenue maximization from samples, which has established itself as an influential line of research in the interface of economics and computation in recent years (see, e.g., [117, 47, 82, 31, 83, 48, 15]). Despite this impressive line of works on revenue maximization, we believe that only the tip of the iceberg has been scratched. only recently have first steps been made in viewing the canonical "prophet inequality" problem through this lens [43], while other canonical settings seem to remain completely unexplored as of yet.

Major research goal: Apply the sample complexity paradigm to canonical AGT settings, including (1) various models of prophet inequality and posted prices, (2) matching markets, (3) bilateral trade, and (4) dynamic pricing in Bayesian markets. In addition, despite the progress made on sample complexity for revenue maximization, quite a few major question still remain, specifically within the realm of multi-parameter settings.

Preliminary results. In [71] we study a complex Bayesian market setting, where agent valuations are drawn from independent priors over submodular valuations (in fact, a superset of submodular, termed XOS). We establish a sampling algorithm that produces static item prices that obtain expected welfare of at least a half of the optimal welfare, for any arrival order and tie breaking of the agents. This paper was later considered as establishing the roots for an elaborate study of prophet inequalities for combinatorial markets and their relation to pricing schemes. In [51] we provide a unified framework for the study of prophet inequalities and posted price mechanisms that extends the result in [71].

Open problem 6. *In the classic prophet inequality setting, with n i.i.d. random variables from an unknown distribution, what is the exact bound that can be obtained with n samples? (it is known from [43] that $O(n^2)$ samples are sufficient to obtain the same bound as knowing the distribution, which is known to be 0.745 [1]. Similarly, what is the exact bound that can be obtained with n samples for non-identical distributions?*

Open problem 7. *Apply the sample complexity paradigm to prophet inequality in combinatorial settings. For example: is it possible to achieve constant approximation for general matroid with n samples? is it possible to achieve constant approximation for combinatorial auctions with submodular valuations (or simpler valuations) with n samples? Note that in [71] the sample complexity is super linear.*

Open problem 8. *Apply the sample complexity paradigm to bilateral trade. Specifically, what is the tradeoff between the sample complexity and the approximation guarantees for the gains from trade in bilateral trade?*

Open problem 9. *Study the sample complexity of mechanisms for revenue maximization in multi-parameter settings. Some preliminary results appear in, e.g., [83].*

Objective 3: Behavioral Considerations

AGT theory has traditionally treated agents as utility maximizers (or expected utility maximizers in settings with uncertainty). Economists today have come to realize the huge contribution of behavioral research to revise their models [155], but this crucial development has skipped the current models of AGT to the most part. In light of the major developments in decision making theory and behavioral economics in the last few decades, AGT needs to update its assumptions about how people make decisions. When dealing with human subjects, there is abundance of empirical evidence from behavioral economics refuting the premise of full rationality and showing that agents in practice are susceptible to many cognitive biases [94], such as risk and loss aversion [95], procrastination [4], anchoring [157], framing, status quo bias [147], and overconfidence, among many others. Indeed, individuals in practice tend to exhibit various cognitive biases that are not captured by the classic model of expected utility maximization.

Quoting Thaler [155]: "*The new approach to economics should include two different kinds of theories: normative models that characterize the optimal solution to specific problems and descriptive models that capture how humans actually behave. [...] By adding these factors such as framing or temptation we can improve the explanatory power of economic models*." It is my view that the same applies to AGT models. Capturing behavioral considerations in AGT models will not only improve their explanatory power, but will also lead to a better design of mechanisms for real-world scenarios, which is a timely objective.

In recent years, some research in AGT has started to look into behavioral considerations. One example is the work of Kleinberg and Oren on time-inconsistent planning and procrastination [104, 105, 106, 84, 5, 6, 7]. Another example is the consideration of the celebrated *prospect theory* of Kahneman and Tversky [95] in devising robust mechanisms for risk-averse agents [33, 76, 100] and in the design of contracts for online labor markets [52].

While one might view cognitive biases as a negative phenomenon that reveals human flaws, these biases could, in some cases, be harnessed to produce better outcomes. A recent paper [14] studied the implications of the *endowment effect* — the tendency to inflate the value of items one owns, studied by the Nobel Laureate Richard Thaler [154, 108] — on the existence of market equilibrium. It is shown that the endowment effect increases market stability in the sense that pricing equilibria exist for a wider set of markets.

Preliminary results on endowment effect. In [69] we showed that allowing both bundle pricing and unsold items leads to a sweeping result, where a relaxed equilibrium with at least a half of the optimal welfare exists for *any* market. Based on our work in [72, 49], none of these relaxations alone is sufficient for such a strong result. We currently study the combined effect of the endowment effect and the natural relaxation of bundle pricing [60], and show that the optimal welfare can be obtained in every market in a bundling equilibrium with a mild endowment effect. This result is very promising, showing that combining the endowment effect with other relaxations can be unexpectedly powerful.

Behavioral assumptions lie also at the heart of the analysis of non-truthful mechanisms. How shall we analyze their performance? To answer this question we must use some behavioral assumption to predict agent behavior. Traditional economics typically models the agents as playing a (Bayesian) Nash equilibrium of the induced game. Consequently, the *price of anarchy* (PoA) was originally defined as the ratio between the performance at a (Bayesian) Nash equilibrium and the optimal outcome [109, 125]. However, the lack of natural dynamics that converge to Nash equilibrium, backed by computational and communication complexity results (e.g., [45, 36, 46, 17]), puts this prediction in doubt. In sharp contrast, there are a host of natural learning-like dynamics (such as regret minimization) that converge to more general notions of equilibria, specifically to correlated equilibria (CE), or the more general coarse correlated equilibria (CCE) [10, 87, 162]. This makes CE and CCE seem to be more natural predictions of agent behavior than Nash equilibria. Consequently, the PoA was studied with respect to more relaxed equilibrium notions that capture more realistic behavior [25, 38, 135, 136, 140, 70, 90].

Preliminary results on price of anarchy. In [73] we studied the PoA with respect to the CCE notion in the simplest imaginable auction model, namely full-information first-price auction of a single item. Even for this simple scenario (which avoids the complexity stemming from multiple items), the PoA with respect to CCE remained unknown until our recent study. Using novel techniques, we resolve this problem and provide tight bounds for the welfare guarantees that can be obtained in this scenario [73].

Major research goal. Conduct a broad study of the effects of behavioral considerations and cognitive biases on the design and analysis of mechanisms and market equilibrium. In particular, we aim to (1) identify key cognitive biases that are relevant to AGT applications and have significant effects on welfare and revenue guarantees, pricing schemes and market equilibrium, (2) identify key behavioral considerations and cognitive biases that are relevant to *dynamic mechanism design* and decision making over time, (3) devise informative models that capture the essence of these behavioral considerations and cognitive biases, yet amenable to rigorous analysis, (4) study the key AGT questions in the obtained framework.

For concreteness, we present a few examples of open problems related to the endowment effect, but they may be equally interesting for other cognitive biases that I aim to model as part of this project.

Open problem 10. *Can the endowment effect be harnessed to produce more stability in additional areas of application? can it be harnessed to produce more revenue? What is the effect of combining the endowment effect with additional relaxations?*

Open problem 11. *Find a natural process that converges to an endowed equilibrium.*

Open problem 12. *Study the price of anarchy and price of stability of simple auction formats under the endowment effect. This problem is particularly interesting for the widely studied simultaneous first- and second price auctions [69, 77, 23, 90]. This problem is interesting both in full information model and in Bayesian settings, and leads to interesting conceptual and modeling challenges.*

An area that has received a lot of attention lately is *dynamic and online mechanism design*. Common models involve dynamic arrivals and departure of agents (e.g., [124, 128, 78]), agents whose private information evolves (e.g., [9, 21, 129], or information that is obtained over time [127].

Open problem 13. *What are the behavioral considerations and cognitive biases that are most relevant within the realm of dynamic mechanism design? devise models that capture them, design mechanisms that are robust to these effects, and understand ways in which cognitive biases can be harnessed to produce stability, efficiency and revenue* *guarantees. Examples of biases that seem particularly relevant to dynamic mechanism design include* pseudocertainty *[132, 158], where multi stage decision making can lead to different decisions than one shot game, even though the payoff is the same, and* gambler's fallacy *[156], which is the tendency to think that future probabilities are altered by past events.*

Objective 4: Divergence from Structural Properties

Classes of set functions are usually characterized by some convenient properties that make them useful in optimization, characterization, learning, etc. For example, *gross substitutes* (GS) valuations [99, 86] (the discrete analog of concave functions) are useful for welfare maximization [118] and for resolving a demand query [111]; similarly, submodularity is useful for minimization [149, 92]. In practice, the input might only approximately adhere to some structural property. For example, in the current *big data* era, where the parameters of a problem are often derived from real-world data, these parameters can be only approximately evaluated [22, 150, 91, 18]). If small divergence from nice property leads to abrupt deterioration in performance, the robustness of the corresponding algorithm is in question.

In AGT, welfare approximation, revenue approximation, existence of market equilibria, price of anarchy, and other important objectives often rely heavily on the structure of the input. To fall within a particular structural class, a valuation must satisfy fairly stringent constraints. *The question is whether the guarantees associated with these stringent constraints continue to hold approximately given that these constraints hold approximately*. For example, the robustness of optimization algorithms has been studied with respect to approximate submodularity [110] and approximate gross substitutes [142].

An additional central object in AGT is the underlying distribution. Here too, some stringent conditions are often assumed (e.g., regular, MHR), and it is of great interest to study the robustness of optimization and learning algorithms to divergence from these stringent conditions. Moreover, suppose the seller knows that the true demand distribution is *close* to a given distribution F. Would the revenue extracted from a mechanism tailored for F be sufficiently close when invoked on the true distribution? This problem was studied with respect to monopoly prices in [20]. Notably, this area is closely related to *sample complexity* (see Objective 2).

Preliminary results My work on approximate structural properties spans many areas of application, including: (1) characterization and learning of *approximately modular* functions [63], (2) welfare guarantees for valuations with *limited complementarity* (often modeled as some notion of approximate substitutability), including computation complexity [62] and communication complexity [61], (3) revenue guarantees of simple auctions for valuations with limited complementarity [59], (4) price of anarchy results for valuations with limited complementarity [62, 67], and (4) welfare guarantees for interdependent valuations satisfying *approximate single crossing* or *approximate submodularity over signals* (see Objective 1).

Major research goal: Conduct a broad study of the robustness of various classes of set functions (such as GS, submodular, etc.), and various classes of distributions. We are interested in two different aspects of robustness: (1) *Robustness with respect to performance measures*: do performance measures provided for valuations adhering to certain constraints continue to hold approximately when these constraints hold approximately? (2) *Robustness with respect to proximity measures*: for properties that admit more than a single characterization, are the different notions "roughly equivalent"? (For example, in [63] we consider two different notions of *approximately modular* functions, and show, building upon [96, 37], that the two notions are roughly equivalent). We seek to study this aspect of robustness with respect to major problems in AGT, including welfare and revenue maximization, price of anarchy, existence of market equilibrium and gains from trade in bilateral trading. Essentially, for every property X, one can ask what notions of approximate X capture tractability with respect to the different areas of applications, and which parameters explain best the deterioration in approximation guarantees. In the study of proximity to distributions, we will use the distance metrics given in [81].

Open problem 14. *The gross substitutes (GS) class is allegedly the frontier beyond which (1) welfare cannot be maximized in polynomial time*[1]*, (2) a demand query cannot be resolved in polynomial time, and (3) market equilibria do not necessarily exist. What notions of approximate GS capture tractability with respect to these applications?*

Open problem 15. *What notions of limited complementarity are useful for revenue maximization for combinatorial auctions? Can the "additivity over limited complementarities" considered in [59] be extended to "subadditivity over limited complementarities"?*

Open problem 16. *What notions of limited complementarity are useful for prophet and secretary problems? Some preliminary work has been done in [74], but settings such as [144] still await generalizations to valuations that are not complement free.*

Open problem 17. *To what extent competition complexity results that apply to* regular/MHR distributions *continue to hold approximately when faced with distributions that are approximately regular/MHR?*

Objective 5: Beyond worst case inputs

While classical theory in AGT is too "optimistic" in the sense that it abstracts away many aspects that exist in real-world applications, it is too "pessimistic" in its worst-case analysis approach. Indeed, in traditional AGT, a mechanism's performance is measured with respect to its worst-case inputs. Worst-case analysis has its advantages: it is often more tractable analytically, and if good results apply, they provide very strong guarantees. However, worst-case scenarios do not necessarily resemble real-world instances, and may consequently lead to bad advice for practical applications. Indeed, some problems which are provably hard in theory are regularly solved quite well in practice. Perhaps the most famous example is that of linear programming. Based on worst-case analysis, the ellipsoid algorithm should be preferred over the simplex algorithm, as the former runs in polynomial time [102], while the latter runs in exponential time [103]. However, the simplex algorithms runs most commonly in *linear* time, and is the most commonly used method for solving linear programs in practice.

A more relaxed analysis approach is termed *beyond worst-case analysis*. This approach refers to a collection of alternative models that are more applicable to real-world scenarios. These models suggest various combinations of random and adversarial models. For example, in semi-random models (e.g., [26, 65]), an initial input is chosen randomly, but is then modified by a (precisely defined) adversary. Similarly, in the celebrated *smoothed analysis* (e.g., [152]), the adversary first chooses an input, which is then modified according to some random perturbation model. These models potentially avoid "fragile" worst-case inputs; they distinguish truly discouraging results from deceptive discouraging results. Various examples of algorithm analysis that takes this approach appear in the recent survey of Roughgarden [139].

AGT is full of discouraging worst-case lower bounds and hardness results. For example, optimal auction theory predicts the need for complex mechanisms for revenue maximization, even for a simple single-buyer scenario [119]. If the buyer's values for the items are independent, then there are numerous results exhibiting simple mechanisms that give good approximation guarantees [16, 34, 35, 59, 89, 145]. Correlated valuations, however, lead to extremely discouraging results, exhibited by carefully crafted inputs [88, 28], showing that the revenue of a simple mechanism can be arbitrarily small relative to the optimal one. This observation has led the community to (almost fully) abandon scenarios with correlated valuations, even though it is clear that real-world valuations are correlated.

As another example, correlated valuations lead to discouraging results with respect to the Bayesian price of anarchy of simultaneous item auctions. While the BPoA is shown to be a small constant even for markets with subadditive valuations over independent items [70], the BPoA is polynomial in the number of agents if valuations are correlated, even for simple markets with unit-demand valuations [136].

---

[1] For example, for the more general family of *submodular valuations*, welfare maximization is computationally hard [79]

As another example, AGT has traditionally concentrated on either complement-free valuations, or the entire set of monotone, non-negative valuations. Going beyond complement-free valuations often lead to strongly discouraging results. Quoting from [2]: "*Complements between goods where one good takes on added value in the presence of another have been a thorn in the side of algorithmic mechanism designers*." However, in real life, complementarities among items are ubiquitous, ranging from toy examples such as coffee and cookies, to prominent settings such as FCC spectrum auctions, takeoff and landing slots, and course selection. Recent work on limited complementarity [62, 64, 59, 2] opens up a wide terrain of unexplored problems appealing to practical scenarios (see more on this in Objective 4).

Yet another example concerns arrival order models. Driven in part by their connection to posted price mechanisms, prophet inequalities [146] have been established for a variety of allocation problems [51, 35, 107, 71] for welfare as well as revenue approximation. Preliminary results establish pricing schemes that give good guarantees for arbitrary arrival order in matching markets [41, 55] and in markets with more complex valuations, such as submodular valuations [71]. However, it remains open whether these strong results extend beyond submodular (or XOS) valuations.

In light of the above discussion, it is of great interest to develop models that go beyond worst-case inputs, be it worst correlation, worst complementarity, or worst arrival order.

Major research goal: Develop *beyond worst-case* input models of "typical instances" suitable for AGT application domains. Use these models to explain the empirical success of certain mechanisms (e.g., generalized second price auction [159, 53], simultaneous item auctions [39], and posted price mechanisms), as well as for predicting the performance of other mechanisms. I seek to employ this approach with respect to: (1) existence of market equilibria, (2) welfare and revenue approximation provided by simple mechanisms, (3) performance of pricing schemes in complex markets, and related prophet inequality setting. To this end, I will use *smoothed analysis* (see above) and *parameterized analysis*, where I will impose natural *parameterization* on the inputs, and conduct my analysis with respect to the chosen parameters. This approach leads to more fine-grained analysis of algorithms and mechanisms than traditional worst-case analysis.

The discussion above suggests to apply semi-random models, smoothed analysis and parameterized analysis with respect to the many discouraging results in AGT theory. Such analysis will help distinguish truly discouraging results from deceptive discouraging results, and settle the gap that sometimes exists between bad prediction in theory vs. good performance in practice.

Open problem 18. *For each one of the following problems: revenue maximization under correlated values, Bayesian price of anarchy of simultaneous item auctions under correlated values, existence of market equilibrium for values beyond GS, are the discouraging results fragile in the sense of smoothed analysis, or would they persist random perturbations?*

Open problem 19. *It is well known by now that pricing mechanisms are closely related to the framework of prophet inequality. One of the main open problems in pricing is whether there exist item prices that guarantee constant-factor welfare approximation for sequential arrival of agents with subadditive valuations. This remains a daunting open problem under worst-case arrival order. Would more relaxed arrival order models lead to positive results?*

Open problem 20. *In what scenarios can online algorithms be converted into dynamic posted price mechanisms, and what are reasonable arrival models for the study of this problem? Can the impossibility result about makespan minimization for unrelated machines ([66]) be circumvented by more relaxed arrival models?*

Open problem 21. *Consider the online weighted bipartite matching. While for unweighted graphs, the famous algorithm of [98] obtains a constant factor* $(1 - 1/e)$ *approximation, in the weighted case the performance of any online algorithm may be arbitrarily bad. On the other hand, Kesselheim et al. [101] show that under a random arrival order, one can obtain* $1/e$ *approximation. Moreover, it was recently shown that this result can be achieved even with a truthful mechanism [131]. Is there a semi-random arrival model that preserves the good guarantees?*

Type of Research and Methodology

The proposed work is mostly theoretical in nature, but is highly motivated and inspired by practical applications. My goal is to develop mathematical foundations for AGT theory that is more nuanced and flexible, and can be tuned to real-world considerations, and prove mathematical theorems within these models. Obviously, I will exert significant effort to ensure that the mathematical models we study will be thoroughly justified, in order that the mathematical results be relevant and meaningful. The proposed project is **interdisciplinary** in nature. Therefore, in overcoming the daunting challenges involved in this project, I will derive theory and mathematical tools from game theory, mechanism design, and theoretical computer science (including computation and communication complexity, graph theory and probability theory).

Some areas suggested in this proposal are more mature than others. In areas that are less explored, the challenges will be much greater, starting by devising an informative model that captures the essence of the problem, followed by developing novel techniques for providing mathematical results within the new models. In the more mature areas, where some research has already been done, our focus would be mostly in studying generalizations, implications, boundaries and limitations, and studying the connections between the different parts.

Work Plan

Our research agenda is structured in five interleaved and complementing objectives, which lends itself neatly to parallel work of many junior researchers working under my guidance and in collaboration with me, each handling an angle of the big picture. Ph.D. students will focus, under my guidance, on a certain class of problems, corresponding to one of the objectives, and understanding the relations between them. M.Sc. students will be guided to study and analyze more specific problems. Collaboration with experienced post-docs will deal with larger classes of scenarios and try to identify further key modeling decisions. The nature of work in my lab has always been exceptionally collaborative, with more senior students or postdocs helping in the guidance of more junior ones, and this will be the case for this project as well. I have in mind small groups of people working on particular problems, but sharing models and insights with one another, complementing each others' skills and particular strengths. Such a collaboration is crucial for the success of this project, due to the strong connections between its different components. Only when viewed together, they will lead to the desired goal of developing a framework for robust AGT theory that will inform the design of practical applications.

High Risk / High Gain

The potential gains of the project are dramatic. Following the growth and proliferation of electronic marketplaces and social applications, the theory of AGT is likely to find increasingly widespread applications in the real world. The state of the art of the AGT theory still lags far behind full applicability in many real Internet applications, largely due to inflated model simplification, overly strong assumptions, and disregard of crucial real-world considerations, such as behavioral ones. This oversimplification may deem the theory non-robust and consequently inapplicable for some real-world scenarios. The approach we outline here considers powerful models of more robust AGT theory, which rely on multiple novel insights. The implications of a robust AGT theory are significant: improving the stability and efficiency of commerce and social applications. From a scientific perspective, we propose to advance the theory beyond the state of the art using novel approaches and techniques.

In terms of risk, given the impressive progress on the topic of this research over the past five years, it is not a priori clear that much improvement can still be made. Indeed, as we have mentioned, new models were developed, as well as insights and techniques. Further improvements will require more significant scientific breakthroughs, and will require harder work. Indeed, overcoming the daunting challenges involved in departure from traditional models is a non-trivial task and will require new ideas. Furthermore, it is unclear a priori that the new models we will devise would be tractable analytically. It is an art to come up with models that are close enough to reality on the one hand, yet amenable to mathematical analysis. Therefore, the specific theoretical approaches proposed here might not yield the desired progress, in which case, I will have to try different approaches. Moreover, a specific direction might need further restrictions to be tractable, which might lead to an iterative process.

However, in order to reduce the risk of insufficient output, we have taken a proactive approach, and as mentioned above, we have already carried out some preliminary research with respect to many of our objectives in preparation for this grant proposal. We believe that our years of experience with this type of work and the impressive preliminary results we already obtained validate that the approaches that we propose have great potential, and that advancing the theory of robust AGT is a possible task. Clearly, the payoff is high enough to make the required hard work more than worthwhile.

Section c. Resources

I am a faculty member of the Blavatnik School of Computer Science at Tel Aviv University. The School of Computer Science is an international leader in CS research, comparable to the best institutions in Europe. Many of its faculty members are internationally renowned experts in theoretical computer science, algorithms, algorithmic game theory, cryptography, computational geometry, bioinformatics, logic, and programming languages. Being located in Tel-Aviv, the heart of Israel's thriving high-tech industry, the school has a rich interaction with industry, as evident from its industrial affiliates program, including Google, Yahoo, Microsoft, IBM, CISCO, Mercury, Checkpoint and Quallcomm. The school enjoys an outstanding international reputation, carries out an extensive research program, with a mix of theoretical and applied areas, and includes winners of prestigious international awards. Being surrounded by a strong group of senior faculty members with close research interests will be very fruitful for the project.

The established PhD and MSc programs of the school have very high standards, attracting and training the most talented students in the country, thus providing a large and well-trained student body from which excellent PhD and masters students can be hired. Faculty members regularly teach a graduate course of their choosing related to their research interests, giving easy access to students. In addition, infrastructure to accommodate postdocs is well developed. Tel Aviv University and the School of Computer Science in it can and will provide the necessary infrastructure (seminar rooms, office space) as well as equipment (computers) for the planned research group. In addition, a competent administrative team assists the researchers and their groups with logistic support, helping with hiring, travel and conferences organization and other secretarial needs.

Budget Justification: The majority of the budget is requested for supporting research personnel. Indeed, it is one of our explicit goals to train high quality AGT researchers, who can continue to contribute to academic research and to research labs in the industry.

Personnel. The core of the research team will consist of the PI, MSc and PhD students, and postdoctoral associates. All recruited personnel will be required to have strong background in algorithms, game theory and algorithmic game theory. Luckily, students with such qualifications are readily available at our institute.

1. **PI:** I am fully committed to the project and will dedicate 60% of my time to the project.
2. **Postdoctoral associates:** Two postdoc fellows are expected to be hired each year. Both Postdocs will come with a strong background in algorithmic game theory. One of them will be familiar with techniques related to "beyond worst case" analysis of algorithms. I would expect these bright young researchers to undertake, in collaboration with me and under my guidance, some of the toughest conceptual challenges in this proposal.
3. **Research students:** Three PhD students are expected to work on the project throughout its lifetime. Two of them will be recruited in the first year and the additional one in the second year. The PhD students will work on specific challenges from those suggested in this proposal. In addition, three MSc students are expected to work on concrete technical questions in the more explored areas.

4. **Project admin:** A part-time administrative assistant will manage organizational issues including personnel, salaries, conference organization, orders, etc.

Travel. Active participation in seminars, workshops and conferences take a large part of the research progress and its dissemination. Therefore, the budget will cover the travel costs for myself and my team to attend scientific conferences and seminars. Having students interact with world class researchers is also a significant boost to their academic education. I will therefore actively invite and support visits by outstanding researchers from abroad to my EC Lab@TAU for short or long-term research periods. Prospective invitees include close collaborator of mine, including Aviad Rubinstein (Stanford), Paul Duetting (London School of Economics), Thomas Kesselheim (U. of Bonn), Matt Weinberg (Princeton), Anna Karlin (U. Washington), and Ariel Procaccia (CMU).

Workshops. I plan to hold two international workshops on topics of relevance to RAGT. The aim of these workshops will be to study areas related to RAGT in depth, to promote research by others in this area, and to disseminate the results of RAGT. At least one of these workshops is intended to follow the successful format of the YoungEC Workshop we held in 2017, which brought together a small number of extremely established central figures in the field together with a larger number of the brightest rising stars worldwide in the area.

Equipment. This will cover the costs of laptops, tablets, and desktops for my groups, as well as dedicated software (such as maple, mathematica) printer, scanner, projector, etc.

Publication Costs. This will cover the costs of open access publications, graphic design, language editing, poster production and printing for conferences etc.

Financial audit. This will cover the costs of an audit certificate.

Other Direct Costs. Additional costs are for maintaining our web site https://en-exact-sciences.tau.ac.il/EC.

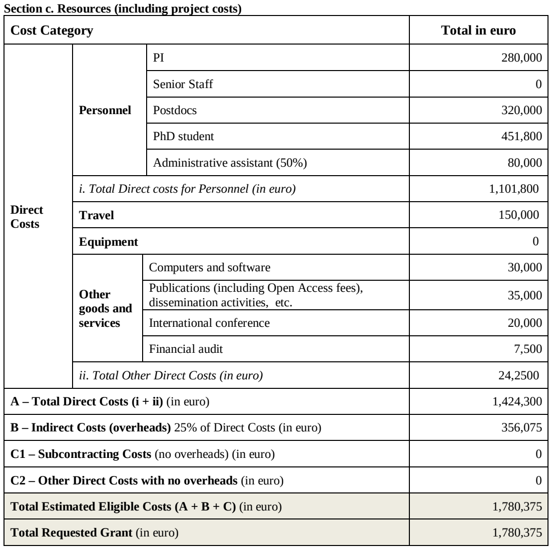

Section c. Resources (including project costs)

Cost Category			**Total in euro**
Direct Costs	**Personnel**	PI	280,000
		Senior Staff	0
		Postdocs	320,000
		PhD student	451,800
		Administrative assistant (50%)	80,000
	i. Total Direct costs for Personnel (in euro)		1,101,800
	Travel		150,000
	Equipment		0
	Other goods and services	Computers and software	30,000
		Publications (including Open Access fees), dissemination activities, etc.	35,000
		International conference	20,000
		Financial audit	7,500
	ii. Total Other Direct Costs (in euro)		24,2500
A – Total Direct Costs (i + ii) (in euro)			1,424,300
B – Indirect Costs (overheads) 25% of Direct Costs (in euro)			356,075
C1 – Subcontracting Costs (no overheads) (in euro)			0
C2 – Other Direct Costs with no overheads (in euro)			0
Total Estimated Eligible Costs (A + B + C) (in euro)			1,780,375
Total Requested Grant (in euro)			1,780,375

Duration of the project in months	60 months
The % of working time the PI dedicates to the project over the period of the grant	60%

References

[1] Melika Abolhassani, Soheil Ehsani, Hossein Esfandiari, MohammadTaghi Hajiaghayi, Robert D. Kleinberg, and Brendan Lucier. Beating 1-1/e for ordered prophets. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 61–71, 2017.

[2] Ittai Abraham, Moshe Babaioff, Shaddin Dughmi, and Tim Roughgarden. Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 3–16. ACM, 2012.

[3] Mohammad Akbarpour, Suraj Malladi, and Amin Saberi. Just a few seeds more: Value of network information for diffusion. Research papers, Stanford University, Graduate School of Business, 2018.

[4] George A Akerlof. Procrastination and obedience. *The American Economic Review*, 81(2):1–19, 1991.

[5] Susanne Albers and Dennis Kraft. Motivating time-inconsistent agents: A computational approach. In *Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings*, pages 309–323, 2016.

[6] Susanne Albers and Dennis Kraft. On the value of penalties in time-inconsistent planning. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, pages 10:1–10:12, 2017.

[7] Susanne Albers and Dennis Kraft. The price of uncertainty in present-biased planning. In *Web and Internet Economics - 13th International Conference, WINE 2017, Bangalore, India, December 17-20, 2017, Proceedings*, pages 325–339, 2017.

[8] Nick Arnosti, Marissa Beck, and Paul Milgrom. Adverse selection and auction design for internet display advertising. *American Economic Review*, 106(10):2852–66, 2016.

[9] Susan Athey and Ilya Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.

[10] Robert Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[11] Lawrence M Ausubel et al. A generalized vickrey auction. *Econometrica*, 1999.

[12] Moshe Babaioff, Yang Cai, Yannai A. Gonczarowski, and Mingfei Zhao. The best of both worlds: Asymptotically efficient mechanisms with a guarantee on the expected gains-from-trade. In *EC*, page 373. ACM, 2018.

[13] Moshe Babaioff, Shahar Dobzinski, and Sigal Oren. Combinatorial auctions with endowment effect. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 73–90. ACM, 2018.

[14] Moshe Babaioff, Shahar Dobzinski, and Sigal Oren. Combinatorial auctions with endowment effect. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 73–90, 2018.

[15] Moshe Babaioff, Yannai A. Gonczarowski, Yishay Mansour, and Shay Moran. Are two (samples) really better than one? In *EC*, page 175. ACM, 2018.

[16] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 21–30. IEEE, 2014.

[17] Yakov Babichenko, Shahar Dobzinski, and Noam Nisan. The communication complexity of local search. *CoRR*, abs/1804.02676, 2018.

[18] Eric Balkanski, Aviad Rubinstein, and Yaron Singer. The limitations of optimization from samples. Working paper, 2016.

[19] Dirk Bergemann and Stephen Morris. Robust mechanism design. *Econometrica*, 73(6):1771–1813, 2005.

[20] Dirk Bergemann and Karl Schlag. Robust monopoly pricing. *Journal of Economic Theory*, 146(6):2527–2543, 2011.

[21] Dirk Bergemann and Juuso Valimaki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.

[22] Dimitris Bertsimas and Aur/'elie Thiele. *Robust and Data-Driven Optimization: Modern Decision Making Under Uncertainty*, chapter 5, pages 95–122. INFORMS PubsOnline, 2014. TutORials in Operations Research.

[23] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 700–709, 2011.

[24] Kshipra Bhawalkar and Tim Roughgarden. Simultaneous single-item auctions. In *Internet and Network Economics - 8th International Workshop, WINE 2012, Liverpool, UK, December 10-12, 2012. Proceedings*, pages 337–349, 2012.

[25] Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 373–382, 2008.

[26] Avrim Blum and Joel Spencer. Coloring random and semi-random k-colorable graphs. *J. Algorithms*, 19(2):204–234, 1995.

[27] Liad Blumrosen and Noam Nisan. Combinatorial auctions (a survey). In *N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, Algorithmic Game Theory. Cambridge University Press*, 2007.

[28] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S Matthew Weinberg. Pricing randomized allocations. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 585–597. Society for Industrial and Applied Mathematics, 2010.

[29] David Buchfuhrer, Michael Schapira, and Yaron Singer. Computation and incentives in combinatorial public projects. In *Proceedings 11th ACM Conference on Electronic Commerce (EC-2010), Cambridge, Massachusetts, USA, June 7-11, 2010*, pages 33–42, 2010.

[30] Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–94, March 1996.

[31] Yang Cai and Constantinos Daskalakis. Learning multi-item auctions with (or without) samples. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 516–527, 2017.

[32] Shuchi Chawla, Hu Fu, and Anna Karlin. Approximate revenue maximization in interdependent value settings. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 277–294, New York, NY, USA, 2014. ACM.

[33] Shuchi Chawla, Kira Goldner, J. Benjamin Miller, and Emmanouil Pountourakis. Revenue maximization with an uncertainty-averse buyer. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2050–2068, 2018.

[34] Shuchi Chawla, Jason D Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 243–251. ACM, 2007.

[35] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320. ACM, 2010.

[36] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *J. ACM*, 56(3), 2009.

[37] Flavio Chierichetti, Abhimanyu Das, Anirban Dasgupta, and Ravi Kumar. Approximate modularity. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 1143–1162, 2015.

[38] George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. In *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part I: Tack A: Algorithms, Automata, Complexity, and Games*, pages 820–832, 2008.

[39] George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. *J. ACM*, 63(2):11:1–11:19, 2016.

[40] Edward Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

[41] Vincent Cohen-Addad, Alon Eden, Michal Feldman, and Amos Fiat. The invisible hand of dynamic market pricing. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 383–400, 2016.

[42] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 243–252, 2014.

[43] José R. Correa, Paul Dütting, Felix A. Fischer, and Kevin Schewior. Prophet inequalities for independent random variables from an unknown distribution. *CoRR*, abs/1811.06114, 2018.

[44] Partha Dasgupta and Eric Maskin. Efficient auctions. *The Quarterly Journal of Economics*, 115(2):341–388, 2000.

[45] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[46] Constantinos Daskalakis and Vasilis Syrgkanis. Learning in auctions: Regret is hard, envy is easy. In *FOCS*, pages 219–228. IEEE Computer Society, 2016.

[47] Nikhil R. Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. *CoRR*, abs/1511.02296, 2015.

[48] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.

[49] Shahar Dobzinski, Michal Feldman, Inbal Talgam-Cohen, and Omri Weinstein. Welfare and revenue guarantees for competitive bundling equilibrium. In *Web and Internet Economics - 11th International Conference, WINE 2015, Amsterdam, The Netherlands, December 9-12, 2015, Proceedings*, pages 300–313, 2015.

[50] Shahar Dobzinski, Hu Fu, and Robert D. Kleinberg. Optimal auctions with correlated bidders are easy. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 129–138, 2011.

[51] Paul Duetting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 540–551, 2017.

[52] David Easley and Arpita Ghosh. Behavioral mechanism design: Optimal crowdsourcing contracts and prospect theory. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 679–696, 2015.

[53] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

[54] Alon Eden, Tomer Ezra, and Michal Feldman. Pricing social goods. In *ESA*, volume 87 of *LIPIcs*, pages 35:1–35:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.

[55] Alon Eden, Uriel Feige, and Michal Feldman. Max-min greedy matching. *CoRR*, abs/1803.05501, 2018.

[56] Alon Eden, Michal Feldman, Amos Fiat, and Kira Goldner. Interdependent values without single-crossing. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, page 369, 2018.

[57] Alon Eden, Michal Feldman, Amos Fiat, Kira Goldner, and Anna Karlin. Interdependent values with submodular valuations. In *Working paper*, 2019.

[58] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. The competition complexity of auctions: A bulow-klemperer result for multi-dimensional bidders. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, page 343, 2017.

[59] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S Matthew Weinberg. A simple and approximately optimal mechanism for a buyer with complements. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 323–323. ACM, 2017.

[60] Tomer Ezra, Michal Feldman, and Ophir Friedler. Endowment effect, stability and welfare. In *Working paper*, 2019.

[61] Tomer Ezra, Michal Feldman, Eric Neyman, Inbal Talgam-Cohen, and S. Matthew Weinberg. Complement-free couples must communicate: A hardness result for two-player combinatorial auctions. *CoRR*, abs/1811.09871, 2018.

[62] Uriel Feige, Michal Feldman, Nicole Immorlica, Rani Izsak, Brendan Lucier, and Vasilis Syrgkanis. A unifying hierarchy of valuations with complements and substitutes. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 872–878. AAAI Press, 2015.

[63] Uriel Feige, Michal Feldman, and Inbal Talgam-Cohen. Approximate modularity revisited. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 1028–1041, 2017.

[64] Uriel Feige and Rani Izsak. Welfare maximization and the supermodular degree. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 247–256. ACM, 2013.

[65] Uriel Feige and Joe Kilian. Heuristics for semirandom graph problems. *J. Comput. Syst. Sci.*, 63(4):639–671, December 2001.

[66] Michal Feldman, Amos Fiat, and Alan Roytman. Makespan minimization via posted prices. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, pages 405–422, 2017.

[67] Michal Feldman, Ophir Friedler, Jamie Morgenstern, and Guy Reiner. Simple mechanisms for agents with complements. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 251–267. ACM, 2016.

[68] Michal Feldman, Ophir Friedler, and Aviad Rubinstein. 99% revenue via enhanced competition. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 443–460, 2018.

[69] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 201–210, 2013.

[70] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 201–210. ACM, 2013.

[71] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 123–135, 2015.

[72] Michal Feldman, Nick Gravin, and Brendan Lucier. On welfare approximation and stable pricing. *CoRR*, abs/1511.02399, 2015.

[73] Michal Feldman, Brendan Lucier, and Noam Nisan. Correlated and coarse equilibria of single-item auctions. In *Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings*, pages 131–144, 2016.

[74] Moran Feldman and Rani Izsak. Building a good team: Secretary problems and the supermodular degree. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1651–1670. SIAM, 2017.

[75] Amos Fiat, Kira Goldner, Anna R. Karlin, and Elias Koutsoupias. The fedex problem. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 21–22, 2016.

[76] Amos Fiat and Christos H. Papadimitriou. When the players are not expectation maximizers. In *SAGT*, volume 6386 of *Lecture Notes in Computer Science*, pages 1–14. Springer, 2010.

[77] Hu Fu, Robert Kleinberg, and Ron Lavi. Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC 2012, Valencia, Spain, June 4-8, 2012*, page 586, 2012.

[78] Alex Gershkov and Benny Moldovanu. Efficient sequential assignment with incomplete information. *Games and Economic Behavior*, 68(1):144–154, 2010.

[79] Shayan Oveis Gharan and Jan Vondrák. Submodular maximization by simulated annealing. In *SODA*, pages 1098–1116. SIAM, 2011.

[80] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.

[81] Alison L. Gibbs and Francis Edward Su. On choosing and bounding probability metrics. *INTERNAT. STATIST. REV.*, pages 419–435, 2002.

[82] Yannai A. Gonczarowski and Noam Nisan. Efficient empirical revenue maximization in single-parameter auction environments. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 856–868, 2017.

[83] Yannai A. Gonczarowski and S. Matthew Weinberg. The sample complexity of up-to-ϵ multi-dimensional revenue maximization. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 416–426, 2018.

[84] Nick Gravin, Nicole Immorlica, Brendan Lucier, and Emmanouil Pountourakis. Procrastination with variable present bias. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, page 361, 2016.

[85] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–31, 1973.

[86] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95–124, 1999.

[87] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.

[88] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 565–566. ACM, 2013.

[89] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. *Journal of Economic Theory*, 172:313–347, 2017.

[90] Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In *Proceedings 12th ACM Conference on Electronic Commerce (EC-2011), San Jose, CA, USA, June 5-9, 2011*, pages 295–296, 2011.

[91] Avinatan Hassidim and Yaron Singer. Submodular optimization under noise. In *COLT*, volume 65 of *Proceedings of Machine Learning Research*, pages 1069–1122. PMLR, 2017.

[92] Satoru Iwata, Lisa Fleischer, and Satoru Fujishige. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *J. ACM*, 48(4):761–777, 2001.

[93] Philippe Jehiel and Benny Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69(5):1237–1259, 2001.

[94] Daniel Kahneman and Amos Tversky. Choices, values, and frames. In *Handbook of the Fundamentals of Financial Decision Making: Part I*, pages 269–278. World Scientific, 2013.

[95] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*, pages 99–127. World Scientific, 2013.

[96] N.J. Kalton and James W. Roberts. Uniformly exhaustive submeasures and nearly additive set functions. *Transactions of the American Mathematical Society*, 278(2):803–816, 1983.

[97] Bala Kalyanasundaram and Kirk Pruhs. Speed is as powerful as clairvoyance. *J. ACM*, 47(4):617–643, 2000.

[98] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 352–358, 1990.

[99] Alexander S Kelso Jr and Vincent P Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society*, pages 1483–1504, 1982.

[100] Thomas Kesselheim and Bojana Kodric. Price of anarchy for mechanisms with risk-averse agents. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, pages 155:1–155:14, 2018.

[101] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *Algorithms - ESA 2013 - 21st Annual European Symposium, Sophia Antipolis, France, September 2-4, 2013. Proceedings*, pages 589–600, 2013.

[102] L. G. Khachiyan. A polynomial algorithm in linear programming. *Soviet Mathematics Doklady*, 20(1):191194, 1979.

[103] V. Klee and G. J. Minty. How good is the simplex algorithm? *In O. Shisha, editor, Inequalities III, Academic Press Inc., New York*, page 159175, 1972.

[104] Jon M. Kleinberg and Sigal Oren. Time-inconsistent planning: a computational problem in behavioral economics. In *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, pages 547–564, 2014.

[105] Jon M. Kleinberg, Sigal Oren, and Manish Raghavan. Planning problems for sophisticated agents with present bias. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 343–360, 2016.

[106] Jon M. Kleinberg, Sigal Oren, and Manish Raghavan. Planning with multiple biases. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, pages 567–584, 2017.

[107] Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities. In *STOC*, pages 123–136. ACM, 2012.

[108] Jack L Knetsch. The endowment effect and evidence of nonreversible indifference curves. *The american Economic review*, 79(5):1277–1284, 1989.

[109] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 4-6, 1999, Proceedings*, pages 404–413, 1999.

[110] Andreas Krause and Volkan Cevher. Submodular dictionary selection for sparse representation. In *Proceedings of the 27th*, pages 567–574, 2010.

[111] Renato Paes Leme. Gross substitutability: An algorithmic survey. *Games and Economic Behavior*, 106:294–316, 2017.

[112] Siqi Liu and Christos-Alexandros Psomas. On the competition complexity of dynamic mechanism design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2008–2025, 2018.

[113] Ausubel Lawrence M and Paul Milgrom. Ascending auctions with package bidding. *The B.E. Journal of Theoretical Economics*, 1(1):1–44, 2002.

[114] Eric Maskin. Auctions and privatization. *Privatization, H. Siebert, ed. (Institut fur Weltwirtschaften der Universitat Kiel: 1992)*, pages 115–136, 1992.

[115] Paul Milgrom. Discovering prices: Auction design in markets with complex constraints. In *Columbia University Press*, 2017.

[116] Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122, 1982.

[117] Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *Proceedings of the 29th Conference on Learning Theory, COLT 2016, New York, USA, June 23-26, 2016*, pages 1298–1318, 2016.

[118] Kazuo Murota. Valuated matroid intersection I: optimality criteria. *SIAM J. Discrete Math.*, 9(4):545–561, 1996.

[119] Roger B. Myerson. Optimal auction design. *Math. Oper. Res.*, 6(1):58–73, February 1981.

[120] Noam Nisan. Algorithmic mechanism design through the lens of multi-unit auctions. In H. Peyton Young and Shmuel Zamir, editors, *Handbook of Game Theory with Economic Applications, Volume 4*, chapter 9, pages 477–515. Elsevier, 2015.

[121] Noam Nisan and Amir Ronen. Algorithmic mechanism design (extended abstract). In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, STOC '99, pages 129–140, New York, NY, USA, 1999. ACM.

[122] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

[123] Y. Shoham P. Cramton and R. Steinberg. Combinatorial auctions, mit press. 2006.

[124] Mallesh M. Pai and Rakesh Vohra. Optimal dynamic auctions and simple index rules. *Math. Oper. Res.*, 38(4):682–697, 2013.

[125] Christos Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, STOC '01, pages 749–753, New York, NY, USA, 2001. ACM.

[126] Christos H. Papadimitriou and George Pierrakos. Optimal deterministic auctions with correlated priors. *Games and Economic Behavior*, 92:430–454, 2015.

[127] Christos H. Papadimitriou, George Pierrakos, Christos-Alexandros Psomas, and Aviad Rubinstein. On the complexity of dynamic mechanism design. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1458–1475, 2016.

[128] David C. Parkes and Satinder P. Singh. An mdp-based approach to online mechanism design. In *Advances in Neural Information Processing Systems 16 [Neural Information Processing Systems, NIPS 2003, December 8-13, 2003, Vancouver and Whistler, British Columbia, Canada]*, pages 791–798, 2003.

[129] Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653, 2014.

[130] Cynthia A. Phillips, Clifford Stein, Eric Torng, and Joel Wein. Optimal time-critical scheduling via resource augmentation. *Algorithmica*, 32(2):163–200, 2002.

[131] Rebecca Reiffenhauser. An optimal truthful mechanism for the online weighted bipartite matching problem. In *SODA*, 2019.

[132] James Rizzo. Newcomb's problem for decision theory and critical theory. *Rethinking Marxism*, 17(3):471–485, 2005.

[133] Kevin Roberts. The characterization of implementable choice rules. *In Jean-Jacques Laffont, editor, Aggregation and Revelation of Preferences. Papers presented at the 1st European Summer Workshop of the Econometric Society. North-Holland.*, pages 321–349, 1979.

[134] Tim Roughgarden. Algorithmic game theory. *Commun. ACM*, 53(7):78–86, July 2010.

[135] Tim Roughgarden. Intrinsic robustness of the price of anarchy. *Commun. ACM*, 55(7):116–123, 2012.

[136] Tim Roughgarden. The price of anarchy in games of incomplete information. In *ACM Conference on Electronic Commerce, EC '12, Valencia, Spain, June 4-8, 2012*, pages 862–879, 2012.

[137] Tim Roughgarden. Approximately optimal mechanism design: Motivation, examples, and lessons learned. *SIGecom Exch.*, 13(2):4–20, January 2015.

[138] Tim Roughgarden. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 1st edition, 2016.

[139] Tim Roughgarden. Beyond worst-case analysis. *Commun. ACM, to appear*, 2019.

[140] Tim Roughgarden, Vasilis Syrgkanis, and Éva Tardos. The price of anarchy in auctions. *J. Artif. Intell. Res.*, 59:59–101, 2017.

[141] Tim Roughgarden and Inbal Talgam-Cohen. Optimal and robust mechanism design with interdependent values. *ACM Trans. Econ. Comput.*, 4(3):18:1–18:34, June 2016.

[142] Tim Roughgarden, Inbal Talgam-Cohen, and Jan Vondrák. When are welfare guarantees robust? In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2017, August 16-18, 2017, Berkeley, CA, USA*, pages 22:1–22:23, 2017.

[143] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.

[144] Aviad Rubinstein and Sahil Singla. Combinatorial prophet inequalities. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1671–1687. SIAM, 2017.

[145] Aviad Rubinstein and S Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. *ACM Transactions on Economics and Computation (TEAC)*, 6(3-4):19, 2018.

[146] E. Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Annals of Probability*, 12:1213–1216, 1984.

[147] William Samuelson and Richard Zeckhauser. Status quo bias in decision making. *Journal of risk and uncertainty*, 1(1):7–59, 1988.

[148] Mark Allen Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

[149] Alexander Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *Journal of Combinatorial Theory*, 80(2):346–355, 2000.

[150] Yaron Singer and Jan Vondrák. Information-theoretic lower bounds for convex optimization with erroneous oracles. In *Proceedings of the 28th*, pages 3204–3212, 2015.

[151] Daniel Dominic Sleator and Robert Endre Tarjan. Amortized efficiency of list update and paging rules. *Commun. ACM*, 28(2):202–208, 1985.

[152] Daniel A. Spielman and Shang hua Teng. Smoothed analysis of algorithms: why the simplex algorithm usually takes polynomial time. In *Journal of the ACM*, pages 296–305, 2001.

[153] Vasilis Syrgkanis and Éva Tardos. Composable and efficient mechanisms. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 211–220, 2013.

[154] Richard Thaler. Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization*, 1(1):39–60, 1980.

[155] Richard Thaler. Behavioral economics: Past, present, and future. *American Economic Review*, 106(7):1577–1600, 2016.

[156] GS Tune. Response preferences: A review of some relevant literature. *Psychological bulletin*, 61(4):286, 1964.

[157] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *science*, 185(4157):1124–1131, 1974.

[158] Amos Tversky and Daniel Kahneman. The framing of decisions and the psychology of choice. *science*, 211(4481):453–458, 1981.

[159] Hal Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.

[160] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

[161] Robert Wilson. Game-theoretic analysis of trading processes. Technical report, STANFORD UNIV CA INST FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES, 1985.

[162] Peyton H. Young. *Strategic learning and its limits*. 2004.