



Engineering the Policy-making Life Cycle

Seventh Framework Programme – Grant Agreement 288147

Deliverable 5.1

Survey of Game Theoretic Tools in Dynamic Environments for Policy Management

Document type:	Deliverable
Dissemination Level:	Public
Editor:	Barry O’Sullivan
Document Version:	1.0.2
Contributing Partners:	(UCC)
Contributing WPs:	WP5
Date of Completion:	31 Jul 2012
Date of Delivery to EC	31 Jul 2012
Number of pages:	27

ABSTRACT

The policy making process by which governments translate political visions and goals into programmes, mechanisms and actions delivers real change with tangible outcomes for society that range from very positive to very negative. The goal of policy makers is to develop policies that are coherent, based on sound information and logic, and endure of long periods. This paper examines the role of managing strategic incentives in policy making so that governing bodies can develop more sound policies. We place an emphasis on economic policy making and describe the boundaries between policy management and implementation via protocols and mechanisms. We survey the field of game theory for existing contributions that are useful to policy makers in public institutions. We outline how strategic decision making by agents participating in a publicly operated protocol or mechanism can obviate the best intentions of mechanism designers by employing self-interested strategies. We then introduce a body of literature from the fields of Economics, Game Theory and Mechanism Design that describes the challenges for policy implementation. Recent advances in these fields have enlightened us as to both the possibilities and impossibilities regarding the implementation of social choice functions. We outline key applications of game theory to policy making decisions in the areas of tax inspection, monetary policy, asset sales and public expenditure. We focus on the most recent developments in Dynamic

Mechanism Design that seek to address scenarios that include uncertainty and stochastic dynamics. Given that policy making is a slow and infrequent process, protocols and mechanisms that implement policy need to be able to adapt to a wide range of scenarios so adaptability is of major interest. We contend that policy objectives for economic transactions can be used to guide implementation and prioritize potentially conflicting objectives.

Copyright © by the ePolicy Consortium

The ePolicy Consortium consists of the following partners: University of Bologna; University College Cork, National University of Ireland, Cork; University of Surrey; INESC Porto, Instituto de Engenharia de Sistemas e Computadores do Porto, Fraunhofer – Gesellschaft zur Foerderung der Angewandten Forschung E.V.; Regione Emila-Romagna; ASTER – Società Consortile per Azioni; Università degli Studi di Ferrara.

Authors of this document:

Alan Holland¹, Barry O'Sullivan¹

¹: Cork Constraint Computation Centre
Department of Computer Science
University College Cork, Ireland.

Contents

1	Introduction	6
2	Game theory and policy-making	7
2.1	Tax Inspection and Competition	7
2.2	Monetary Policy	8
2.3	Auctions for Procurement or Asset Sales	9
2.3.1	English auction.	9
2.3.2	Dutch auction.	9
2.3.3	Sealed-Bid first-price auction.	10
2.3.4	Sealed-Bid Second-Price (Vickrey) auction.	10
3	Modelling Agent Behavior	10
3.1	Game-Theoretic Behavior	11
3.2	Revenue Equivalence and Policy Objectives	12
4	Learning in Games	12
4.1	Fictitious Play	13
4.2	Convergence Properties	13
5	Mechanism Design	15
5.1	Mechanism Design Goals	15
5.2	Mechanism design for economic policy-making	16
5.3	Mechanism Design Formalism	17
5.4	Combinatorial Auction as a Policy Tool	18
6	Example: Efficient Government Grant Auctions for Renewable Device Subsidies	19
6.1	Model	21
6.2	Allocation Algorithms	21
6.2.1	Inefficiency and a Non-monotone Algorithm	21
6.2.2	Efficiency and Monotonicity	22
7	Dynamic games and Stochastic Dynamics	22
7.1	Online VCG	22
7.2	Dyanmic VCG and its Application to Energy Storage Management	23
7.2.1	Example: Commodity Storage	23
7.2.2	The Societal Benefits of a Smarter Regulatory Regime	23
8	Conclusion	24

This page has been intentionally left blank.

1 Introduction

A policy is typically a set of rules that is designed to facilitate the achievement of certain goals or objectives on the part of a country or organization. A policy should aid decision making and should evolve as the objectives change over time. A protocol is more specific than a policy because it defines a set of procedures to be followed for the accomplishment of an identified task. There is a methodology or well defined procedure that controls how tasks are achieved. This well-defined procedure is called as protocol.

The adherence to a protocol associated with the completion of a task brings clarity and certainty to the state of a task. Protocols tend to be observed repeatedly and can be refined based upon delivery against an over-arching policy that itself also changes as an organizations' needs evolve. Protocols are often considered to be an effective way to establish standard and repeatable processes for large organisations whose management requires mechanisms for controlling large groups of individuals. Policies and protocols are thus interwoven management tools with strong dependencies.

The game-theoretic analysis of deliberation and negotiation and the normative theory of deliberative democracy both view contention for resources from different perspectives but have developed in mutual isolation. A study by Landa and Meirowitz [32] confronted the arguments raised by normative theorists opposed to the perceived relevance of underlying assumptions in game-theoretic work. They found that the game-theoretic approach is particularly well-suited for providing insights about the feasibility of deliberative institutions and practices. Game theory is improving our understanding of decision making and, in particular, how economic agents react to a set of rules. The central solution concept surrounds an equilibrium in which agents do not have an incentive to unilaterally deviate from a specific action [41]. Recent research has extended the range of solution concepts to address broader environments that include uncertainty, stochastic dynamics and other complicating factors. We survey extant work related to game theory as a framework for assessing the efficacy of policies and protocols.

In particular, the field of mechanism design is a branch of economics whose primary application is the design of protocols for the sale or procurement of items. It is particularly relevant to policy making because it concerns the design of protocols for implementing policy objectives in specific settings. The Nobel Prize for Economics was awarded to Leonid Hurwicz, Eric Maskin and Roger Myerson in 2007 for having laid the foundations of mechanism design theory. This brought recognition to the founders of a field that has contributed enormously to policy-making and governance [33]. Leonid Hurwicz initiated research in this field in the 1960s when he examined how a planner should reach a decision when the quality of the decision relies on information spread among numerous people. Mechanism design theory formulates this problem mathematically and studies properties of allocation and payment rules. Among Hurwicz's key insights is the idea that the self-interested agents must find it in their interest to reveal private information. This insight informed a contemporary intellectual debate concerning the relative merits of Capitalism and Socialism. It helped governments understand the importance of incentives, private information and to consider effective regulation of capitalist economies [33]. Key contributions from Maskin and Myerson extended the theory further in the 1970s and 1980s. Mechanism design theory now provides a general framework to study collective decision problems and many specific sub-fields have been studied. We examine the

emergence of mechanism design as a rapidly evolving economic tool that aims to improve the effectiveness of economic protocols given a model of game theoretic rational decision making by agents. We shortlist relevant advances in the field that may support policy making capabilities and we focus upon dynamic mechanisms with uncertainty.

2 Game theory and policy-making

Policy makers frequently make decisions regarding income gathering and expenditure. The most complex challenges for policy makers include tax compliance management and efficient expenditure of funds to support social objectives. We first need to understand the imperatives of the players in a setting ruled by policy makers. We use Game Theory to model their actions and reactions in this environment.

Game theory is a mathematical theory of strategic interaction where multiple players must make decisions that may affect the interests of other players. An auction is an example of a game in which bidders are competing agents, each of whom is seeking to maximize their utility [27]. The bid-taker sets the rules for the game in such a way as to achieve his objective, which is often revenue maximization but may also be the fulfilment of some social objective. The bidders, on the other hand, will strategize so that their expected surplus is maximized [31, 38, 39, 43].

A strategic equilibrium is a profile, or combination, of strategies such that if other players conform to the equilibrium strategies (i.e. other players are rational), no player has an incentive to unilaterally deviate from his equilibrium strategy [56]. Game theory provides several solution concepts to compute the outcome of a game with self-interested agents. A solution concept is used to predict the strategies agents will choose in order to maximize their utility, thus determining an equilibrium position for the game. These concepts assume knowledge about agent preferences, rationality, and shared information about one other. The best known concept is that of a Nash equilibrium, which states that in equilibrium every agent will select a utility-maximizing strategy given the strategy of every other agent [41]. A Nash-equilibrium is self-referential and constitutes a profile of strategies that form “optimal reactions” to other agents’ “optimal reactions”. Nash equilibrium is the pure form of the basic concept of strategic equilibrium; as such, it is useful mainly in normal form games with complete information. When allowing for randomized strategies, at least one Nash equilibrium exists in any game with regular payoff functions [56]. A game may possess one or more Nash equilibria. We discuss two example auction forms below and outline their equilibrium bidding strategies.

2.1 Tax Inspection and Competition

The Inspection Game is a multistage game between a smuggler and a customs inspector. It was first identified by Dresher [13] and studied by Maschler [36]. Sakaguchi then examined a repeated game scenario [51]. Ferguson and Melolidakis also studied a related problem of inspecting an industrial manufacturer for illegal dumping of a toxic waste inventory [15]. Avenhaus et

al. conducted a survey of other research in inspection games [3].

The basic premise is that a smuggler (or tax evader) can choose when to commit an unlawful action and the inspector must choose how much effort to expend inspecting would-be criminals. This investment decision is taken with a view to the benefits the smuggler attains for the action, the level of deterrent once caught and also their opponents view on inspection rates. Likewise, the smuggler conditions his action on the incentives for gain, penalties for capture and inspection rates. The body of research mentioned above is useful to policy makers when deciding on budgetary levels for inspection tasks.

Greenberg proposed an optimal auditing scheme for tax collection authorities faced with certain tax and penalty functions and a finite budget. The problem parameters were the probability of being audited and the expected payoffs for reporting truthfully and cheating, respectively. It is then shown that there is a choice of these parameters so that in equilibrium, the percentage of individuals that cheat is arbitrarily small [22]. Tax compliance within the confines of an enclosed game is well studied.

Corporation Tax and International Competition: Game theory is also useful to policy makers when choosing tax rates. It can be viewed in terms of both internal compliance and also international competition. In particular, low corporation tax can be viewed as a tactical tool to attract foreign direct investment from footloose industries. The benefit for countries accrues from income tax paid by employees of these companies.

Fourçan and Warin used a game theoretical approach to analyze tax harmonization, or competition, in the European Monetary Union [16]. They found that without harmonization, free-riding behaviors may appear, leading to a sub-optimal tax equilibrium. Tax competition may also create budgetary problems and the objective of a balanced budget may not be attained. But national tax autonomy has one main advantage: as monetary policy is federalized, and as fiscal policy is constrained by the Stability and Growth Pact, taxation becomes the last macroeconomic instrument of individual governments to counteract asymmetric shocks. The authors argue against other commentary criticizing tax autonomy and free-riding outcomes. In equilibrium, the competition could shift to the lowest tax rate of all countries, in theory forcing countries to reduce their spending. Fourçan and Warin's analysis in the dynamic case where harmonization costs are not incurred produces an equilibrium that may be of a higher welfare level. Coordination would occur without the need for strict rules. Their conclusion is that if countries maintain sound public finance, tax competition would not lead to a race to the bottom. This work is illustrative of the fact that outcomes to games depend upon modelling choices, such as harmonization costs in this case.

2.2 Monetary Policy

The Barro Gordon model of monetary policy captures how short-sighted policy makers succumb to the temptation of inflation to temporarily boost an ailing economy. The repeated version of the Barro-Gordon model illustrates the usefulness of the repeated game paradigm in applied economics [5]. The Static Barro-Gordon Model of Inflation operates as follows:

- In the first period firms choose their expectations of inflation, π^e .

- In the second period the central bank observes π^e and chooses actual inflation π . This gives the central bank the power to boost GDP by fueling inflation by surprising the market. This devalues pensions and savings but boosts output thereby presenting a dilemma for policy makers. The timing reflects the concept that the central bank has an information advantage over the firms.

This simplified model considers all firms as a single player. Firms maximize their utility function $-(\pi - \pi^e)^2$ and hence optimally choose expected inflation to be actual inflation π . Firms lock in wages and prices given this expectation so surprises in inflation can boost or cut demand. The solution concept that describes the expected outcome of the above sequential game is a subgame perfect equilibrium (SPE). This is determined via backward induction from final stage game Nash equilibria to prior subgame equilibria in which expected actions are known given the solutions just calculated [20].

This model helped to identify an optimal government strategy of appointing an independent agent with a long term contract charged with maintaining low inflation. The lengthy employment period provides an incentive to be patient and avoid the temptation of increasing inflation for short term gain.

2.3 Auctions for Procurement or Asset Sales

An auction is a market mechanism in which an item, or items, are exchanged on the basis of bids submitted by participants. The auction mechanism consists of a set of rules that govern the sale, or purchase in a reverse auction, of an item according to the most favorable bid. Auction theory concerns the design of auctions and how their rules should be determined so that desired goals are achieved. The goals of auction design, e.g. revenue or social welfare maximization, may differ according to the requirements of the bid-taker and the bidders' valuations for the item(s) on sale.

Auction theory has emerged as an important field of research in recent years [29, 28]. The proliferation of large-volume, state-run auctions for items such as spectrum-rights has focussed attention on how auctions should be designed [37, 39]. Subtle changes to auction rules can cause marked differences in outcomes.

2.3.1 English auction.

This is the classical auction type that is typically used in the sale of property or collectibles. Participants bid openly against one another, with each successive bid being higher than the previous one. The auction continues as long as there is someone willing to outbid the current asking price. The auction then ends when no participant is willing to bid further, at which point the highest bidder pays his declared bid amount. The seller may set a reserve price so that if the auctioneer fails to raise a bid higher than this reserve the sale may not go ahead, but the seller typically pays a fee to the auctioneer in any case.

2.3.2 Dutch auction.

The Dutch auction is also known as an "open-outcry descending price" auction or a "clock auction" [39]. The auctioneer starts at a high price and subsequently lowers the price repeatedly

as time passes. This is sometimes achieved by using a clock which gradually decrements the price. The first bidder who communicates that he will accept the current price wins the item at that price for a quantity they may specify publicly to the auctioneer. The Dutch auction has been made famous by the Amsterdam flower auctions, where speed is essential because of the perishable nature of the products for sale.

2.3.3 Sealed-Bid first-price auction.

In this type of auction all bidders simultaneously submit bids in such a way that no bidder knows the bid of any other participant. The highest bidder pays the price he submitted. In this form of auction, bidders strategize about the valuations of other agents when evaluating what they should bid in order to maximize their expected surplus.

Bidders in the traditional Dutch auction and sealed first-price auction will bid below their true valuation for an item so that they can maximize their expected profit. This tactic is known as bid shading. These two auctions are also theoretically equivalent, but in practice Dutch auctions will produce less revenue than sealed first-price auctions. This is one of the most important results of experimental economics [54, 55].

2.3.4 Sealed-Bid Second-Price (Vickrey) auction.

This auction is conducted in the same manner as a sealed-bid first-price auction with the difference that the winner only pays the amount equal to the second highest bidder. The auction format has appealing strategic properties in that it is a dominant strategy for each bidder to bid his true valuation. A recent auction of note was the initial public offering (IPO) of Google shares, that used a modified multi-unit, homogenous, Vickrey auction [12, 26, 60]. In a pure Vickrey auction the optimal strategy for bidders is to bid their true maximum value. So if a bidder is willing to pay \$100 per share then that should constitute his bid. If the Google IPO auction was a pure Vickrey (assuming all potential bidders were rational and permitted entry to the auction) then the post-IPO market price would equal the IPO clearing price. This means that the post-auction share price should not fluctuate excessively in the immediate aftermath of the flotation. However, the auction was not a pure Vickrey auction and all the bidders were not completely rational. For example, some bidders may not have been able to calculate their optimal strategy correctly. Google also reserved the right to set the offering price lower than the auction-clearing price.

3 Modelling Agent Behavior

The behavior of agents often depends upon whether they can accurately predict the actions of other agents. In an independent private-values model, each bidder knows how much he values the object for sale but his value is not dependent upon the bids of others [18, 58]. Each bidder derives a value from only his own personal tastes and not external factors such as re-sale value.

The revelation of other agents' types does not influence each agent's private type.

The common-values model was subsequently introduced, where the true actual value is the same for everyone, e.g. the oil in a drilling rights auction, but bidders have different private information ("signals") about the true actual value [50, 61, 63]. In this model a bidder would change his estimate of the value if he learned another bidder's signal. Common valuations often occur in auctions for rights to natural resources [8]. If a bidder's signal was significantly more than all other bids for example, he may re-estimate the value of the item, therefore, his ex-post valuation may be decreased. If it decreases to below his bid amount, he is then a victim of the "winner's curse", a term first coined by Capen et al. [8]. The winners in common-value auctions are necessarily the most optimistic bidders when payment is conducted using a first-price scheme. This can sometimes result in winning an item at a cost of more than the ex-post valuation [7, 31, 62], which may result in serious losses for the winner [25].

Environments with asymmetric information describe situations in which some agents hold private information that is relevant to all parties [61]. This information can be directly relevant in that it directly affects the payoffs of the bidders. For example, when the bid-taker knows the quality of the items for sale but the bidders do not. Asymmetric information can also be indirectly relevant by helping each agent to gauge the expected rational behavior of others and thereby solve his strategic uncertainty.

3.1 Game-Theoretic Behavior

Second-price (Vickrey) auctions with private values. In this type of auction, the winner is the bidder with the highest bid but he only pays the second highest bid amount for the item. The optimal bidding strategy is to submit a bid equal to one's actual true value for the item. This is a weakly dominant strategy, i.e. no matter what other bidders do, no other strategy can achieve a superior outcome. Note that a dominant strategy in a game for a player gives a better payoff than any other strategy, regardless of what the other players are doing. It weakly dominates another strategy if it is always at least as good [20]. In an English auction the winner pays the amount at which the second highest bidder dropped out, suggesting that it lends itself towards analysis via the second-price auction format [57]. Vickrey's analysis can, therefore, be applied to such auctions when attempting to predict bidder behavior [58].

Sealed-bid first-price auctions with common values. Consider the sale of drilling rights in an oil-field. A strategic bidder recognizes that winning the auction implies that other bidders estimated a lower valuation for the oil in the field. These other estimates would cause the bidder to be nervous about their valuation and result in a lower estimate of the value of the oil. The equilibrium bidding strategy is, therefore, to reduce the bid amount slightly to take the "winner's curse" into consideration. This is an example of a Bayesian game in which the players have incomplete information about the game. This can be described in terms of a game in which a player may have one of many (or even infinite) "types", and that the type of any player is known to that player but unknown to others [20]. The type of a player determines the payoffs that player receives from any outcome of the game. The common equilibrium notion for such games is Bayes-Nash Equilibrium (BNE). In a BNE, each player picks a strategy function,

rather than a simple strategy [57, 20]. The strategy function then selects a particular strategy for the player's type. A BNE is a profile of strategy functions such that no single player can improve his expected utility by changing his function. BNE is a solution concept that is often applied to auctions [10, 46].

3.2 Revenue Equivalence and Policy Objectives

An important result in auction theory is the Revenue Equivalence Theorem (RET) which states that, if all bidders are risk neutral, and have independent private values for the auctioned items, then the standard single unit auctions have the same expected revenue.

Theorem 1 (Revenue Equivalence Theorem (RET) [40, 49]). Assume each of a given number of risk neutral potential buyers has a privately-known valuation independently drawn from a strictly-increasing atomless distribution, and that no buyer wants more than one of the k identical indivisible prizes. Then any mechanism in which (1) the prizes always go to the k buyers with the highest valuations and (2) any bidder with the lowest feasible valuation expects zero surplus, yields the same expected revenue (and results in each bidder making the same expected payment as a function of her valuation).

Proof. See [40] for a proof. □

The RET applies broadly to many auction types, beyond just the English, Dutch, first- and second-price auctions, to include many other auction mechanisms. It is a remarkable result, as various auctions may have completely different strategies and rules. It is important to note, however, that although the expected revenue is the same for auctions that satisfy the conditions in Theorem 1, it may vary in different specific instances. Some auctions return revenues with greater variance even though the mean revenue is the same. A risk averse bid-taker may not welcome variance and would, therefore, prefer a more predictable outcome [59]. This offers a partial explanation for the prevalence of first price auctions over second price auctions.

4 Learning in Games

As Fudenberg and Levine [19] pointed out, there are some important conceptual and empirical shortcomings of classical game theory.

- If there are multiple equilibria how do agents coordinate their beliefs about each other's play by pure introspection?
- Common knowledge of rationality and about the game itself can be difficult to establish.
- Equilibrium theory can better explain likely outcomes in later rounds of a repeated game than early rounds. This shift from non-equilibrium to equilibrium play is impossible to reconcile with a purely introspective theory.

Learning can explain the tendency for play to move towards an equilibrium in repeated games. Players imagine that the other agents are pre-programmed to play either a certain strategy (Cournot adjustment) or to randomize over a set of strategies according to some fixed probabilities (fictitious play). Additional data gathered in repeated rounds can provide a more accurate picture of the initially unknown intentions of other agents that may allow a more accurate prediction of the (mixed) strategy played by the other player. By playing a best response to

this increasingly accurate estimate of other firms actions, an equilibrium can be reached over time.

Learning models assume rationality but not common knowledge of this rationality. The basic assumption that other players are a black box ignores the strategic interaction in each game. Learning models can suffer from circular dependencies in which repeated patterns of play can be observed that never converge. However, learning models are a useful means of gauging how convergence may or may not occur.

4.1 Fictitious Play

The fictitious play process involves an assumption by players that their opponents strategies are drawn from some unknown distribution that remains stationary. For clarity of exposition we restrict our attention to a two-player setting and assume that the strategy sets of both players are finite.

Players choose the best response to their belief of their opponent's strategy. Each player has a weighting function $\theta_i^0 : S_{-i} \rightarrow \mathbb{R}^+$ that counts the number of times that an opponents strategy profile is played.

$$\theta_i^t(s_{-i}) = \begin{cases} \theta_i^{t-1}(s_{-i}) & \text{if } s_{-i} \neq s_{-i}^{t-1}, \\ \theta_i^{t-1}(s_{-i}) + 1 & \text{if } s_{-i} = s_{-i}^{t-1}. \end{cases} \quad (1)$$

Agent i estimates the probability $\gamma_i^t(s_{-i})$ that an opponent plays strategy profile s_i .

$$\gamma_i^t(s_{-i}) = \frac{\theta_i^t(s_{-i})}{\sum_{\tilde{s}_{-i} \in S_{-i}} \theta_i^t(\tilde{s}_{-i})} \quad (2)$$

Agent i selects a pure strategy that is a best response to his belief of other players' strategy profiles. There is not necessarily a unique best-response to every assessment so fictitious play is not always unique.

4.2 Convergence Properties

Strict Nash equilibria are absorbing so once a learning process arrives at such an outcome it will remain there and thus converge. However, fictitious play cannot converge when there are only mixed strategy Nash equilibria. For example, in Figure 1 Shapley described a game with payoffs akin to that of rock-paper-scissors [52]. The unique Mixed Strategy Nash Equilibrium is for both players to mix $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. He showed that it is possible that given an off-diagonal starting point, fictitious play will cycle around off-diagonal outcomes indefinitely and not converge to the true equilibrium.

Furthermore, it is possible to show that sometimes convergence to an unstable Nash Equilibrium is possible. For example, in the game described in Figure 2, where both players gain if they choose different actions, there are three equilibria. The first two are the absorbing pure strategy Nash equilibria where the both choose opposite sides. However, the third equilibrium is for the players to mix over A and B. Fictitious play can lead to a cycle in which both players continuously choose the same side of the coin and move between (A,A) and (B,B).

	L	M	R
T	0,0	1,0	0,1
M	0,1	0,0	1,0
D	1,0	0,1	0,0

Figure 1: Cycling in a game of Rock-Paper-Scissors (or top-middle-bottom, left-middle-right).

	A	B
A	0,0	1,1
B	1,1	0,0

Figure 2: Game that may produce pathological behavior with learning agents.

This example is proof that convergence to a Nash Equilibrium is not always an indication of a likely outcome of a game because pathological behavior that is unlikely in practice may be predicted.

5 Mechanism Design

Mechanism Design can be considered as “inverse game theory” whereby the rules of the game are decided by an authority so as to fulfil some objective. Two typical goals of auction design are either revenue maximization or maximization of social welfare. The goal of maximizing revenue is an obvious one and features in auctions where the identity or private valuations of the winning bidder(s) matter little when compared to the revenue received by the bid-taker. In some circumstances, however, the bid-taker may wish to achieve certain social objectives, but because these individuals’ actual preferences are not publicly observable, the analysis of such auctions can become more complicated. The mechanism design problem is to elicit these preferences so that they may be aggregated into social preferences to form a collective decision [20, 35].

5.1 Mechanism Design Goals

The traditional goal of mechanism design is to determine the rules of a game in which an overall equilibrium (or equilibria) is reached according to some desirable system-wide properties, given that all participating agents are self-interested [35]. A social choice function (SCF) describes the properties that the outcome should possess. Some typical desirable properties include the following:

- Individual Rationality: No agent attempts to take part in a trade that fails to increase, or at least leaves constant, his own utility [34]. This is an important property if agent participation is voluntary.
- Efficiency: The outcome must maximize overall agent utility, thereby maximizing social welfare.
- Revenue Maximizing: A single agent, an auctioneer for example, maximizes her revenue (utility).
- Budget Balance: The sum of all agent payments equals zero, therefore, no money is extracted or injected into the system. This is particularly important for any self-sustaining mechanism where no external benefactor exists to subsidize the system.

However, these desirable properties may directly conflict with one another. For example, budget balance and efficiency conflict in Vickrey auctions, which achieve only the latter property.

The Revelation Principle [20, 35] is a fundamental tenet of mechanism design and implies that in a wide variety of settings, only “truthful revelation mechanisms” in which agents truthfully announce their types need to be considered. This result is not immediately obvious but means that there are no manipulable mechanisms (i.e. in which untruthful behaviour is advantageous) that, when agents strategically report their types, attain superior outcomes according to the social objective than any non-manipulable mechanism. Conitzer and Sandholm [11]

have demonstrated, however, how this principle may fail in certain extreme instances when computational or communication complexity hinders strategic manipulation.

5.2 Mechanism design for economic policy-making

Mechanism design theory is ubiquitous and affects almost all aspects of policy with implications at two levels. Firstly, it tells us when markets can be expected to lead to desirable outcomes and whether other institutions should be considered instead. Secondly, mechanism design theory offers useful design guidance for alternative institutions when markets fail. The most conspicuous policy areas that have benefitted from mechanism design are listed below.

Regulatory Economics: In Baron and Myerson's seminal work on regulatory policy-making, they used mechanism design to derive optimal regulatory schemes ensuring the provision of public services at least cost [4]. Later research showed that, when cost realizations are observable, simple mechanisms can achieve this objective. Such results have improved actual regulatory schemes and the design of contracts between international institutions.

Auction: Previously, we saw how game theory is used to analyze expected outcomes of auctions. Similarly, mechanism design (also sometimes known as inverse game theory) can be used to design the rules of an auction so that a desirable equilibrium from the auctioneers perspective is reached. Auctions are one of the first and most prominent applications of mechanism design theory. They benefitted greatly from Nobel Prize winners Maskin and Myerson's contributions regarding the challenge of how to allocate some item(s) among bidders when the value of the item(s) to a bidder is private (known only to that bidder). The objective of the auction designer may be to maximize revenue, or to ensure that the items are awarded to those who value them the most thus maximizing economic efficiency. Governments use auctions to allocate a country's natural resources that include mineral deposits, exploration rights, timber, frequency spectrum or property. Governments also use reverse auctions to procure goods and services from private sector suppliers. Mechanism design theory has been instrumental when guiding the design of a set allocation and payment rules for auctions across many applications.

Environmental Policy: Global co-ordination of pollution control is essential given that we all share a single atmosphere and ocean system. Mechanisms for internalizing the externalities of pollutants such as carbon dioxide form an integral aspect to any solution that will curtail pollution. Economic theory has informed efforts such as the Kyoto Protocol but more effort is required in order to overcome political hurdles. In related applications, mechanism design theory also informs the design of sustainable management of natural resources such as fishing or tree-felling.

Development Programs: Mechanism design theory has also heavily influenced the design of aid programs in poor countries [48]. Traditional solutions to community problems such as lending, land sharing arrangements and resource management have been improved following the contributions of mechanism design theory [33]. For example, mechanism design helps evaluate the relative performance of different cross-reporting and joint liability in microfinance arrangements [1].

5.3 Mechanism Design Formalism

Mechanism design is the crafting of a set of rules to determine allocations of items and payments so that participating agents behave in an expected manner. It is a game of private information in which a central agent, “center” or “principal” chooses the payoff structure. The agents know instinctively their own information relevant to payoffs. For example, a message may contain information about their preferences or the value of an item for sale. This information is referred to as the agent’s “type” $\theta \in \Theta$. Agents report a type to the center $\hat{\theta}$ that they may choose strategically so that it is different from their true value θ . After the reporting phase, the center determines an outcome that consists of an allocation and payments.

The outcome o consists of an allocation and payoff $o(\theta) = \{x(\theta), p(\theta)\}$, $x \in X, p \in P$ where x and p are allocation and monetary transfer functions, respectively. The center typically wishes to fulfil a social choice function $f(\theta)$ to map the true type profile directly to the allocation of goods transferred,

$$f(\theta) : \Theta \rightarrow X$$

whereas a mechanism maps the reported type profile to an outcome

$$y(\hat{\theta}) : \Theta \rightarrow O.$$

There are some key design considerations for mechanisms that seek to achieve a social choice function in equilibrium. The most important one lets us confine our attention to truthful mechanisms when seeking to implement a social choice function¹. The revelation principle described above informs us that no matter the mechanism a designer can confine attention to equilibria in which agents truthfully report types. The revelation principle is as follows: “For any Bayesian Nash equilibrium there corresponds a Bayesian game with the same equilibrium outcome but in which players truthfully report type.” This allows us to solve for a Bayesian equilibrium by assuming truthful reporting of types subject to an incentive compatibility constraint.

Truthfulness and The Generalized Vickrey Auction The seminal contributions of Vickrey [58], Clarke [9] and Groves [24] (VCG) to the field of mechanism design provide an important standard. VCG is a general method of designing truthful mechanisms and a detailed proof of its truthfulness may be found in [35].

For a given mechanism, a solution concept is used to predict the strategies that agents will choose in order to maximize their utility, thus determining an equilibrium position for the game. The Generalized Vickrey Auction (GVA) uses the VCG mechanism to determine payments and has a dominant strategy equilibrium. This means that the best response strategy for each agent remains the same, irrespective of its knowledge about other agents or their actions. Agents bid their truthful valuation for an item and do not profit from under-estimating or exaggerating their valuation even if they know all other bids. This is a powerful solution concept and makes it unprofitable for agents to concern themselves with other agents’ bids.

The GVA also ensures optimal efficiency, thereby maximizing social welfare according to some objective such as fairness. However, it is not budget balanced, so a benevolent external party,

¹These games contains private information and thus entail Bayesian equilibria.

such as a government, may be required to supplement the budget. To determine payments to each agent participating in the overall optimal solution, the overall revenue is determined without each agent present in turn. This involves $m + 1$ optimization problems if m agents participate in the optimal solution to the allocation problem.

Although optimization of a social objective becomes possible using such a truthful mechanism, there are several notable disadvantages of the GVA. If non-optimal solutions are found to the optimization problems that determine the prices paid (based on the Vickrey-Clarke-Groves mechanism [9, 24, 58]) the mechanism is no longer guaranteed to be truthful. Whilst various polynomial-time approximation algorithms can provide good or near optimal solutions very quickly, Nisan and Ronen [42] showed, constructively, that a non-optimal solution can in fact result in payments arbitrarily far from optimum. If an auctioneer seeks to approximate optimal solutions in a GVA using polynomial-time algorithms the results may not be reliable and agents may have an incentive to lie.

Other limitations include reduced revenue compared with other auctions and susceptibility to a fraudulent auctioneer. It is possible for an auctioneer to introduce fake bids just below the value of the winning bids to increase revenue. For this reason a trustworthy auctioneer is imperative in a GVA.

A GVA is also sensitive to bidder collusion. Bidders may coordinate their bid prices so that the bids remain artificially low. In this manner, the bidders get the item at a lower price than they normally would. The GVA self-enforces some typical collusion agreements and makes it easier for agents to conspire by allowing the coordination of bids. Therefore, from the perspective of deterring collusion, Dutch or first-price sealed-bid auctions are preferable because colluding agents require more trust in one another for the collusion to succeed.

5.4 Combinatorial Auction as a Policy Tool

In a setting where multiple distinguishable items are bought or sold, bundling of related items that exhibit sub-additive or super-additive value may be desirable for competing agents. For example, a pair of flights that allow a connection to be made in a timely manner to a desired destination has value for an agent, whereas each flight has little value on its own if the connection cannot be made. If agents are not permitted to communicate such dependencies there is an exposure problem when items are sold separately possibly leading to a subset of items with little value being acquired [39]. This encourages cautionary bidding tactics that depress bids. Combinatorial auctions (CAs) may alleviate the exposure problem by permitting bids on bundles of items that match bidders' requirements. CAs can improve economic efficiency where items exhibit complementarities/substitutabilities. Government have used this approach to maximize revenues from the sale of telecom spectrum rights. They are also using combinatorial auctions to improve the efficiency of procurement across various spend categories such as vehicle fleet and courier services.

Such auctions have been used by public bodies to maximize economic efficiency in many diverse scenarios such as the sale of spectrum licences in America's Federal Communications Commission (FCC) auctions [37] and London Transport Authority's procurement of bus services from private operators [45, 39]. The Chilean government have also adopted CAs for the supply of

school meals to children [14]. In the latter case, the quality of suppliers was considered as well as the bid amount in deciding the winner, and the system also ensured there was no monopoly in any individual region. The reported supply costs have fallen by 22% since the adoption of the program. In auctions where complementarities or substitutabilities are exhibited between items, there is a compelling argument for the introduction of combinatorial bidding to improve overall efficiency.

The combinatorial auction mechanisms used by public bodies are usually first price schemes in which winning bidders pay what they bid. These schemes are not incentive compatible and do not guarantee efficiency. They do, however, tend to improve efficiency when there are notable super-additive or sub-additive value functions for items.

6 Example: Efficient Government Grant Auctions for Renewable Device Subsidies

Let us consider a more specific example. There are m households whose suitability for receiving a subsidy for a photovoltaic device depends upon the pitch of their roof, orientation, horizon profile and location. The crux of the problem for government bodies is that they lack sufficient information regarding the cost and benefit of device installations for all possible domestic and industrial users. A market mechanism that elicits this private information can facilitate an efficiency maximizing outcome. A summary of the assumed setting is as follows:


- there is an altruistic central government planner,
- a finite budget b to subsidize renewable energy micro generation,
- a set of self-interested agents I (seeking subsidies),
- private information held by each agent i ²
- common knowledge regarding the price to acquire and install a device r .

The private information that informs the value of a device (see Figure 3) is

- the pitch of roof p_i
- the orientation o_i within the location
- the latitude l_i
- the expected cashflow stream v_i


The agents compute their private value given this information (Figure 4) and this is used to determine the maximum amount they may be willing to spend on a device. This, in turn, informs the minimum subsidy they will require from the government. If the subsidy required was zero or less, then it should be optimal for them to invest in a device without support. For this reason, government bodies should not disclose an imminent subsidy scheme until shortly before it is acted upon. Otherwise, agents willing to invest in a device without subsidy would wait to receive a subsidy.

²Smart phone applications are available to help agents determine their independent private value.



(a) Roof pitch.

(b) Orientation.



(c) Latitude.



(d) Information capture.

Figure 3: Key inputs for determining value.




Figure 4: Maximum Cost an agent is willing to pay

6.1 Model

We first identify the set of parameters that can feed into a single value describing the private value of a solar device per unit of power, v^i for agent i . This value reflects the current time-discounted value of the expected stream of cashflows given their location, financial and building information. The social choice function is to assign the panels p_j , $j \in \{1, \dots, J\}$ to agents in a manner that minimizes the maximum cost for any agent. The imposition (or cost) for agent i if panel j is received is the price of the device minus the value per unit of power multiplied by the power output of the device,

$$c_j^i = r_j - v^i \phi_j, \forall i \in \{1, \dots, I\}, j \in \{1, \dots, J\}$$

where r_j is the purchase price of an installed panel.

6.2 Allocation Algorithms

This problem can be transposed to a makespan minimization problem denoted in the scheduling theory literature as $Q||C_{max}$, where C_{max} refers to makespan in scheduling theory. We wish to allocate device acquisition and hosting responsibility (jobs) across houses (machines) that each perceive a private cost associated with acceptance of that job. The minimization of the maximum time to wait for all jobs to complete is comparable to the minimization of the maximum cost imposed on any house-owner so that inconvenience is bounded as tightly as possible.

6.2.1 Inefficiency and a Non-monotone Algorithm

Consider the following simple algorithm: order the panels from highest to lowest power and greedily assign each device in turn to the household that has received lowest cost imposition thus far in the partial allocation. This algorithm is a 2-approximation that is non-monotone. Consider an example with 3 devices $\{d_1, d_2, d_3\}$ and 2 agents $\{h_1, h_2\}$ that illustrates non-monotonicity. Let the publicly known power ratings for the devices be $\phi_1 = 10W$ and $\phi_2 = \phi_3 = (9 + \epsilon)W$ and all devices cost $r_j = 60\text{€}$. This is common to all agents. However, let each house-owner's value per unit of power be $v_1 = 5\text{€/}W$ and $v_2 = (5 - \epsilon)\text{€/}W$. This is the private information that we wish to elicit. Our greedy algorithm first assigns: $d_1 \rightarrow h_1$, $d_2 \rightarrow h_2$ and $d_3 \rightarrow h_2$ resulting in costs $c_1 = 60 - 10 \times 5 = 10\text{€}$ and $c_2 = 2(60 - (45 - 4\epsilon - \epsilon^2)) \cong (30 + 8\epsilon)\text{€}$. But if we increase v_2 so that $v_2 = (5 + \epsilon)\text{€/}W$ then it receives only the first device. The first (highest power device) will be assigned to the second agent because she now has a higher value per unit of power. The second (lower power) device is assigned to the first agent because this agent has received a lower cost imposition thus far. The third (lower power) device is also assigned to the first agent because this agent has received a lower cost imposition thus far because the previous lower power device imposed less cost than the high power device did for the agent that values renewable power more. So this algorithm is not monotone and a direct consequence of this is that it cannot be used within any truthful mechanism for allocating devices to agents [2].

6.2.2 Efficiency and Monotonicity

There exists a randomized 3-approximation that is truthful in expectation. Kovacs et al. [30] developed an approximation scheme for scheduling n jobs to m machines of different speeds so that the makespan is minimized. This problem is sometimes referred to as $(Q||C_{max})$. A fast, deterministic monotone 3-approximation algorithm exists for this problem. The importance of monotonicity is very relevant to our setting and the context of truthful mechanisms in general. When each agent knows its own value for hosting a device, it is necessary to design an incentive for declaration of true values to enable efficient allocation. Archer et al. [2] demonstrated that such motivation is possible only if the allocation algorithm within the mechanism is monotone.

7 Dynamic games and Stochastic Dynamics

Policy decisions are often made infrequently and in the absence of information as to environmental conditions in future periods when specific protocols or mechanisms to implement policies are in effect. For this reason, it is useful to consider mechanisms that are dynamic or reactive to future circumstances. The topic of dynamic games is very relevant to policy makers that are often dealing with challenges for which long term solutions are required. Parkes et al. provided a recent survey of the dynamic mechanism design framework that embraces coordinated decision making in an external environment with uncertainty and agents with dynamics preferences [44].

7.1 Online VCG

In order to deal with a scenario in which the number of agents may be increasing or decreasing over time, it is necessary to consider how a central planner deals with such temporal issues in a mechanism that incentivises truthfulness. This is applicable to congestion games such as traffic management, vehicle recharging or access to spectrum and is thus important from a policy makers perspective [17, 21].

A generalization of the VCG mechanism to dynamic environments is described by Parkes and Singh [53]. The online VCG mechanism provides agents with an incentive to follow a decision policy that maximizes the sum of expected total discounted value to all participants. The mechanism's incentive properties extend to any problem in which each agent's value, conditioned on a sequence of actions by the mechanism, is not dependent upon the private type of others. Payments are collected so that each agent's expected payment is precisely the expected externality imposed by the agent on others following its arrival [44]. The expected externality is the decrease in the total expected discounted value to the other agents under the updated optimal policy following agent i 's arrival. It is necessary that the mechanism employs a correct probabilistic model of the environment with explicit expectations of agent arrivals and departures.

7.2 Dyanmic VCG and its Application to Energy Storage Management

A revenue maximizing payment scheme that is a dynamic analogue of the VCG mechanism was developed by Bergemann and Valimaki [6]. This dynamic version of the VCG mechanism requires that at each stage each agent pays the amount she inhibits other agents from obtaining value (now and in the future) by her current report. Agents hold a view that is probabilistic and evolves over time. This scheme is very relevant for domains in which uncertainty and volatility are significant. We outline an application for regulatory management of an energy resource.

7.2.1 Example: Commodity Storage

Commodity storage is as an important buffer between uncertain supply and demand volumes and provides a cushion against volatile commodity prices. For bulky commodities such as natural gas there is a notable scarcity of storage. The problem of regulating access to a managing third party access to commodity storage where multiple self-interested agents share space and flow capacity in a single physical store. The rationale for strategic manipulation in equilibrium arises because firms can “steal” injectability and deliverability from others by acting slightly before a joint optimal co-ordinated action. This occurs because the pressure in the store is shared and premature injection causes an externality whose cost is shared. Efficiency maximizing economic mechanisms for commodity storage can introduce penalties for socially harmful actions (i.e. causing the pressure in the store to be suboptimal).




Figure 5: Cross section of underground gas storage facility.

7.2.2 The Societal Benefits of a Smarter Regulatory Regime

Natural gas is a scarce resource and vital to both heating and electricity generation. An increase in the economic efficiency in natural gas storage facilities has several notable side-effects. It can be viewed in a positive light when one compares increased competitiveness against other fossil fuels. Natural gas is widely regarded as the most environmentally friendly fossil fuel because it releases less carbon dioxide and other pollutants than oil or coal. A more efficient storage mechanism would help gas to compete against coal in the market for domestic and commercial heating as well as for electricity power generation. Also, there is a correlation between the propensity of an electricity network to support intermittent renewable energy and

the gas-fired generation capacity. Gas-fired generators react quickly and thus offer synergies for wind generation [47]. A report from MIT provides a forward looking study of gas as a fuel for future decades and thus informs policy makers on the need for protocols and mechanisms to guide the marketplace [23].

8 Conclusion

Strategic decision making by self-interested agents is an integral aspect of how individuals and firms make decisions to improve their own utility or wealth. Policy makers in public institutions need to consider the effects of such behavior when designing legislation, protocols or mechanisms for sharing or regulating access to public goods. The advent of sophisticated modeling capabilities and equilibrium concepts from game theory can help policy makers better understand possible outcomes of policy implementations. Game theory can offer important guidance to ensure that actual outcomes match expected behavior at the time of policy design.

Furthermore, the development of scientific methodologies for constructing economic mechanisms to effectively co-ordinate actions can improve social welfare considerably. Mechanism design theory (or inverse game theory) provides a general framework to study collective decision problems. We examined the emergence of mechanism design as a rapidly evolving economic tool that aims to improve the effectiveness of economic protocols given a model of game theoretic rational decision making by agents. We illustrated sample applications that serve to highlight the relevance of these advances. We also summarized relevant advances in the field that may support policy making capabilities and we focused upon dynamic mechanisms with uncertainty. This subfield of mechanism design is particularly relevant when catering for domains in which uncertainty and market dynamics can pose significant challenges for public bodies when attempting to design robust and non-manipulable rules.

References

- [1] Zulkefly Abdul Karim. Microfinance and mechanism design: The role of joint liability and cross-reporting. MPRA Paper 23934, University Library of Munich, Germany, July 2009.
- [2] Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In Proceedings of FOCS, October 2001.
- [3] R. Avenhaus, B. von Stengel, and S. Zamir. Inspection games. Handbook of Game Theory, 3, 1995.
- [4] David P Baron and Roger B Myerson. Regulating a monopolist with unknown costs. *Econometrica*, 50(4):911–30, July 1982.
- [5] Robert J. Barro and David B. Gordon. Rules, discretion, and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12:101–121, 1983.
- [6] Dirk Bergemann and Juuso Valimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010. doi: 10.3982/ECTA7260.
- [7] Jeremy Bulow and Paul Klemperer. Prices and the winner’s curse. *Rand Journal of Economics*, 33(1):1–21, 2002.
- [8] Edward C. Capen, Robert V. Clapp, and William M. Campbell. Competitive bidding in high-risk situations. *Journal of Petroleum Technology*, 23:641–651, 1971.
- [9] E.H. Clarke. Multipart pricing of public goods. *Public Choice*, XI:17–33, 1971.

- [10] Vincent Conitzer and Tuomas Sandholm. Applications of automated mechanism design. In *The UAI Bayesian Modeling Applications Workshop*, Acapulco, Mexico, 2003.
- [11] Vincent Conitzer and Tuomas Sandholm. Computational criticisms of the revelation principle. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 262–263, New York, NY, USA, 2004. ACM Press.
- [12] Javier David and Nicole Maestri. Google IPO seen as prize catch for NYSE, Nasdaq. *Forbes Magazine*: <http://www.forbes.com/reuters/newswire/2004/04/28/rtr1351468.html>, April 2004.
- [13] Melvin Dresher. A sampling inspection problem in arms control agreements: a gametheoretic analysis. Technical Report Memorandum RM-2972-ARPA, The RAND Corporation, Santa Monica, California, 1962.
- [14] Rafael Epstein, Lysette Henríquez, Jaime Catalán, Gabriel Y. Weintraub, and Cristián Martínez. A combinatorial auction improves school meals in Chile. *Interfaces*, 32(6):1–14, 2002.
- [15] Thomas S. Ferguson and Costis Melolidakis. On the inspection game. *Naval Research Logistics*, 45, 1998.
- [16] André Fourçans and Thierry Warin. Tax harmonization versus tax competition in europe: A game theoretical approach. *Cahiers de recherche CREFE / CREFE Working Papers 132*, CREFE, Université du Québec à Montréal, April 2001.
- [17] Eric J. Friedman and David C. Parkes. Pricing wifi at starbucks: issues in online mechanism design. In *ACM Conference on Electronic Commerce*, pages 240–241. ACM, 2003.
- [18] Lawrence Friedman. A competitive bidding strategy. *Operations Research*, 4:104–112, 1956.
- [19] Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*, volume 1 of MIT Press Books. The MIT Press, 1998.
- [20] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, August 1991.
- [21] Enrico H. Gerding, Valentin Robu, Sebastian Stein, David C. Parkes, Alex Rogers, and Nicholas R. Jennings. Online mechanism design for electric vehicle charging. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2, AAMAS '11*, pages 811–818, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems.
- [22] Joseph Greenberg. Avoiding tax avoidance: A (repeated) game-theoretic approach. *Journal of Economic Theory*, 32(1):1–13, 1984.
- [23] MIT Study Group. The future of natural gas: An interdisciplinary MIT study. Technical report, MIT, <http://web.mit.edu/mitei/research/studies/natural-gas-2011.shtml>, 2011.
- [24] Theodore Groves. Incentives in teams. *Econometrica*, 61:617–631, 1973.
- [25] Ron Harstad, John Kagel, and Dan Levin. Comparative static effects of number of bidders and public information on behavior in second-price common-value auctions. *International Journal of Game Theory*, 24:297–319, 1995.
- [26] Edward Helmore. Google takes high road to market. *Guardian Weekly*, May 06 2004.
- [27] Alan Holland. *Risk Management for Combinatorial Auctions*. PhD thesis, University College Cork, July 2005.
- [28] Paul Klemperer. *The Economic Theory of Auctions*. Edward Elgar, Cheltenham, United Kingdom, 2000.

- [29] Paul Klemperer. Why every economist should learn some auction theory. In Mathias Dewatripont, Lars Peter Hansen, and Stephen Turnovsky, editors, *Advances in Economics and Econometrics: Invited Lectures to 8th World Congress of the Econometric Society*. Cambridge University Press, 2003.
- [30] Annamária Kovács. Fast monotone 3-approximation algorithm for scheduling related machines. In *Proceedings of the 13th annual European conference on Algorithms, ESA'05*, pages 616–627, Berlin, Heidelberg, 2005. Springer-Verlag.
- [31] Vijay Krishna. *Auction Theory*. Academic Press, 1 edition, March 2002.
- [32] Dimitri Landa and Adam Meirowitz. Game theory, information, and deliberative democracy. *American Journal of Political Science*, 53(2):427–444, 2009.
- [33] Patrick Legros and Estelle Cantillon. The Nobel Prize: What is mechanism design and why does it matter for policy-making? voxeu.org/article/nobel-prize-what-mechanism-design-and-why-does-it-matter, October 2007.
- [34] R.D. Luce and H. Raiffa. *Games and decisions*. John Wiley & Sons, New York, 1957.
- [35] Andreu Mas-Collel, Micheal D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [36] Michael Maschler. A price leadership method for solving the inspector’s non-constant sum game. *Naval Research Logistics Quarterly*, 13:11–33, 1966.
- [37] John McMillan. Selling spectrum rights. *Journal of Economic Perspectives*, 8(3):145–62, 1994.
- [38] Paul Milgrom. Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3:3–22, 1989.
- [39] Paul Milgrom. *Putting Auction Theory to Work*. Cambridge, March 2004.
- [40] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, pages 58–73, 1981.
- [41] John F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.
- [42] Noam Nisan and Amir Ronen. Computationally feasible VCG mechanisms. In *ACM Conference on Electronic Commerce*, pages 242–252, 2000.
- [43] Axel Ockenfels and Alvin E. Roth. The timing of bids in internet auctions: market design, bidder behavior, and artificial agents. *AI Magazine*, Fall:79–87, 2002.
- [44] David C. Parkes, Ruggiero Cavallo, Florin Constantin, and Satinder Singh. Dynamic Incentive Mechanisms. *Artificial Intelligence Magazine*, 31:79–94, 2010.
- [45] Martin Pesendorfer and Estelle Cantillon. Combination bidding in multi-unit auctions. *Harvard Business School Working Draft*, 2003.
- [46] Ryan Porter and Yoav Shoham. On cheating in sealed-bid auctions. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 76–84, New York, NY, USA, 2003. ACM Press.
- [47] Meysam Qadrdan, Modassar Chaudry, Jianzhong Wu, Nick Jenkins, and Janaka Ekanayake. Impact of a large penetration of wind generation on the GB gas network. *Energy Policy*, 38(10):5684–5695, October 2010.
- [48] S. Rashid, J. Shorish, and N. Sobh. *Dynamic Contracting for Development Aid Projects: Mechanism Design and High Performance Computation*. Reihe Ökonomie. Inst. für Höhere Studien (IHS), 2006.
- [49] John G. Riley and William F. Samuelson. Optimal auctions. *American Economic Review*, 71:381–392, 1981.

- [50] Michael H. Rothkopf. A model of rational competitive bidding. *Management Science*, 15:362–373, 1969.
- [51] Masutani Sakaguchi. A sequential allocation game for targets with varying values. *Journal of Operations Research Society Japan*, 20:182–193, 1977.
- [52] Lloyd S. Shapley. On the nonconvergence of fictitious play. *RAND Memorandum*, 1962.
- [53] David C. Parkes Satinder Singh. An MDP-based approach to online mechanism design. In *Proceedings of the 17th Annual Conf. on Neural Information Processing Systems*, 2003.
- [54] Vernon L. Smith. An experimental study of competitive market behavior. *Journal of Political Economy*, 70:111–137, 1962.
- [55] Vernon L. Smith. *Bargaining and Market Behavior: Essays in Experimental Economics*. Cambridge University Press, 2000.
- [56] Sonderforschungsbereich 504. <http://www.sfb504.uni-mannheim.de/glossary/> .
- [57] Theodore L. Turocy and Bernhard von Stengel. Game theory. CDAM Research Report: Number LSE-CDAM-2001-09, Centre for Discrete and Applicable Mathematics, London School of Economics, October 2001.
- [58] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(32):8–37, 1961.
- [59] Richard Weber. Mathematics for operational research. <http://www.statslab.cam.ac.uk/~rrw1/mor/s.pdf> , 2003.
- [60] Ben White. Aiming to auction its way to a more 'inclusive' IPO: Complex scheme could confuse small investors. *Washington Post*, April 2004.
- [61] Robert B. Wilson. Competitive bidding with asymmetric information. *Management Science*, 13:816–820, 1967.
- [62] Robert B. Wilson. Competitive bidding with disparate information. *Management Science*, 15:446–448, 1969.
- [63] Robert B. Wilson. A bidding model of perfect competition. *Review of Economic Studies*, 44(3):511–18, 1977.