

Demange, Gabrielle

Working Paper

Mechanisms in a Digitalized World

CESifo Working Paper, No. 6984

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Demange, Gabrielle (2018) : Mechanisms in a Digitalized World, CESifo Working Paper, No. 6984, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/180246>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Mechanisms in a Digitalized World

Gabrielle Demange

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Mechanisms in a Digitalized World

Abstract

Due to computing and communication facilities, formal procedures, often referred to as ‘algorithms’, are now extensively used in public, economic and social areas. These procedures, currently at the forefront of criticisms, share some features with mechanisms as defined by economists, following Hurwicz. My aim is to investigate these relationships and to discuss some of the economic risks generated by the power of algorithms.

JEL-Codes: D440, D470, D710, D820.

Keywords: mechanisms, algorithms, algorithmic pricing and trading, social choice rule, data, Admission post-bac.

Gabrielle Demange
Paris School of Economics-EHESS
48 boulevard Jourdan
France – 75014 Paris
demange@pse.ens.fr

April 3, 2018

This is written for the volume ‘Social Design: Essays in Memory of Leonid Hurwicz’ edited by Walter Trockel. I thank him for his encouragement. This work is partially based on a talk I gave in the stimulating workshop ‘Social Responsibility of Algorithms’ organized by Alexis Tsoukias at Paris-Dauphine University, December 2017.

1 Introduction

Social institutions were the primary interests of Leonid Hurwicz. He developed a theory of how to analyze institutions and economic systems in terms of their incentives and enforcement properties in (1960) and (1973). In putting the emphasis on the crucial role of information for allocating resources efficiently, he formalized ideas from Hayek and Mises on the market as an aggregator of dispersed information. In doing so, as argued by Myerson (2009), he shed light on an old debate about socialism and central planning. The tools he introduced, relying on an analytical modeling of incentives, have a fundamental influence on current economics, both theoretical and applied. They paved the way to mechanism design, which considers how a designer (planner, institution, firm) who aims at achieving certain goals should choose the rules applied to individuals who act strategically. Mechanism design plays a critical role in the development of new market allocation procedures. To name a few, the allocation of students to schools, kidney exchanges, or auctions are all determined by mechanisms. As such, Leonid Hurwicz can be considered as a precursor of the new area of market design. Furthermore, the formal procedures computed by algorithms for solving public, economic or social problems share common features with mechanisms as defined by economists.¹ As these procedures are developing fast due to computing and communication facilities, the impact of Hurwicz's work now extends to an even broader area than expected.

After pioneering the mathematic modeling of incentives and introducing mechanism design, Hurwicz became interested in the 'human side' of mechanisms. The article 'But who will guard the guardians?' (2008) refers to at least two human sides. The first one is alluded to in the title of the article, which is a question raised by Juvenal, described by Hurwicz as follows

In posing the famous question, the Roman author, Juvenal, was suggesting that wives cannot be trusted, and keeping them under guard is not a solution because the guards cannot be trusted either.

In the mechanism context, Juvenal's question could be rephrased as: Who are the guardians of the institutions? Who watches whether the announced constitution (or regulation) is correctly applied? In other words, Hurwicz raises the incentives' issue on the mechanism's designer rather than on the individuals on which the mechanism applies. Apart from the designer's incentives, one may also add: Who checks whether the mechanism is correctly computed? Who checks that errors in the mechanism do not generate large risks? Such issues are especially important for the complex mechanisms that are now computed through algorithms by computers.

The second human side raised by Hurwicz (2008) pertains to the individuals and their 'illegal'

¹Naturally, the relationships do not apply to all algorithms, in particular to those computed in artificial intelligence.

strategies. This aspect is in line with his primary objective of studying the functioning of economic systems, impossible to describe fully by an analytical approach. While the analytical apparatus of mechanism design is suitable for studying institutions that must be precisely defined, such as electoral rules, it is too constrained to describe most institutions because of the numerous individuals' possibilities of action (the illegal or secret strategies).² Research has not much considered illegal strategies, but rather developed in studying well-defined settings ranging from the implementation literature to market design, in which case the issue of illegal strategies is irrelevant.

A third human consideration differs from the ones referred to by Hurwicz. A reluctance to mechanisms has been revealed now that they are implemented in the real world at a large scale through computers and algorithms. Though, the reluctance is only partly explained by the use of these computerized tools. The very genuine feature of a mechanism is to be mechanical, and this feature *per se* might be perceived as non-human, whether the mechanism is computed by hand or a computer.

The plan is as follows. Section 2 describes rules and mechanisms, introduces some basic insights from the theoretical literature, and presents two examples of mechanisms run by a governmental agency -spectrum auctions ie US n thand assignments of students to universities in France. When the designer is a governmental agency, there is a legitimate demand for explanation. I will argue that the approach called 'axiomatization' developed by social choice theory may help policymakers in providing such an explanation. Section 3 discusses algorithms (in the broad meaning used currently) in the economic and social areas and their links with rules and mechanisms. It presents mechanisms used by private firms and ends with the economic risks generated by their computing power.

2 Rules and mechanisms

It is useful to consider rules before introducing mechanisms. A typical rule is a voting procedure for the election of a president among several candidates. The rule assigns the winner to the votes. No specific assumption is made on how voters vote. A mechanism instead -in the precise sense of the mechanism design literature following Hurwicz- assumes that the votes are cast strategically and considers the rule that results of these strategic votes. As the term mechanism is more broadly used in practice, I will refer in some places to a mechanism even when no specific strategic assumption is being made.

² Such argument has the same flavor as the one saying that contracts are typically incomplete.

2.1 Rules

The rules considered here aim at solving some social issues between a group of 'units' such as consumers, workers or citizens. Specifically, a rule solves the issue by choosing an outcome, such as an allocation of resources, an assignment of tasks or a president. Crucially, the outcome is based on a profile of data, which specifies data for each unit, representing for example the unit's preferences, skills, resources. In formal terms, a rule assigns an outcome to each set of data. The rule thus starts with units' data without making assumptions on how they have been gathered. There are many rules in a variety of contexts, aiming at answering questions such as:

1. How to select a candidate? A voting procedure is a rule that assigns the selected candidate (the outcome) to the expressed votes over the candidates (the data). For some rules, a vote is a single name, while for others it is a list of admissible candidates, or a full ranking of the candidates.
2. How to rank a set of alternatives? A rule here assigns a full ranking of the alternatives, as considered by Arrow (1950):

By a "social welfare function" will be meant a process or rule which, for each set of individual orderings R_1, \dots, R_n , for alternative social states (one ordering for each individual), states a corresponding social ordering of alternative social states, R .

A social welfare function may be used to select a single alternative, the one at the top of the social ordering. As an illustration, Arrow considers a community that has to repeatedly choose between three alternative modes of social action, e.g., disarmament, cold war, or hot war. In that case a rule assigns a ranking of the alternatives as a function of individuals' preferences at the time of the decision.

3. How to rank a set of Websites?³ Consider the search engine PageRank of Google, described to 'bring order on the Web' by Page *et al.* (1999). PageRank rates the Websites corresponding to a query: the units are the Websites and the outcome is their rating, which determines in which order Websites are displayed on the screen. As described in 1999, PageRank is a method for rating Web pages *objectively and mechanically*, mainly based on the hyperlink structure: PageRank is a rule that determines the rating of the Websites as a function of data, where the data for each Website is composed of the list of Websites that point to it.
4. How to assign students to schools, to universities? An assignment procedure is a rule that assigns the students to universities based on students' preferences and universities' priorities,

³This question is very much related to the previous one, as noted by Dwork, Kumar, Naor, and Sivakumar (2001).

as discussed in Section 2.2. Similar assignment problems arise in other contexts, such as the allocation of social housing.

5. How to assign kidneys between receivers and donors? Here a rule defines an ordering and a matching between receivers and donors based on observable individuals' characteristics (compatibility, age, health status...) in the waiting list (Roth, Sönmez, Ünver 2004).
6. How to allocate a painting? Here a rule determines who wins the painting, the price paid by the winner and the compensations to others (if any), as a function of the valuations of the potential buyers (their data). An auction procedure is such a rule. More complex auctions for multiple goods are now run, as seen in the example presented in Section 2.2.

The above questions can all be answered in a discretionary way. Instead, a rule specifies the answer for all possible data *prior* to the knowledge of the current data. A voting procedure, for example, is written in a constitution before knowing the citizens' votes.

The designer of a rule may be a state, a public agency or a private firm. In the latter case, a firm is not required to explain its rule, which might not be transparent. For example the procedure used by Google to rank the Websites (question 3 above) has evolved and made less transparent than the original PageRank. Section 3 will discuss mechanisms used by private firms. When the designer is a state or a public agency, as is the case in our two next examples, there is a legitimate demand for explanation. As said in the introduction, the axiomatization approach developed by social choice theory may help policymakers in providing such an explanation (Section 2.3).

2.2 Two examples

I describe here in more detail two examples, a successful and a failed one.

Allocating spectrum rights The sale of spectrum licenses over the US illustrates a successful procedure organized by the Federal Communications Commission (FCC). The procedure dramatically changed in 1993 (see McMillan 1994). Prior to 1993, the sale of spectrum licenses over the US was an administrative decision, based on hearings and lotteries. There were obvious inefficiencies: some licenses were left unassigned and it happened that a winner of a license re-sold it quickly at a much higher price than the acquisition one. After 1993, thousands of licenses were sold through auctions. The FCC goals were multiple: to avoid monopoly, facilitate contiguous territories, favor access to certain minorities, avoid collusion. The auctions were carefully designed with the help of game theorists to satisfy these goals.⁴ They turned out to be a big success on various grounds. They

⁴As described in McMillan (1994), options were multiple, for example simultaneous versus sequential auctions, open versus sealed bids, or royalties versus reserve prices.

raised significantly higher revenues than previously and the absence of immediate resale or bilateral exchanges witnessed the efficiency of the allocation. Such types of large scale multi-item auctions are now conducted in many areas, like the millions of Internet ad-auctions. The interaction between game theorists and policy-makers to set up new allocation procedures thus proves to be fruitful.

Admission post-bac (APB) in France The procedure called APB put in place in 2009 for assigning students at their entrance to the french universities turned out to be a failure, resulting in its replacement in 2018. The high education system in France is mostly public and any student who gets the second-degree diploma called baccalaureat is entitled to a seat. As a result, the admission system has to deal with a high number of candidates (in 2016, 335 696 entered). The APB procedure is based on the centralized deferred-acceptance algorithm introduced by Gale and Shapley (1962). Students state their preference orders over the slots (a slot specifies the education curriculum and the university), the universities theirs over the students (often called priorities) and the algorithm computes an assignment, based on a virtual process of successive applications-rejections defined by the preferences. The algorithm is the building block of many successful matching mechanisms dealing with school choice or labor market clearing (for a presentation see Roth 2004).

A main property of the deferred-acceptance algorithm is that it produces a *stable* assignment, which was the main purpose of Gale and Shapley (1962). They defined an assignment to be unstable if there are two applicants α and β who are assigned to colleges A and B , respectively, although β prefers A to B and A prefers β to α . As argued by Gale and Shapley, if this situation did occur, applicant β could indicate to college A that he would like to transfer to it, and A could respond by admitting β , letting α go to remain within its quota. The original assignment is therefore 'unstable' in the sense that it can be upset by a college and applicant acting together in a manner which benefits both.

A second property of the deferred-acceptance algorithm has been shown by Dubins and Freedman (1981):

Suppose a student, called Machiavelli, lies, that is, does not apply to the universities in the order of true preference. Can this help Machiavelli? The answer is no, not if the others continue to tell the truth. Similarly for coalitions of student liars.

Such a property, according to which students have no incentives to lie, is now called *strategy-proofness*. Accounting for incentives is one of the main concerns of Hurwicz, as discussed in Section 2.4. People may be doubtful about Dubins and Freedman's claim, especially when the method is applied to a large population, as in the APB mechanism, because there are many opportunities to

lie. The argument is indeed not trivial, although it does not rely on any knowledge in mathematics.⁵

The APB procedure nevertheless revealed to be a fiasco in the last years. But it modified the deferred-acceptance algorithm in an important way, which might explain the failure. To cope with the required ‘no selection’ principle, according to which any student with the baccalaureat is entitled to a seat in any field, no priority was set for the universities. When the number of applicants to some slots largely exceeded the number of seats, students were allocated at random to satisfy the no selection principle. The result was that some students lacking the background for succeeding in a field and almost certain to fail got a seat while some others, much better qualified, did not. The problem was exacerbated by the huge increase in the number of applicants to Universities due to the 2000 baby-boom and the policy of proposing many new variants of the baccalaureat. But at the same time, neither the number of seats nor new curriculums appropriate to the background of the new students’ population followed the trend. The absurdity of the system led to its rejection in 2017 after many debates and careful examinations, which culminated in a meeting organized by the Field medal C. Villani for French policymakers and deputies.

The blame has been put on the ‘non-human’ aspect of the procedure, in particular to the random draws, due to the fact that it was implemented by an algorithm. As said previously, this non-human aspect is not the one referred to by Hurwicz in 2008, but is related to the mechanical aspect of a mechanism/algorithm.⁶ The new system that replaces APB starting 2018 instead is quite opaque with unclear specification of the universities’ objectives, and the overall process might last a very long time.⁷

The result of APB’s failure is a clear defiance towards ‘algorithms’ from the French population. In my view, the failure is due to the absence of consistency and transparency in the policy, not in the way it is computed. Taking the viewpoint of social choice theory described in the next section would have been beneficial: explain the desirable properties the government wants to achieve and make explicit the constraints. It would have made clear that the random draw was resulting from

⁵The APB procedure was not strategy-proof because it modified the deferred-acceptance algorithm so that universities’ choices depended on the students preferences. Though, students’ misreporting does not seem to have been a big issue.

⁶Another example of reluctance to automatic systems is Centrelink put in place in Australia to recover social security overpayments. Though the system, dubbed ‘Robodebt’ by users, had some flaws initially -people being unable to complain or reach the service- the main objection relied on the automatic nature of the procedure. Information can be found in <https://www.humanservices.gov.au/organisations/about-us/publications-and-resources/government-response-community-affairs-references-committee-report>

⁷Each student will first select a list of applications without order; each application in the list is sent to the corresponding university. Then there will be a succession of rejection or acceptance by universities and students (on their list). The number of applications in a list is a priori constrained to 10, but each application may regroup up to 20 slots.

the absence of selection and the space constraint. But this was not politically admissible.

2.3 Justifying a rule: Axiomatization

The huge benefit of thinking of solving problems through rules is to specify the desirable properties -called 'axioms' following Arrow (1950)- one would like a rule to satisfy. Ideally, these properties can be stated in words. This was basically the approach for designing the spectrum rights auction (along incentives issues). The 'axiomatization' approach compares the rules on the basis of the properties they satisfy.

One may distinguish two types of properties: those bearing on the outcome specified by the rule for a given data profile and those bearing on the behavior of the rule when data varies, i.e. how the outcome varies with the data.

Here are a few representative examples of properties that bear on the outcome for a given data profile: the outcome should be efficient, envy-free (for example, an allocation of tasks is envy-free if each person prefers his/her bundle of tasks and compensation to that of anyone else), anonymous (neutrality with respect to labeling), stable in a suitable sense, as the assignments reached by the deferred-acceptance algorithm.

Here are properties on the behavior of a rule when data varies. Some reflect a monotonicity with respect to data. For example, in a representative election, a party who sees its number of votes to increase at the expense of another party should obtain at least as many seats. In the assignment problem, no student is worse off if more seats are available at universities. Other properties reflect a consistency 'principle' (also called uniformity), which underlies studies in fair division, bankruptcy problems (Young 1987), apportionment problems (Balinski and Demange 1989). In fair division for example, the principle says that if an allocation among a group of individuals is fair, then it should be perceived as fair when restricted to each subgroup of individuals (for a general presentation see Thomson 1990).

To sum up, the axiomatization approach is as follows: define desirable properties on the rule associated to the context under consideration and characterize the rules that satisfy them. Why is it a relevant and delicate question? In most settings, no rule enjoys all properties that sound desirable, as stated by the impossibility theorems, following Arrow (1950) and (1951). A rule has to make a choice between properties.

2.4 Mechanism: Introducing incentives

A rule needs units' data such as preferences, characteristics to compute its outcome. How to learn them? When people provide their preferences, are they truthful? We saw that students have no

incentives to lie when the deferred-acceptance algorithm is used: it is strategy-proof. But many rules are not. Consider for example plurality voting with more than 2 candidates: a voter might be better off by not voting for her preferred candidate.

A major vision of Hurwicz has been to account systematically for the incentives of individuals to *provide* their data, possibly *strategically*. Formally, this is studied through a game in which individuals' strategies are the (non-verifiable) announcement of their data. The outcome due to the strategic players may result in a very different outcome than the one prescribed by the rule. There are however difficulties to address this issue if the rule is not strategy-proof: How do people behave? What do they know about others' behavior? Do they need to anticipate others' behavior? I give here a simple example that will be useful to illustrate the role of information in data collection.

Simple buyers-seller games Let us consider the exchange of a painting between a seller S and two potential buyers B_1 and B_2 . Each attaches a value to the painting. Let S 's valuation be 70 (meaning that S benefits from selling at a price larger than 70), B_1 's and B_2 's be respectively 100 and 80 (meaning that they benefit from buying at a price respectively lower than 100 and 80).

Consider two rules, known as first price and second price auctions, assuming the above valuations known. In each rule, B_1 obtains the painting, but B_1 pays the highest valuation, here 100, in the first price auction and the second highest valuation, here 80, in the second price auction. These two rules prescribe an efficient outcome since the painting is acquired by B_1 , the person whose valuation is the highest, but they produce a different share of the surplus.

Let us now assume that an auctioneer asks the buyers to announce their valuations. The buyers' incentives to announce their true valuations dramatically differ in the two auction rules. In the second price auction, B_1 does not benefit from lying about her valuation because as long as she wins, she will pay 80 (if B_2 does not lie). B_2 does not benefit from lying either because he can win the object only by bidding more than 100 (if B_1 does not lie), in which case the price becomes 100, higher than 80, his valuation. This argument holds more generally whatever the valuations: the second price auction is strategy-proof for the buyers (Vickrey 1961). In the first price auction instead, B_1 benefits from lying and bidding just above B_2 's bid, which is surely less than 80. Though, this strategy assumes that B_1 knows B_2 's bid; if this is not the case, finding which amount B_1 should bid starts to be quite complex as it depends on B_1 's expectation on B_2 's bid, and vice-versa.⁸

Let us now consider the seller. S also may have incentives to lie. In the second price auction for example, S benefits in posting a reserve price larger than 80, so as to increase the price paid by B_1 , at the risk of not selling the painting and missing a benefit opportunity if S does not perfectly

⁸An equilibrium as defined by Nash solves this feedback. Wilson (1967) provides a first analysis of the bidders' strategies in an equilibrium context, further developed by Milgrom and Weber (1982).

know B_1 's valuation. Note, however, that *under perfect knowledge* of the bidders' valuations, the seller extracts all the surplus whatever the auction.

Neither auction thus elicits all three players to reveal their valuations. Furthermore, strategic behavior may result in inefficiency when valuations are unknown due to foregone opportunities to trade. Such analysis and results on auctions extend to situations with more buyers and general valuations or to multiple sellers and buyers (Demange and Gale 1985). Currently, auctions are being applied at a huge scale on the Web, say for selling the ads and their positions on a Webpage (Varian 2007). The inefficiency in auctions illustrated by the above example is a robust phenomena: Inefficiency is unavoidable when valuations are unknown, as first shown by Myerson and Satterthwaite (1983) in the case of one seller and one buyer: no exchange mechanism is efficient because strategic play induces foregone opportunities of exchange. Though no one is efficient, some mechanisms are better than others. Finding them is the main issue of mechanism design. Finally, players' information is crucial in determining the outcome and a player may benefit from his knowledge of others' data.

2.5 Developments

Due to the facilities offered by Internet and computers to run mechanisms, a tremendous amount of developments is being conducted by researchers in computer sciences. Spidit for example is a not-for-profit Website that provides algorithms to compute the fair division of goods, credits, or tasks (see the description in Goldman and Procaccia 2015). A line of research, at the boundary of game theory and computing science, is referred to as algorithmic mechanism design (Nisan and Ronen 2001). A main objective is to adapt the classic analysis by considering complexity issues (in the computational sense) or by analyzing complex settings. Strategy-proofness is weakened by accounting for the cost and complexity in lying. Dynamical aspects are introduced. Complex 'combinatorial' auctions are considered, in which buyers bid for a combination of items (a package) such as a package made of take-off and landing slots at airports (see the book edited by Cramton, Shoham and Steinberg 2006). Allowing bidders to bid for a package helps reaching efficiency when the goods in a package are 'complements', meaning that a buyer values two goods more than than the sum of each. For example, complementarities arise in the FCC auction (section 2.2) when a mobile phone operator values licenses in two adjacent cities more than the sum of the individual license values, due to roaming between the cities.

3 Algorithms

In its original scientific sense, an algorithm is a computational tool solving a well-defined class of problems, such as computing a solution to a linear optimization problem or finding a stable assignment as the deferred-acceptance algorithm does (section 2.2). It is composed of a list of instructions⁹ that can be applied to different data sets, for example different students' preferences and universities' priorities. In everyday life, the usage of the word 'algorithm' has spread and includes any computerized tool that relies on data in order to provide an outcome in social life, such as an assignment, recommendations to consumers, fiscal controls, or facial recognition. In the settings involving choices and interactions between users, an algorithm can often be viewed as a way for computing the outcome of a rule (as defined as in Section 2.1). As such, it is useful to distinguish the rule from the tool for computing it. Abstracting from the computational aspect, lessons can be drawn from the theoretical and strategic analysis of rules roughly described in the previous section.

3.1 Algorithm as computing a rule

Considering the rule computed by an algorithm, I retain three main lessons from the previous analysis.

Firstly, in most contexts, an algorithm has to operate a selection between fundamental desirable properties that cannot be satisfied simultaneously. Speaking in terms of a rule -prior to the algorithm used to compute it- helps making the selection explicit and thinking in terms of the priority goals one wants to achieve.

Secondly, exhibiting the properties satisfied (or not) by the rule underlying an algorithm helps explaining its rationale.

These two points are related to transparency. They are especially relevant when the algorithm is designed by a public agency, but less so when the designer is a private firm. To illustrate, consider PageRank, the search engine of Google (section 2.1). The engine's computation of the ranking no longer follows the principles described in Page *et al.* (1999) and is largely unknown. According to the European Commission, Google has abused its market dominance as a search engine by giving an illegal advantage to another Google product, its comparison shopping service. As a result, the European Commission has fined Google 2.42 billion of euros for breaching EU antitrust rules. The judgement is based on statistics showing the bias in favor of Google's product, not on the fact that the firm does not explain how its ranking is computed.¹⁰ In other words, a firm is not required to

⁹This is not true for new algorithmic methods such as machine learning with evolving data, neuronal networks and all the techniques referred to as 'artificial intelligence'.

¹⁰More information can be found at http://europa.eu/rapid/press-release_IP-17-1784_en.htm.

provide detailed information on its algorithms. When the environment is complex, the rule that is computed by the algorithm might be difficult to decipher.

Thirdly, information is crucial. When individuals provide their data voluntarily, the design of an algorithm should account for their incentives to distort them. When an algorithm instead extracts data for its designer, possibly without the individuals' consent, new issues arise. While data collection and processing bring social benefits in many areas, they also raise privacy concerns. Both public and private designers can use sophisticated techniques to extract data but often for different purposes. I examine in the next section the case of private firms.

3.2 The power of algorithms

This section discusses some economic aspects that pertain to the power of algorithms/mechanisms used by firms.

Collecting personal data: Extracting the new oil? As we have seen on a simple example between a seller and buyers (Section 2.4), the information the participants have on each others' valuations dramatically affect the exchanges. In particular, a seller informed on the buyers' valuations can extract the surplus of the exchanges, provided there is no competing seller. When Internet users search or visit Websites, they provide accurate information to search engines or companies on their intent to purchase, thereby allowing for discriminate pricing and surplus extraction. This is part of the extraction of the 'new oil'.¹¹

This point resonates with current 'privacy' concerns. The EU General Data Protection Regulation (GDPR), to be enforced in May 2018, regulates the use of data.¹² A primary goal is to improve the users' knowledge about how their data are processed, not about the economic implication of data collection. Article 12, which deals with 'Transparent information, communication and modalities for the exercise of the rights of the data subject' mainly specifies what the collecting entity should reveal to the subject whose data is collected. Though useful in some contexts, this information comes too late in others. Once an Internet user's intent to buy a product is known, a firm can charge him/her a higher price in a world where prices can be changed very quickly and be personalized. This is a major economic problem, complex to solve. So far, one could argue that firms and users are engaged in a kind of tacit barter agreement where firms deliver a valuable service (search engine, targeted proposals for products, targeted ads, access to friends, maps) in exchange for the uploading of users' data.¹³

¹¹Clive Humby, a data scientist, claimed in 2006 'Data is the new oil', now a popular maxim.

¹²<https://www.eugdpr.org/>

¹³Bloch and Demange (2018) study whether tax instruments are useful in reducing excessive level of data collection.




Figure 1: **An example of the Buy Box on Amazon Marketplace**
 from Chen, Mislove, Wilson (2016)

Organizing market exchanges and algorithmic pricing Internet platforms (AirBnb, Uber, Amazon ...) organize exchanges through computerized tools. To illustrate, consider the Amazon platform Marketplace, which allows sellers (including Amazon) to display and sell their products. A main tool of Marketplace is the 'Buy Box', which displays the price and the name of a particular seller corresponding to a search, as shown in Figure 1 (see Chen, Mislove, Wilson 2016 and the references therein). Being the Buy Box winner yields a significant advantage since 82 % of the sales go through the Buy Box. The Buy Box winner, chosen by a mechanism/algorithm designed by Amazon, is not necessarily the seller with the cheapest product. The mechanism has unknown features, resulting in an informational gap between Amazon and the other sellers. Though, as a private firm, Amazon cannot be enforced to reveal its mechanism. Other sellers have two options: decipher Amazon's strategy or leave Marketplace.

Another gap exists between the sellers, specifically between those who set their prices using computer algorithms, known as *algorithmic pricing*, and those who do not. Sellers can adjust their prices at any time and have access to all the prices posted on Marketplace. Algorithmic pricing thus allows sellers to react very quickly to changes in the prices set by other sellers. For example, they can 'track' others' prices, as illustrated in Figure 2 between Amazon and seller 2 (resp. the red and green lines). There are periods of sharp decrease in the posted prices -each seller trying to be the cheapest but at the minimum rebate- and, as soon as the competitor raises its price, the other follows the increase. This can explain why there is no evidence that algorithmic pricing pushes prices down, on average: As first shown by Stigler (1964), dynamic pricing together with the price observation facilitates collusion.

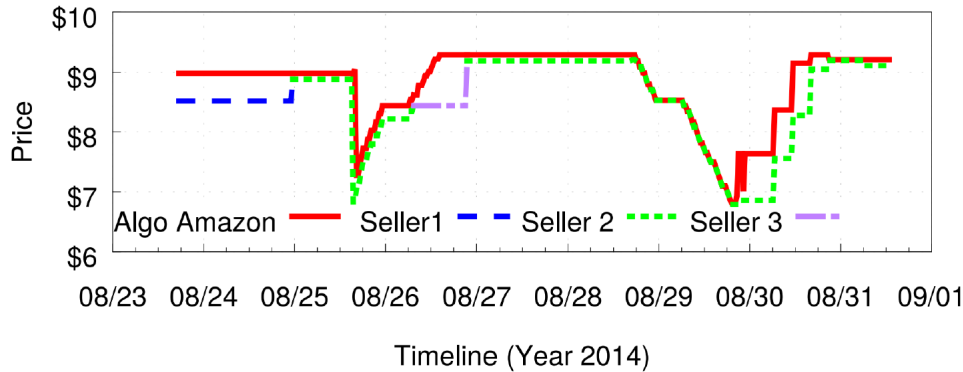


Figure 2: **Pricing strategies of Amazon and 3 sellers**
from Chen, Mislove, Wilson (2016)

Note finally that the sellers who do not rely on algorithmic pricing are unable to exploit all the available information on time. One may guess that they will tend to be eliminated.

Algorithmic trading Algorithmic trading techniques, among which is high-frequency trading, are becoming prevalent on financial markets. In 2016, high-frequency trading is estimated roughly at 55 % of the volume in the US equity markets and 80% of the volume in the foreign exchange futures (Miller and Shorter 2016).

Each algorithmic trading here is a mechanism, which takes at each instant of time the observable market data, in particular the participants' orders, to generate orders, post prices and execute trades. These automatic mechanisms interact between themselves and may result in snowball effects difficult to control, thereby calling for regulation. Pointing to the potential large and negative externalities generated by algorithmic trading, the EU Markets in Financial Instruments Directive (MiFID) prescribes:

An investment firm that engages in algorithmic trading shall have the effective systems and risk controls suitable to the business it operates to ensure that its trading systems are resilient [...] and prevent the sending of erroneous orders [...] that may contribute to a disorderly market (Article 17(1) MiFID II).

The main rationale here for the regulation is to avoid a systemic event due to algorithm trading. Such event occurred in the 'Flash Crash' of May 6 2010, which was generated by an erroneous order followed by a sequence of automatic reactions. The possible drawback of automatization first appeared in the stock market crash of October 19 1987, with the S&P 500 stock market index falling about 20 percent. At that time, a new technique, called portfolio insurance, was introduced, gener-

ating automatic sales when the market was falling, and some argued that it had a large role in the amplitude of the fall. But, according to the designers, the amplitude was due to a misunderstanding of portfolio insurance by the market's participants: Interpreting the initial automatic selling orders as bad news, the market's participants started to sell, triggering further automatic sales and creating a downward spiral (see Carlson 2007).

Currently, automatic trading has developed at a much larger scale, so algorithms and their interactions might create extreme disruptions difficult to control. There is a debate about the benefits (increased liquidity) and the costs of high frequency trading. The costs involve not only short-term price disruptions but also unfair competition because techniques such as (extremely fast) order cancellation allow high frequency traders to get advanced information on other traders' intentions. This led some to call for regulating the markets, by introducing a 'Tobin' tax or by changing the price mechanism.¹⁴

4 Conclusion

The approach pioneered by Hurwicz and few other researchers in economics and game theory has a profound and growing impact in a variety of economic and social decisions. Such impact has been multiplied by the huge development in communications and computing facilities. Though, the use of heavily technical tools and the scale at which they are applied raise new challenging issues. Keeping Hurwicz's viewpoint, a new issue is now 'But who will guard the algorithms'?

References

- Arrow, K. J. (1950). A difficulty in the concept of social welfare. *Journal of political economy*, 58(4), 328-346.
- Arrow, K. J. (1951) *Social choice and individual values* (Vol. 12). 3d edition 2012. Yale university press.
- Balinski, M.L. & Demange, G., 1989. An Axiomatic Approach to Proportionality between Matrices. *Mathematics of Operations Research*.14(4), 700-719.
- Bloch, F., & Demange, G. (2018). Taxation and privacy protection on Internet platforms. *Journal of Public Economic Theory*, 20(1), 52-66.
- Carlson, M. A. (2007). A brief history of the 1987 stock market crash with a discussion of the federal reserve response.

¹⁴I am grateful to Carmine Ventrone for calling my attention on this point. For more information, see e.g. the book edited by Easley, Prado and O'Hara (2014), in particular Chapter 4 by Golub, Dupuis and Olsen and Chapter 10 by Linton, O'Hara and Zigrand.

- Chen, L., Mislove, A., & Wilson, C. (2016). An empirical analysis of algorithmic pricing on amazon marketplace. In Proceedings of the 25th International Conference on World Wide Web (pp. 1339-1349).
- Cramton P., Shoham Y., & Steinberg R. (2006) eds. Combinatorial Auctions. Cambridge and London: MIT Press, 2006.
- Demange, G., & Gale, D. (1985). The strategy structure of two-sided matching markets. *Econometrica*, 873-888.
- Dubins, L. E., & Freedman, D. A. (1981). Machiavelli and the Gale-Shapley algorithm. *The American Mathematical Monthly*, 88(7), 485-494.
- Dwork, C., Kumar, R., Naor, M., & Sivakumar, D. (2001). Rank aggregation methods for the web. In Proceedings of the 10th international conference on World Wide Web (pp. 613-622). ACM.
- Easley, D., Prado, M. L. D., & O'Hara, M. (2014). High-frequency Trading: New Realities for Traders, Markets and Regulators. Risk Books.
- Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1), 9-15.
- Goldman, J., & Procaccia, A. D. (2015). Spliddit: Unleashing fair division algorithms. *ACM SIGecom Exchanges*, 13(2), 41-46.
- Hurwicz, L. (1960). Optimality and informational efficiency in resource allocation processes, 27-46. Stanford, CA: Stanford University Press.
- Hurwicz, L. (1973). The design of mechanisms for resource allocation. *The American Economic Review*, 63(2), 1-30.
- Hurwicz, L. (2008). But who will guard the guardians? *The American Economic Review*, 98(3), 577-585.
- McMillan, J. (1994). Selling spectrum rights. *Journal of Economic Perspectives*, 8(3), 145-162.
- Milgrom, P. R., & Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*. 1089-1122.
- Miller, R. S., & Shorter, G. (2016). High-frequency trading: Overview of Recent Developments. CRS Report, 44443.
- Myerson, R. B., & Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2), 265-281.
- Myerson, R. B. (2009). Fundamental theory of institutions: a lecture in honor of Leo Hurwicz. *Review of Economic Design*, 13(1-2), 59.
- Nisan, N., & Ronen, A. (2001). Algorithmic mechanism design. *Games and Economic Behavior*, 35(1-2), 166-196.
- Page, L., Brin, S., Motwani, R., & Winograd, T. (1999). The PageRank citation ranking: Bringing order to the web. Stanford InfoLab.

- Roth, A. E. (2008). Deferred acceptance algorithms: History, theory, practice, and open questions. *International Journal of Game Theory*, 36(3-4), 537-569.
- Roth, A. E., Sönmez, T., & Ünver, M. U. (2004). Kidney exchange. *The Quarterly Journal of Economics*, 119(2), 457-488.
- Stigler, G. J. (1964). A theory of oligopoly. *Journal of Political Economy*, 72(1), 44-61.
- Thomson, W., (1990). The consistency principle. *Game Theory and Applications*, 187, 215.
- Varian, H. R. (2007). Position auctions. *International Journal of Industrial Organization*, 25(6), 1163-1178.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1), 8-37.
- Wilson, R. B. (1967). Competitive bidding with asymmetric information. *Management Science*, 13(11), 816-820.
- Young, H. P., (1987). On dividing an amount according to individual claims or liabilities. *Mathematics of Operations Research*, 12(3), 398-414.