

Building Economic Machines: The FCC Auctions.

Francesco Guala

Studies in History and Philosophy of Science 32 (2001), pp. 453-77.

Keywords: rational choice theory; experiments; social mechanisms; methodology of economics.

Abstract: The auctions of the Federal Communication Commission, designed in 1994 to sell spectrum licences, are one of the few widely acclaimed and copied cases of economic engineering to date. This paper includes a detailed narrative of the process of designing, testing and implementing the FCC auctions, focusing in particular on the role played by game theoretical modelling and laboratory experimentation. Some general remarks about the scope, interpretation and use of rational choice models open and conclude the paper.

Building Economic Machines: The FCC Auctions.

Francesco Guala

Centre for Philosophy of the Social Sciences, University of Exeter, Amory Building, Exeter EX4 4RJ, UK. Email: f.guala@ex.ac.uk

Just as chemical engineers are called upon not merely to understand the principles which govern chemical plants, but to design them, and just as physicians aim not merely to understand the biological causes of a disease, but their treatment and prevention, a measure of the success of microeconomics will be the extent to which it becomes the source of practical advice, solidly grounded in well tested theory, on designing the institutions through which we interact with one another (Roth, 1991, p. 113).

Experimentation, and especially laboratory experimentation, can be used to provide experience with completely new modes of organisation that have no ‘natural’ origin. Such ‘synthetic’ processes cannot be studied historically since none have ever existed (Plott, 1981, p. 133).

Debates in the philosophy of the social sciences rely too often on a priori argumentation, or take off from considerations of purely theoretical scientific practice. In contrast, the philosophy of the natural sciences has been massively influenced by the detailed studies of concrete scientific practice done by historians and sociologists during the past three decades.¹ Descriptive work of the latter kind can affect philosophical views of science (for example a realist or an instrumentalist viewpoint) without necessarily engaging directly in abstract philosophical arguments. This paper is intended as a contribution to this sort of literature, setting the empirical ground for philosophical theorising about the social sciences.

The case study comes from a sub-field of rational choice theory known as the theory of ‘mechanism design’. The second ambition of this paper, then, is to try to intervene in the debate on rational choice modelling and the metaphysics of social mechanisms.² The argument stems from the assumption that we know best what we have built ourselves, and therefore whenever possible interpretations of a scientific theory (in the natural *and* in the social sciences) should take applied science as their point of departure. Given that successful technological applications have the power to drive entire research programmes at once, it is particularly important to ask what kind of interpretation of rational choice theory is warranted by its applied branches.

I shall examine one of the most publicised applications of rational choice theory to date: the construction in 1993-94 of a new market institution, by means of which the Federal Communication Commission (an agency of the US government, FCC from now on) would allocate a peculiar kind of goods. I shall therefore focus on auction theory, a branch of game theory that “is closer to applications than is most frontier mathematical economics” (McAfee and McMillan, 1987, p. 700). As it will turn out, ‘being close to applications’ is not quite the same as ‘being straightforwardly applicable’.

1. Background of the FCC auction

A ‘decentralising wave’ hit the American economy in the late eighties. A new political imperative prescribed the replacement, whenever feasible, of centralised, bureaucratic systems of allocation by market processes. Before this wave reached the telecommunications industry, licences for wireless Personal Communication Systems (PCS) – providing the right to use a portion of the spectrum for radio communication, or telephones, portable faxing machines, etc. – were assigned *via* an administrative hearing process. Each potential user had to apply to the FCC and convince them of their case; a commission would then decide whom the licence should have gone to. Such a method had a number of drawbacks: above all, it was slow, cumbersome, non-transparent, and gave licences away for free instead of selling them for their market value. In 1982 the Congress decided to reform the system and make it faster by using lotteries: each licence was randomly assigned to one of the companies which had applied for it.

The lottery system was quicker, but had other defects: in particular, some companies would participate even though they were not interested in a licence, just to resell it and make large profits out of it. A secondary market was thus created in which licences were resold by lucky winners to those who would really use them. The lottery system generated an unjust and unjustified distribution of income from the controller of the airwaves (the USA) to individuals who had done nothing to deserve it.³ In July 1993 the Congress decided that the lottery system had to be replaced by a market institution, and the FCC was faced with the problem of identifying and implementing within ten months the best auction system for selling their licences.

The problem of devising a suitable kind of auction was far from trivial, and previous experiences in New Zealand and Australia had shown that it is also a delicate one: a badly conceived reform could lead to disastrous results.⁴ As a journalist put it, “when government auctioneers need worldly advice, where can they turn? To mathematical economists of course”, and “as for the firms that want to get their hands on a sliver of the airwaves, their best bet is to go out first and hire themselves a good game theorist”.⁵ In September 1993 the FCC issued a ‘Notice of Proposed Rule Making’ setting the goals to be achieved by the auctions, tentatively proposing a design, and asking for comments and suggestions from potential bidders.

Soon a number of economists (Robert Wilson, Paul Milgrom, Charles Plott, Preston McAfee, Jeremy Bulow, Mark Isaac, Robert Weber, John Ledyard and many others) got involved in the problem as companies’ advisors. The aims to be pursued in auctioning PCS acted as constraints on the work of the consultants. The auction was intended to achieve an efficient allocation (making sure that the spectrum rights went to those companies that most valued them and could make a better use of them); to prevent monopolies; and to promote small businesses, rural telephone companies, minority-owned and women-owned firms (as prescribed by the Government and the FCC policy). Moreover, it was understood that the volume of revenue raised by the auctioneer (the FCC) was an important factor to be taken into account. The target was thus set right at the beginning of the enterprise. It took the form of an economic ‘phenomenon’ to be created from scratch, with certain specific characteristics that made it valuable in the FCC’s eyes. The following story is a *tour de force* from this preliminary identification of the target to the final product (the FCC auction as it was eventually implemented in July 1994), through a series of theoretical models, experimental systems, simulations, public and private demonstrations of the properties of different economic machines.

2. Mechanism design

The FCC enterprise is a typical case of ‘*mechanism design*’ problem. The term ‘mechanism’ is widely used in the social sciences, in various and often inconsistent ways.⁶ In this paper I shall stick to the conventional usage in philosophy of science:

mechanisms are abstract or concrete structures endowed with their own specific causal properties. Putting mechanisms at work generates ‘*processes*’.⁷ According to a widely shared view of scientific knowledge, the main task of the theorist is to explain spontaneously occurring and experimental processes, by designing an appropriate model for each kind of mechanism and the phenomena it generates. The FCC case belongs to an altogether different kind of scientific activity, proceeding in the opposite direction *from models to mechanisms and processes*.

Designs are motivated by a [...] a mathematical model, a body of theory [...] that is perhaps completely devoid of operational detail. The task is to find a system of institutions – the rules for individual expression, information transmittal, and social choice – [...] that mirrors the behavioural features of the mechanism. The theory suggests the existence of [mechanisms and] processes that perform in certain (desirable) ways, and the task is to find them. This is a pure form of institutional engineering (Plott, 1981, p. 134).⁸

Theory can be used to produce new technology, by shaping the social world so as to mirror a model in all its essential aspects. The ‘idealised’ character of a theory may thus a virtue rather than a defect, as the explicit role of theory is to point to a possibility. Theory *projects*, rather than describing what is already there.

The mechanisms we shall be concerned with are market institutions. Mechanism design is often motivated by the will to replace centralised, expensive, or inefficient systems of regulation with ‘better’ (i.e. decentralised, cheaper, or more efficient) ones. The problem is analogous to deciding whether to direct traffic using policemen rather than traffic lights or roundabouts. Each system has advantages and drawbacks in terms of cost, ambiguity of the rules, propensity to generate confusion and mistakes, costs of enforcement, and so on.

The theory of mechanism design involves both the study of the functioning of different institutions, and their evaluation. It is an enterprise between theoretical and applied economics, which requires first stating clearly the goals to be achieved by the mechanism, and then finding the best means to achieve them in the case at hand.⁹ Rational choice theory is an extremely valuable analytical tool in this enterprise. Once

the environment (agents' preferences) is defined, it is possible to think of institutional rules as defining a game, which the agents are facing and trying to solve rationally. Ideally, it should be possible to predict exactly what outcome will be achieved by a given mechanism in a given environment by means of equilibrium analysis.

3. The role of theory

When the FCC started auctioning in 1994, the results were immediately hailed as a major success, if only for the huge sums of money gained by the Federal Government (\$23 billions between 1994 and 1997). Most of the glory went to the theorists who helped design the auction mechanism. The FCC auction was claimed “the most dramatic example of game theory’s new power”, “a triumph not only for the FCC and the taxpayers, but also for game theory (and game theorists)”, and so on.¹⁰ It would be mistaken, however, to think of the FCC auction’s design as entirely theory-driven, dictated by auction theory. Auctions like those for PCS are in fact a typical example of what game theory is *not* very good at modelling.

Game-theoretic accounts of auction mechanisms date back to the sixties, thanks mainly to the pioneering work of William Vickrey (1961). Vickrey solved an auction game known as the ‘independent private values model’, where each bidder is supposed to be aware exactly of the value of the auctioned item, but does not know its value to other bidders. Such an assumption seems to be satisfied in auctions of, e.g., antiques, that will be privately enjoyed by buyers who do not intend to resell them. Wilson (1977), and then Milgrom and Weber (1982) extended Vickrey’s private value model to other important cases.¹¹ Auctions are modelled as non-cooperative games played by expected-utility maximising bidders. The players are assumed to adopt equilibrium strategies - in the standard sense of a Bayes-Nash equilibrium in which, given everyone else’s moves, no player can do better than she is presently doing by changing her strategy.¹²

After more than two decades of theoretical development, the theory of auctions still relies on a number of restrictive assumptions, and can by no means be applied to all cases. The most important and disturbing feature of the commodities (the PCS licences) to be auctioned by the FCC is their being, to say it in economists’ jargon,

sometimes ‘complementary’ (e.g. licences to provide the same service to different contiguous regions) and sometimes ‘perfect substitutes’ (licences for different spectrum bands which can provide the same service) for one another. The value of an individual licence, thus, may be strictly dependent on the buyer’s owning one or more of the other items: the value of a ‘package’ could differ from the sum of the values of the individual items that are in it. This is due to a number of characteristics of the airwaves industry, from fixed-cost technology to customer-base development, from problems of interferences to the use of different standards by different companies (McMillan, 1994, p. 150). For all these reasons a licence for transmitting in a certain region is in general more or less valuable depending on whether one owns the licence for a neighbouring area or not.

Complementarities are one of economists’ nightmares, because models of competitive markets with goods of this kind in general do not have a unique equilibrium and are unstable. No theorem in auction theory tells you what kind of institution will achieve an efficient outcome. The theory is from this point of view incomplete: there does not exist a prediction for auctions with interdependent goods. The first issue to be tackled by the consultants when they started to work in 1993 was whether to use a traditional bidding mechanism or to create something new, an institution designed ‘ad hoc’ for the specific problem at hand. Although conservative considerations of reliability pulled in the first direction, the peculiar features of the airwaves industry seemed to require the second approach. In the Notice of Proposed Rule Making of September 1993, the FCC suggested the implementation of an auction system in two stages (a ‘*combinatorial auction*’), in which goods are initially auctioned in packages (with a sealed-bid mechanism) and, later on, on an individual basis. This procedure seemed to solve the problem of aggregation in a straightforward way: the items are assigned to the winners of either the first or the second auction depending on which one guarantees more revenue to the seller. If bidders really value a package more than the individual items in isolation (as the complementarity hypothesis suggests), then the procedure should provide an incentive for aggregation.

Paul Milgrom and Robert Wilson on the one hand, and Preston McAfee on the other (consulting for Pacific Bell and AirTouch Communications respectively) objected to such a solution, and proposed an alternative mechanism called

‘simultaneous ascending-bid auction’. In a *simultaneous* auction several markets are open at the same time, and bidders can operate in all of them at once. In an *ascending* auction bidders continue to make their offers until the market is closed - which usually means until no new offers are put forward. Simultaneity and the ascending form allow each bidder to collect valuable information about the behaviour of other firms as well as about her own chance to construct the aggregation of items she most prefers. Bidders can thus switch during the auction to less-preferred combinations as soon as they realise they will not be able to achieve their primary objective. Moreover an ascending bid (as opposed to a sealed-bid) system is supposed to reduce the ‘winner’s curse’ phenomenon,¹³ because by keeping an eye on each other’s bids buyers can make a better conjecture about what the evaluations of the other bidders (and thus the real value of the licences) are. The proposed mechanism however had some worrying features: notably, it looked at first sight rather complicated and had never been put at use before.

On behalf of Pacific Bell, Milgrom and Wilson sent a document to the FCC arguing that a combinatorial process in two stages like the one proposed in the September 1993 Notice was to be discarded. Their arguments are representative of the kind of reasoning used by theorists in order to select the ‘right’ design for the FCC case. It is worth spending some time reviewing at least one of them, because it sheds interesting light on the role played by game theory in the design process.

Milgrom and Wilson argued, among other things, that a combinatorial institution may give rise to free-riding situations.¹⁴ Suppose there are three bidders (One, Two and Three); One’s willingness to pay for item *A* is 4, whereas he is not eligible to buy item *B* nor the package consisting of *A* and *B* together. Bidder Two’s willingness to pay is symmetrical: 4 for *B*, and not eligible for *A* and *AB*. Bidder Three, in contrast, is willing to pay $1+\epsilon$ for *A*, $1+\epsilon$ for *B*, and $2+\epsilon$ for *AB* (with ϵ small and positive). The payoffs are represented in the following matrix:

	<i>A</i>	<i>B</i>	<i>AB</i>
One	4	----	----
Two	----	4	----
Three	$1+\epsilon$	$1+\epsilon$	$2+\epsilon$

The only efficient allocation is in this case the one assigning *A* to One, and *B* to Two. In an ascending simultaneous auction, bidder Three bids 1 for each item, bidders One and Two bid just enough for Three to give up, and then acquire *A* and *B* respectively: the efficient outcome is a sub-game perfect equilibrium of the simultaneous ascending mechanism.

Milgrom and Wilson turned then to the two-stages combinatorial design. Under this institutional arrangement, bidder Three does not have an incentive to bid on *A* and *B* individually; he just participates in the auction for package *AB*, where he bids $2+\epsilon$. Player One can win the individual auction for *A* by bidding 1, and so can Two win the auction for *B*. But then the package would go to player Three. One and Two have therefore an interest to raise the total value of *A* and *B* by bidding more on at least one of them, but would like the other to do so in order to minimise their own costs. Milgrom and Wilson showed that such a free-rider situation has a mixed-strategy equilibrium that is inefficient. Bidders One and Two each face a sub-game which can be represented by the following payoff matrix:

	Raise bid	Don't raise
Raise bid	2,2	2,3
Don't raise	3,2	0,0

By backward induction, it can be proven that this sub-game has an equilibrium in which each bidder plays 'Raise the bid' with probability $2/3$ and 'Don't raise' with probability $1/3$. But then there is a $1/9$ probability of Three getting both *A* and *B* by paying just $1/4$ of what bidders One and Two would jointly pay for them (Milgrom, 1998, p. 15).

Such an argument clearly makes use of game theory, but does not follow from a *general* game-theoretic model of combinatorial auctions. Lacking a comprehensive theory of these processes, the economists engaged in the FCC enterprise relied on a number of independent theoretical insights and local analyses of how players are supposed to behave when solving certain tasks in isolation. “The spectrum sale is more complicated than anything in auction theory”, as two of the protagonists admit (McAfee and McMillan, 1996, p. 171). The relation between theoretical reasoning and the final implementation of the auction is well summarised by the following remarks:

The FCC auctions provide a case study in the use of economic theory in public policy. They have been billed as the biggest-ever practical application of game theory. Is this valid? A purist view says it is not. There is no theorem that proves the simultaneous ascending auction to be optimal. The setting for the FCC auctions is far more complicated than any model yet, or ever likely to be, written down. Theory does not validate the auction form the FCC chose to implement. The purist view, however, imposes too high a standard. The auction form was designed by theorists. The distinction between common-value and independent-value auction settings helped clarify thinking. The intuition developed by modelling best responses in innumerable simple games was crucial in helping the auction designers anticipate how bidders might try to outfox the mechanism (McMillan *et al.*, 1997, p. 429).

Theory played a crucial role in the design, but only if we interpret the term ‘theory’ in a loose sense.¹⁵ “The auction designers based their thinking on a range of models, each of which captures a part of the issue” (McAfee and McMillan, 1996, p. 171). It is true that “Game theory played a central role in the analysis of rules” and “ideas of Nash equilibrium, rationalisability, backward induction, and incomplete information [...] were the real basis of daily decisions about the details of the auction design” (Milgrom, 1995, pp. 19-20), but

[...] the real value of the theory is in developing intuition. The role of theory, in any policy application, is to show how people behave in various circumstances, and to identify the tradeoffs involved in altering those circumstances. What the

theorists found to be the most useful in designing the auction and advising the bidders was not complicated models that try to capture a lot of reality [...]. Instead, a focused model that isolates a particular effect and assumes few or no special functional forms is more helpful in building understanding (McAfee and McMillan, 1996, p. 172).

The other side of the coin is that it is impossible to define the exact form of the rules to be used, and theory never gives you the whole picture of the complicated process at any time. For this reason, it is true that “the auctions would not have taken the shape they did were it not for the economic knowledge brought to the design process” (McMillan *et al.*, 1997, p.429) – but only if we extend the meaning of ‘economic knowledge’ well beyond the theoretical realm. Indeed, today “much of what we know about multi-unit auctions with interdependencies [i.e. complementarities] comes from experiments” (McAfee, 1994, p. 151, n. 6).

4. Constructing and controlling ‘microeconomic systems’

Although its roots go back to the eighteenth century, experimental economics as practised nowadays is entirely a post-Second World War phenomenon. It is worth introducing some of its basic methodological tenets, and not only because they have an independent interest and will provide some taste of what an economic laboratory experiment is; the methodology of experimental economics is illuminating also in order to understand when, where, and how game theoretic models can be made to work in the lab and – by extension – in the ‘real world’ once the latter has been properly engineered. The key concept here is that of a ‘microeconomic system’.

The idea that heterogeneous knowledge is required in order to run and interpret correctly an experiment is hardly new. Like many other philosophical insights, it goes back at least to Pierre Duhem, who first noticed that a ‘cluster’ of independent hypotheses about the functioning of instruments, measurement procedures, etc. is involved in experimental science. More recently, it has become common to stress the non-theoretical, practical, tacit character of much experimental knowledge.¹⁶ In economics, a ‘correctly performed’ experiment is one that the experimenter is able to control and interpret properly. A set of precepts and rules have been devised, which

define the concept of a ‘good’ experimental system. These rules are partly independent from the theories under test, and rely heavily on the experimenter’s judgement for their application.

Pacific Bell hired in 1993 a group of economists from Caltech led by Charles Plott, to run a series of experiments that would test some theoretical insights, thus helping to choose the best design and to implement it operationally. Initially, the role of experiments was to help choose among the two institutions (the ‘combinatorial’ and the ‘continuous ascending’ auctions), for which neither a comprehensive theoretical analysis nor empirical data existed. The main concern, from the FCC and the advisors’ point of view, was with efficiency. By theoretical means, it is impossible to prove that the allocation generated by the continuous ascending auction is efficient. By observing a ‘real world’ process, very few hints about whether the outcome is efficient or not can be gained, because of the unobservable nature of bidders’ valuations. Experiments, in contrast, allow the ‘induction’ of known values on subjects, and thus enable to check whether an institution really allocates the goods to those who value them most. The Caltech team made use of what they call ‘testbed’ experiments:

An experimental ‘testbed’ is a simple working prototype of a process that is going to be employed in a complex environment. The creation of the prototype and the study of its operation provides a joining of theory, observation, and the practical aspects of implementation, in order to create something that works (Plott, 1996, p. 1).

Experimental testbeds have of course to be designed in an appropriate way. They must, first of all, fulfil the requirements of ‘Induced Value Theory’, an approach to experimentation originally due to the economist Vernon Smith (1976; 1982; 1987). Induced Value Theory consists of five ‘*precepts*’, “a proposed set of sufficient conditions for a valid controlled microeconomic experiment” (1982, p. 261). Two of them, *saliency* and *non-satiation*, are supposed to “allow us to assert that we have created a microeconomic system [...] in the laboratory” (1987, p. 108). A ‘microeconomic system’ is defined by the characteristics of the agents (their tastes, knowledge, and commodity endowments), and by some institution regulating their

decisions. The saliency requirement prescribes that the outcomes of an experiment be a function of agents' decisions only, and that the rewards be a (monotonically increasing) function of the outcomes. Non-satiation prescribes to set the rewards so that the agents strictly prefer more rewards to less in all cases.

The concept of microeconomic system is quite interesting from a philosophical point of view. The construction of a laboratory microeconomy enables one to test a theory in an 'ideal' situation, in which the agents' utility functions have only experimental rewards as arguments and are monotonically increasing with them. A microeconomic system is supposed to be a system that economists know how to interpret, which is 'transparent' to economic analysis. Furthermore, this notion plays a rhetorical role in suggesting that the laboratory systems are 'real' economic systems, in the sense that they do not differ from 'real-world' economies in any respect that is deemed relevant by standard economic theory.¹⁷

Two other precepts of 'Induced Value Theory' aim at making sure that subjects' preferences for the reward medium are not 'disturbed' by any other interfering motive. As Smith puts it, they "assure that we have created a controlled microeconomy" (1987, p. 108). The simplest incentive method, one used for instance in auction experiments, consists in telling each subject her own reservation price, and then pay (or get from) her the amount she has been able to gain (or lose) during the market experiment. Control on preferences is achieved by carefully designing the experiment so as to shield it from possible disturbances. Experimenters for example make sure that the subjects do not care about what others gain or lose (precept of *privacy*) by isolating one subject from another; they try to compensate for the cost of decision-making by paying them extra-money on top of what subjects gain or lose; they prevent subjects from attaching game-value to transactions (*dominance*) by making possible losses relevant, and so on. Notice that whereas the goals to be achieved (as specified by the precepts) are rather precise, *how* to do it is not. A lot is left to intuition, personal judgement, experience, and practical knowledge. Not surprisingly, then, whether the precepts have been fulfilled in particular cases is often a matter of controversy.¹⁸

Privacy and dominance impose constraints on agents, which are rather severely disciplined. The precepts of induced value theory define an ‘artificial’ situation: it is simply not true that saliency, non-satiation, privacy and dominance are in general instantiated in non-laboratory economic situations. These requirements, however, are crucial in the process of applying highly abstract models to concrete cases: they help to build on the one hand the experimental counterpart of the theoretical restrictions that make demonstrations from economic models possible; and, on the other, the (extra-theoretical) ‘background conditions’ which enable economic relationships to hold (when they hold at all, of course) in the real world. Let me try to be more precise.

The rigour of economics’ *theoretical* demonstrations is usually achieved by imposing a high number of assumptions, postulates, premises, etc., which collectively contribute to the derivation of some precise economic proposition. Thus, for example, the existence of efficient equilibria may be proved by assuming that the market is perfectly competitive, populated by rational and self-interested economic agents, that such agents have consistent preferences, that a Walrasian auctioneer is there to match supply and demand, that the costs of transacting are negligible, and so on. Even a simple model of supply and demand requires a rather fine-grained description of the circumstances, in order to come out with the claim that an efficient equilibrium price will clear the (fictional) market.

But theoretical models do not provide the rules *materially* to construct the markets they (partially) describe. First of all, the basic components of partial equilibrium explanations, like for instance the law of demand (agents demand more quantity as the price diminishes) are *ceteris paribus* in character.¹⁹ The quantity demanded is for example typically assumed to depend only on the price of the goods, and an accurate design of the experiment (namely one that follows the precepts of induced value theory) is supposed to fulfil such conditions. Secondly, some concepts that are used as restrictions in the model need *other* background conditions for their material instantiation. The *homo oeconomicus* of most Neoclassical models, for example, can operate only in an appropriate environment. Real human beings are more likely to behave in a self-interested manner if they do not know what the consequences of their actions will be for the other players (see the requirement of privacy above). They tend

to act more rationally if the problem-situation is well specified and they have had enough time to practice with it, etc. The important point, to repeat, is that not *all* such conditions are listed by the generic, basic theory. In some cases, economic wisdom guides experimental practice in identifying the features which allow abstract economic concepts to become instantiated in real, flesh-and-bones economic entities. Saliency, non-satiation, privacy and dominance, as well as the prescription to use experienced subjects belong to this body of knowledge. These guidelines are not *part of* economic theory, but rather tell us (roughly) how to construct real economic systems faithful to our abstract models.

To sum up, at the most abstract level, we have highly theoretical concepts such as competitive equilibrium. These are embedded in a structure of deductive reasoning from a set of strict assumptions identifying the conditions under which those concepts may be deductively demonstrated. It must be stressed that the description of the sufficient structure for effects like efficient equilibria to obtain is still abstract. The ‘real world’ counterparts of theoretical entities like the rational economic men of our models are instantiated only under further restrictive arrangements. These arrangements define the bottom level of concreteness for the applicability of economic models. Economists are guided by experimental, practical, as well as theoretical knowledge in designing their experiments so that these background conditions are satisfied. The circumstances in which an economic system maintains its own structural properties may however be quite narrow and fragile. This is of course of great concern to the economic engineer, whose machines are supposed to work for several years, in different contexts and without the constant supervision of their manufacturer.

5. Testing the robustness of alternative designs

In early 1994 the Caltech group ran a series of comparative efficiency tests of the simultaneous ascending auction versus a combinatorial sealed-bid plus sequential continuous design. The results were interpreted as favourable to the former (Plott, 1997).²⁰ In this first stage, experimenters were particularly interested in observing the functioning of the two bidding institutions in their ‘bare-bone’ versions, in order to become familiar with their fundamental properties and problems. Comparative testing

of institutions in their simplest versions, however, was also instrumental to let the first operational details and problems emerge. The mechanisms were for the first time transported from the abstract world of ideas to the laboratory. Some ‘flesh’ was added to transform the fundamental insights into working processes, flesh which might have modified, impeded, or hidden some of the systems’ structural dispositions. These experiments could therefore be seen as tests of the combinatorial and the ascending continuous auctions - but ‘tests’ of a peculiar kind, because the mechanisms themselves, rather than any theoretical model, were subject to examination. The issue was not whether some theory about an institution was to be discarded, but whether one rather than another institution should have been chosen for its properties, despite the fact that no thorough theoretical understanding of its functioning was available.

At the first step of ‘concretisation’, not surprisingly, experimentalists spotted problems that the theorists had not (and probably could not have) anticipated. Plott and his collaborators implemented the combinatorial design by running so-called ‘Japanese auctions’ combined with a sealed bid pre-auction. In the first round a sealed bid offer is made for a package of items, which are then auctioned individually in the second round: in a Japanese auction the auctioneer announces higher and higher prices and the bidders ‘drop out’ one after another; the last one to remain in the auction wins the item and pays the second highest bid value. If the aggregate value of the individually-sold items exceeds that of the package, the results of the sealed-bid auction are disregarded; otherwise, the items are sold as a package.

One problem with this procedure is that some bidders have an interest to stay in the Japanese auction well above their reservation price, in order to raise the prices and overcome a sealed bid pre-offer. This may push the price ‘too high’ and cause a bubble effect. The risk of staying in the auction above your reservation price, however, increases as the number of bidders that are participating diminishes. For this reason, information is crucial: players must not know how many others are still competing. Not to communicate explicitly how many bidders are still in the game, however, does not constitute a sufficient form of shielding: one has to be careful that *any* possible information flow be stopped. The click of a key, a ‘blink’ on the computer screen in coincidence with bidders’ drop-out or even a door slamming when they leave the room is enough to give valuable hints to the other participants. These

problems were actually discovered by Plott and his team only thanks to laboratory work, and the Japanese auction design was in fact abandoned for the continuous auction one also for these practical difficulties.

The problem above is one of ‘robustness’. The combinatorial auction, in other words, is not only outperformed by the continuous ascending auction in terms of efficiency, but is also difficult to implement correctly. It is a ‘delicate’ process, and small imperfections (in shielding the flow of information, for example) may cause it to break down altogether. One can distinguish cases of ‘environmental’ robustness like the one above, from cases of ‘personality’ robustness:²¹ economists’ assumptions in modelling economic agents may be *for the most part* accurate, but in more or less marginal instances real behaviour may diverge from the model’s predictions. Environmental robustness, in other words, is a function of an institution’s capacity to work properly in a number of different environments. Personality robustness, in contrast, depends on its capacity to work with a range of real players, who may behave slightly differently from the ideal rational maximising agents postulated by game theoretic models.

Confusion and misunderstanding of the rules are also sources of concern in the implementation of mechanisms, not the least because they may generate legal quarrels. Testbed experiments allowed the spotting of those critical moments when subjects may need help from the auctioneer in order to understand some detail of the auction. The problem is not just whether a certain abstract model fits reality (or whether reality can be made to fit the model), but also *how* it fits. The question is how fragile such a fit is, how sensitive to little imperfections, mistakes, etc.

Another important role of laboratory experiments consisted in helping to develop the appropriate software. Game theory, of course, does not specify whether an auction should be run with mere pencil and paper support, or with computers. Electronic auctions are in general privileged because they facilitate analysis of data and enforcement of the rules: one can, for instance, design a software which does not accept bidding below the highest standing bid of the previous round, thus controlling automatically for errors and saving precious time. But no specific software for the continuous ascending auction was available at the time, and a new one had to be

created ad hoc. The data of testbed experiments, elaborated by means of an independently designed programme, were used as inputs into the final software to be used in the real FCC auctions. By means of such ‘parallel checking’ (Plott, 1997, pp. 627-631) the consultants made sure that the FCC software worked properly and delivered correct data in a reliable fashion. Trained students were employed in order to investigate the properties of the software. They used diaries and notebooks to keep track of all the problems that could arise, and then were asked to answer a questionnaire surveying the most likely sources of ‘bugs’.²² The process must be ‘idiot-proof’: ‘robust’ to the people who use it, who may create a number of bizarre problems (up to crashing the whole network) or make unpredictable moves in playing their game.

The enterprise of building a ‘properly working’ institution in the laboratory thus shows that small variations can make a big difference indeed. “The exact behaviour of the auctions can be sensitive to very subtle details of how the auction process operates” (Plott, 1997, p. 620). To understand why a mechanism may be chosen for reasons that differ from its theoretical properties, one has to recall that those properties are defined at a certain (usually high) level of abstraction. An institution producing desirable allocations when correctly built may nevertheless be very difficult to build correctly.²³

Compared to other applied disciplines, economic engineering is peculiar from at least two respects: first, once the basic causal structure has been chosen (e.g., the combinatorial auction design), the economic engineer has very little room for manoeuvre, in terms of the corrections, modifications, and interventions which are allowed. The materials are for the most part given because agents should not, for instance, be forced to change their preferences – although some kinds of preferences (e.g. collusive or altruistic ones) can be *ruled out* by design (see above). The principal way in which the economist can intervene on the problem-situation is by defining the rules of the institution. Secondly, as we have seen, in the FCC case the designers were short of theory guiding the implementation. It was not just a matter of ‘adding matter’ to an abstract causal structure, but rather of adding matter while understanding and shaping the causal structure itself.

6. Testing the rules

The Caltech experiments were used differently in different phases of the project. First, as we have seen, they were instrumental in choosing among the Milgrom-Wilson-McAfee design and the initial FCC proposal. Secondly, testbeds were used to transform the abstract design into a real process that could perform the required task reliably in the environment where the auction was to be implemented. Finally, experimental data were most valuable in interpreting and checking the results of the first *real* auctions ran by the FCC. Before coming to that, however, it is worth looking at some of the experiments run by Plott and his team in order to check the joint effect of the rules that would have regulated the FCC auctions.

Mechanism designers see the rules as a device for defining the strategic situation which bidders face and to which they are supposed to react rationally. Unfortunately, in the case of spectrum licences, theorists were unable to solve the whole game, and thus had to rely on the analysis of single sub-games in a piecemeal manner. How the pieces would have interacted once put together remained an open issue, about which theorists had relatively little to say. Experiments can be used to move gradually from the world of paper and ideas to the real world, without losing the desired structural properties of a mechanism on the way. The enterprise is similar to that of designing a space probe: it would be too costly to proceed on a trial-and-error basis, and perform a series of full-scale tests.

The structure of the continuous ascending auction has been so far sketched at a most abstract level. The details of the design were in fact much more complicated. Preston McAfee, Robert Wilson and Paul Milgrom were mainly responsible for writing the detailed rules that were to regulate bidding in all its various aspects, and which eventually were put together in a lengthy document. The most important – and debated – rules concerned increments, withdrawals, eligibility, waivers, and activity.

The simultaneous ascending auction proceeds in rounds. At every round, the players offer their bids, which are scrutinised by the auctioneer, and then are presented with the results. The feedback includes a list of the bids presented at that round and by whom, the value of the ‘standing high bid’, and the minimum bid allowed for the next round. The minimum allowed bid is calculated by adding to the standing high bid a fixed proportion (5% or 10%, usually) called the *bid increment*. A bid is said to be eligible if it is higher than the minimum bid allowed in that round and it is presented by an *eligible bidder*. Each bidder, in fact, must at the beginning of the auction make a deposit proportional to the number of licences she wants to bid for (each item is auctioned on a different market). Such a deposit establishes her ‘initial eligibility’, i.e. the number of markets she can enter.

The idea of eligibility was introduced not only to prevent problems such as those that occurred in the New Zealand auctions (where bidders who were not really interested in the licences participated and then resold the licences to others), but also to regulate the time-length of the auction. Eligibility constrains a bidder’s activity in an ingenious manner. A bidder is said to be *active* if either she has the standing high bid from the previous round, or is submitting an eligible bid. The activity cannot exceed a bidder’s initial eligibility, but bidders have also an incentive not to remain below a certain level of activity. In the first round, in fact, a buyer must bid at least a certain fraction of her initial eligibility; if she does not, her eligibility is reduced in the next round. Such a procedure is supposed to increase the transparency of the auction (by forcing companies to commit early on), and to speed it up by preventing ‘wait and see’ strategies on bidders’ part. The possibility of mistakes is taken into account by providing bidders with five ‘*waivers*’ of the activity rules. Bidders could also *withdraw* a bid, but with the prospect of paying the difference between the final selling price as elicited in a further round and the withdrawn bid, in case the latter exceeded the former.

The rules regulating activity were motivated by the worry that bidders could have a strategic interest in slowing down the pace of the auction.²⁴ The time-length of an auction depends on two factors: the number of rounds played and the interval between rounds. According to the Milgrom-Wilson design, each auction is supposed to stop after a certain period and start again the next day, until no new eligible bids are

received. The idea of having subsequent ‘rounds’ was motivated by the thought that companies may need time to revise their strategies, reflect on their budgets, and thus avoid ‘winner’s curse’ phenomena. As Paul Milgrom put it, “there are billions of dollars at stake here, and there is no reason to rush it when we are talking about permanently affecting the structure of a new industry”.²⁵

The FCC, however, was quite concerned about the time problem, not least for practical reasons (booking a big hotel to run an auction, for instance, costs money). Whereas theory does not tell you how long a continuous ascending auction may go on for, experiments allowed the testing of different rules, with rounds taking place at different intervals. One possible solution was to impose big increments above the highest standing bid, so as to identify the winners quickly. In experiments, however, it was observed that big increments sometimes eliminated bidders too quickly, causing their eligibility to drop and therefore creating a ‘demand killing’ effect (Plott, 1997, p. 633). Such an interaction between the increment rule and the eligibility rule could have hardly been predicted without experiments. Without a general theory of simultaneous ascending auctions, theorists could rely only on a number of independent insights about the effects of different rules, but could not exactly figure out what would happen were all the rules implemented at the same time.

The concepts of withdrawals, eligibility, increments, and announcement of stage changes, all involve reasonable sounding concepts when considered alone, but there remain questions about how they might interact together, with other policies, or with the realities of software performance. Can one waive and bid at the same time? What happens if you withdraw at the end of the auction - can the auction remain open so the withdrawal can be cleared? How shall a withdrawal be priced? (Plott, 1997, p. 629).

The answers to these questions were partly sought in the laboratory. “The complex ways the rules interact, and the presence of ambiguities, do not become evident until one tries to actually implement the rules in an operational environment” (Plott, 1997, p. 628). More time between rounds might allow bidders to put forward sensible bids, but more frequent rounds might also shorten the process considerably. And would either of these solutions affect efficiency? Plott and his collaborators eventually found

that total time was principally a function of the number of rounds, and auctions going on for more than one hundred rounds were observed in the laboratory. The Caltech team thus tried to vary the interval between rounds, and concluded that to have more frequent rounds did not affect efficiency in their laboratory experiments (Plott, 1997, pp. 632-633).

The Milgrom-Wilson-McAfee rules also involved the possibility of withdrawals, in case a winner decided *a posteriori* that the price was not worth the item bought, or that she could not achieve the preferred aggregation: the item would be ‘sold back to the market’ and the withdrawing bidder would pay the difference between her bid and the final price to the FCC. Withdrawals and losses can however cause ‘cycles’: an item may be bought and resold, and bought and resold and so on until someone is satisfied with the price. Experiments were used to create cycles in the laboratory (see Figure 1) in order to see whether they could arise in practice, to study the conditions (the parameters) under which they can be generated, and how they behave (for how long they go on before disappearing, for instance).

[Figure 1 ABOUT HERE]

Theoretical and empirical insights, to sum up, provided initially just a few rough ‘causal stories’ about the functioning of certain bits of the institution. Thanks to previous experiments, it was known that more transparent mechanisms (like the English system) tend to prevent winner’s curse effects better than ‘less transparent’ ones (e.g., sealed-bid auctions). It was also known that imposing no entry fee would have encouraged opportunistic bidders to participate. It was conjectured that under some circumstances some bidders might have an interest (opposite to the FCC’s) to delay the auction – and so on. These insights do not jointly pin down the structure of the best possible institution in all its aspects. They just convey information about what certain components do in certain cases, when acting in isolation. Instead of laying down the structure of the auction on paper, it was displayed at work in the laboratory. The job of checking that the different components had been put together in the ‘right’ way (that they *could* be put together in the right way, to begin with), and that they worked smoothly in practice was done by running demonstrations in the lab.

7. Checking external validity

The consultants eventually monitored a *real* FCC auction (which took place in Washington, DC, in October 1994)²⁶ to check whether the transposition of the Milgrom-Wilson simultaneous ascending auction mechanism from the world of theory to the real world *via* the laboratory had been successful. Or, to put it differently, whether ‘external validity’ had been achieved.²⁷ The expertise achieved in the lab was invaluable for this. Experiments had taught that especially in the first rounds “regardless of the amount of preparation and testing, things happen”, and “decisions must be made on the spot from experience and judgement” (Plott, 1997, p. 631). A committee of consultants was thus formed, to supervise the FCC auction and intervene at any time to assist the participants. Any inconvenience that arose in testbeds was likely to be spotted by the trained eye of the experimenters.

Laboratory tests with similar parameters to those supposed to be the case in the real auction were run beforehand so that the results could be compared after the event. The experimental auctions had been constructed so as to have the same number of players, the same number of items auctioned, complementarities such as those presumably existing in the ‘real’ market, a similar number of rounds,²⁸ similar (although fictional) values, etc. Then, a large amount of data collected both in the lab and in the real auction was systematically analysed and compared. This data included bidding patterns, pricing trajectories, the rise of ‘bubbles’, the formation of sensible licence aggregations, the fact that similar items sold for similar prices, the absence of post-auction resale, and several other phenomena of interest.²⁹ To give the taste of this procedure, the ‘equilibrating’ trajectory of an experimental testbed auction is represented in Figure 2.

The dots on the curve stand for the revenues obtained in the laboratory auction, whereas the horizontal line represents the revenues predicted by the competitive equilibrium model, computed on the basis of experimentally induced demand and supply curves. The adjustment process towards equilibrium is one of the most replicated effects in experimental markets: it takes time and repetition to ‘find’

equilibrium prices. In Figure 3 the revenue of the real FCC auction of October 1994 is represented.

[Figure 2 ABOUT HERE]

[Figure 3 ABOUT HERE]

The path is remarkably similar to that of the experimental auction. (The downward sloping curve represents so-called ‘excess bids’ and can be ignored for our purposes.) The price trajectories in the two auctions evolved in a similar way, and since experimenters *knew* that the laboratory auctions delivered approximately efficient outcomes in the long run, they felt confident to argue that an efficient allocation was achieved in the Washington auction as well. The reasoning behind such a conclusion is straightforward:

If indeed the same principles were operating in the FCC environment then the FCC auction converged to near the competitive equilibrium and exhibited high efficiency (Plott, 1997, p. 637).

The external validity argument takes the form of a ‘bootstrapping’ from a set of background assumptions and the observed data to the underlying generating process that produced it. The structure of such a (fallible) inference can be reconstructed as follows:

1. If all the directly observable features of the target and the experimental system are similar in structure;
2. And if all the indirectly observable features have been controlled in the laboratory;
3. And if there is no reason to believe that they differ in the target system;
4. And if the outcome of the two systems at work (the data) is similar;
5. Then, the experimental and target systems are likely to be structurally similar mechanisms (or data-generating processes).

The strength of Plott’s argument lies on the work he and the other consultants did to ensure that the same processes took place in reality as those they had observed in their laboratory. The same causes are supposed to operate because experimenters built

the two systems so as to be structurally similar to one another. The transportation of the mechanism outside the laboratory was as smooth, gradual, and carefully monitored as possible. The bootstrapping inference was based on a solid ground of theoretical, practical, and experimental knowledge.

8. Rational choice technology

The FCC auctions were a success. They were successful not just politically,³⁰ but also from a scientific viewpoint. They provided a solution to a complicated problem, and have inspired other projects as a benchmark example of economic engineering.³¹ This does not mean that the implemented institution was absolutely flawless: it could be improved and some work is still being done on rules, for example to prevent collusive bidding.³² But the solutions devised by the FCC consultants stood on firm ground, probably as firm as our best science and technology can be.

As stated in the introduction, this article is not intended to discuss the various philosophical interpretations of rational choice theory in any depth. That is a project worth a separate paper. However, some tentative remarks can be provided that go in that direction. First, several aspects of this case study seem to rule out an instrumentalist interpretation of the local models that guided economists' work as FCC consultants. To build a successful auction one needs to pay attention to the computational abilities and preferences of its users. One has to make sure that the tasks the bidders face are not too complicated or the rules unclear. Bidders' reactions to possible strategic situations must be analysed in the light of realistic cognitive capacities at the individual level. One cannot just presume that buyers behave 'as if' they were rational. Bidders must react adequately to new situations and sometimes be creative in devising new strategies, as opposed to just relying on established routines.³³ The economic engineer must design a mechanism keeping individuals' *real* capacities in mind. On the other hand, it is by designing and implementing an adequate mechanism that the engineer ensures that rational choice models can work. Since it is partly in virtue of the structure of the situation that economic agents behave rationally, a great part of economic engineering is devoted to ensure that the structure is 'right' (and experiments, as we have seen, are invaluable for that).³⁴

One should also keep in mind that bidding teams in the real FCC auctions included professional game theorists hired by firms in order to maximise their chances of putting together a profitable aggregation.³⁵ This way, the classical game-theoretic assumptions of rationality and common knowledge of rationality (everybody knows that everybody is rational, and that everybody knows that everybody knows...) were most likely to be satisfied. Which suggests a more general point: due to the partly self-referential character of social concepts, adequate self-perpetuating mechanisms must be set in place for the construction of a stable and reliable socio-economic mechanism to be possible at all.³⁶ The institution assumes behaviour with certain formal characteristics on the agents' part, but the each agent must also know these assumptions, and must be confident that the other agents are willing to and capable of fulfilling the mechanism's requirements.

Finally, it is a common strategy to attack rational choice theory by pointing out that it is either false or void of empirical content as a (general) theory of human behaviour. 'General' is between brackets because it is usually taken for granted that absolute generality must be the goal of every respectable scientific theory. But either charge does not do justice to the rational choice approach.³⁷ There *can* be reliable economic mechanisms, as stable, fine-tuned, and robust as we need them to be. But they seem to be hardly spontaneous nor common – they are mostly constructed with ingenuity in the abstract realm of theory, sometimes even repeatedly tried in a university lab. Perhaps there exists some more general account of human behaviour and decision making, which can include rational choice theory as a sub-case.³⁸ Consider that rational choice models seem to possess just enough detail to be translatable into operational mechanisms. But nevertheless a lot of 'filling out' is needed before the mechanism is ready to operate. And that seems to be the case with any account of human behaviour. The question is not merely one of choosing between accounts with a general vs. limited domain of application,³⁵ but rather one of trading off between abstract and concrete models, which need more or less filling out in order to be transformed into useful and robust social technology.

Acknowledgements: I should thank two anonymous referees, Barry Barnes, Phil Mirowski, audiences at the History of Economics Society 1998 and the Political

Theory Seminar of the University of Exeter for their comments and critiques on earlier versions of this paper. Special thanks go to Andrea Salanti, with whom I have started working on this topic (see our 2001 paper). My intellectual debt towards Nancy Cartwright, John Dupré, and the ‘Models as Mediators’ project at the LSE should be clear from the main text above, although I obviously take full responsibility for all the remaining mistakes.

REFERENCES:

- Ayres, I. and P. Cramton (1996) ‘Deficit Reduction through Diversity: A Case Study of How Affirmative Action at the FCC Increased Auction Competition’, *Stanford Law Review* 48, 761-815.
- Barnes, B. (1983) ‘Social Life as Bootstrapped Induction’, *Sociology* 17, 524-45.
- Cartwright, N. (1989) *Nature’s Capacities and Their Measurement* (Oxford: Oxford University Press).
- Cartwright, N. (1999) *The Dappled World* (Cambridge: Cambridge University Press).
- Collins, H.M. (1985) *Changing Order* (London: Sage).
- Cramton, P.C. (1995) ‘Money Out of Thin Air: The Nationwide Narrowband PCS Auction’, *Journal of Economics and Management Strategy* 4, 267-343.
- Cramton, P.C. (1997) ‘The FCC Spectrum Auctions: An Early Assessment’, *Journal of Economics and Management Strategy* 6, 431-95.
- Cramton, P.C. (1998) ‘The Efficiency of the FCC Spectrum Auctions’, *Journal of Law and Economics* 41, 727-36.
- Cramton, P. and J. Schwartz (2000) ‘Collusive Bidding: Lessons from the FCC Spectrum Auction’, *Journal of Regulatory Economics* 17, forthcoming.
- Dupré, J. (2001) ‘Economics without Mechanism’, in U. Mäki (ed) *The Economic World View* (Cambridge: Cambridge University Press), forthcoming.
- Friedman, J. (ed 1996) *The Rational Choice Controversy* (New Haven: Yale University Press).
- Guala, F. and A. Salanti (2001) ‘Theory, Experiments and Explanation in Economics’, *Revue internationale de philosophie*, forthcoming.
- Hacking, I. (1983) *Representing and Intervening* (Cambridge: Cambridge University Press).

- Harrison, G. (1989) 'Theory and Misbehavior of First-Price Auctions', *American Economic Review* 79, 749-62.
- Hausman, D.M. (1989) 'Supply and Demand Explanations and Their *Ceteris Paribus* Clauses', *Review of Political Economy* 2, 168-87.
- Hedström, P. and R. Swedberg (eds 1998) *Social Mechanisms* (Cambridge: Cambridge University Press).
- Ingrao, B. and G. Israel (1987) *La mano invisibile* (Bari: Laterza); Engl. tr. *The Invisible Hand* (Cambridge, Mass.: MIT Press, 1990).
- Lawson, T. (1997) *Economics and Reality* (London: Routledge).
- Ledyard, J.O., D. Porter and A. Rangel (1997) 'Experiments Testing Multiobject Allocation Mechanisms', *Journal of Economics and Management Strategy* 6, 639-75.
- Machamer, P., L. Darden and C.F. Craver (2000) 'Thinking about Mechanisms', *Philosophy of Science* 67, 1-25.
- Mas-Colell, A., M.D. Whinston and J.R. Green (1995) *Microeconomic Theory* (Oxford: Oxford University Press).
- McAfee, R.P. and J. McMillan (1987) 'Auctions and Bidding', *Journal of Economic Literature* 25, 699-738.
- McAfee, R.P. and J. McMillan (1996) 'Analyzing the Airwaves Auction', *Journal of Economic Perspectives* 10, 159-75.
- McMillan, J. (1994) 'Selling Spectrum Rights', *Journal of Economic Perspectives* 8, 145-62.
- McMillan, J. (1995) 'Why Auction the Spectrum?', *Telecommunications Policy* 19, 191-9.
- McMillan, J., M. Rotschild and R. Wilson (1997) 'Introduction', *Journal of Economics and Management Strategy* 6, 425-30
- Milgrom, P. (1989) 'Auctions and Bidding: A Primer', *Journal of Economic Perspectives* 3, 3-22.
- Milgrom, P. (1995) 'Auctioning the Radio Spectrum', in *Auction Theory for Privatization* (Cambridge: Cambridge University Press), forthcoming (available on-line at [www. market-design.com/library.html](http://www.market-design.com/library.html)).
- Milgrom, P. (1998) 'Game Theory and the Spectrum Auctions', *European Economic Review* 42, 771-8.

- Milgrom, P. (2000) 'Putting Auction Theory to Work: The Simultaneous Ascending Auction', *Journal of Political Economy*, forthcoming.
- Mirowski, P. (1989) *More Heat than Light* (Cambridge: Cambridge University Press).
- Pickering, A. (ed 1992) *Science as Practice and Culture* (Chicago: University of Chicago Press).
- Plott, C.R. (1981) 'Experimental Methods in Political Economy: A Tool for Regulatory Research', in A.R. Ferguson (ed) *Attacking Regulatory Problems* (Cambridge, Mass.: Ballinger), pp. 117-43.
- Plott, C.R. (1991) 'Will Economics Become an Experimental Science?' *Southern Economic Journal* 57, 901-19.
- Plott, C.R. (1996) 'Laboratory Experimental Testbeds: Application to the PCS Auction', *Social Science Working Paper 957*, Pasadena, California Institute of Technology.
- Plott, C.R. (1997) 'Laboratory Experimental Testbeds: Application to the PCS Auction', *Journal of Economics and Management Strategy* 6, 605-38.
- Roth, A.E. (1991) 'Game Theory as a Part of Empirical Economics', *Economic Journal* 101, 107-14.
- Salmon, P. (1998) 'Free Riding as a Mechanism', in R.E. Backhouse, D.M. Hausman, U. Mäki and A. Salanti (eds) *Economics and Methodology: Crossing Boundaries* (London: MacMillan), pp. 62-87.
- Salmon, W.C. (1984) *Scientific Explanation and the Causal Structure of the World* (Princeton: Princeton University Press).
- Satz, D. and J. Ferejohn (1994) 'Rational Choice and Social Theory', *Journal of Philosophy* 91, 71-87.
- Schotter, A. (1995) 'A Practical Person's Guide to Mechanism Selection: Some Lessons from Experimental Economics', working paper, C.V. Starr Center for Applied Economics, New York University.
- Smith, V.L. (1976) 'Experimental Economics: Induced Value Theory', *American Economic Review* 66, 274-7.
- Smith, V.L. (1982) 'Microeconomic Systems as an Experimental Science', *American Economic Review* 72, 923-55.
- Smith, V.L. (1987) 'Experimental Methods in Economics', in J. Eatwell, M. Milgate, and P. Newman (eds) *The New Palgrave: Allocation, Information and Markets*, (London: Macmillan), pp. 94-111.

Thaler, R.H. (1988) 'The Winner's Curse', *Journal of Economic Perspectives* 2, 191-202.

Vickrey, W. (1961) 'Counterspeculation, Auctions, and Competitive Sealed Tenders', *Journal of Finance* 16, 8-37.

¹ See Hacking (1983) and Cartwright (1999) for some representative examples of the new, practice-based, philosophy of science of the eighties and nineties.

² A good overview of the rational choice theory debate can be found in Friedman (ed. 1996). On social mechanisms, see note 7 below. I'll leave it to the reader to appreciate to what extent these two debates suffer from the defects highlighted in the main text above.

³ According to the US Department of Commerce, more than 45 billion dollars were gained by opportunistic lottery winners in the eighties (cfr. McMillan, 1994, p. 147, n. 3). For a comparison of the auction method of allocation with other methods such as lotteries or administrative hearings, as well as for a discussion of the pros and cons of auctioning licenses, see McMillan (1995).

⁴ McMillan (1994, p. 147) and Milgrom (1995, pp. 4-11) tell the stories of these earlier design failures.

⁵ *The Economist*, July 23, 1994, p. 70; quoted in McAfee and McMillan (1996), p. 159.

⁶ One of the sources of confusion is whether 'mechanisms' are social entities or representations. The ambiguity is generalised, even in advanced textbooks. Mas-Colell *et al.*, for instance, in the very same page define a mechanism as "the formal representation of [...] an institution" (more precisely: a collection of strategy sets and an outcome function from the Cartesian product of the strategy sets to the set of alternatives), to claim shortly after that "a mechanism can be seen as an *institution* with rules governing the procedure for making collective choice" (1995, p. 866 *my italics*).

⁷ The views on economic mechanisms and machines that inform this essay are closest to Cartwright's (1999) and Dupre's (2001). The distinction between mechanisms and processes is due to Salmon (1984), but see also Machamer *et al.* (2000). On socio-economic mechanisms in particular cf. Lawson (1997) and the essays collected in Hedström and Swedberg (eds. 1998), although Pierre Salmon (1998) and his commentators are perhaps more illuminating.

⁸ I had to amend Plott's text to make it consistent with the chosen terminology.

⁹ Mechanism design is an activity in which normative and positive economics intersect. Despite the philosophical interest of the normative aspects of mechanism

design, I shall in this paper be concerned with the positive ones only, and take normative considerations mostly as given.

¹⁰ The above quotes are from *Fortune*, 6 February 1995, p. 36, cited by McAfee and McMillan (1996), p. 159, who report on other similar reactions. Other references can be found in Milgrom (1995).

¹¹ Notably, to the ‘common value’ case, in which the exact value of each item is the same for every bidder but unknown to all. (Auctions for oil leases seem to be of this kind.)

¹² For an introduction to auction theory, cf. Milgrom (1989); for a more comprehensive survey, see McAfee and McMillan (1987).

¹³ The ‘winner’s curse’ is a much discussed phenomenon in recent economic literature. Winners of certain kinds of auction seem to have a tendency to bid too much, and thus experience losses in the long run. This phenomenon is puzzling for standard game theory, which models bidders as perfectly rational players that should not be affected by this ‘curse’. See Thaler (1988) for an introduction.

¹⁴ The example is presented in Milgrom (1998, 2000).

¹⁵ McAfee and McMillan for instance claim that “there is a direct link between game theory’s laureates and the spectrum auction. The ideas with which Nash, Harsanyi and Selten are associated - Nash, Bayesian and Perfect Equilibrium - are the basic tools of the theory used in designing the auction” (1996, p. 171, n. 15).

¹⁶ Cf. e.g. Hacking (1983), Collins (1985), and the essays in Pickering (ed. 1992).

¹⁷ To question the realisticness of a certain experiment then, leads naturally to question the realisticness of standard economic models. The concept of ‘microeconomic system’ plays the rhetorical role of ‘shifting the burden of proof’ to the theorist’s side.

¹⁸ See for instance the debate initiated by Harrison (1989) and then carried on in the *American Economic Review*, 82 (1992).

¹⁹ Hausman (1989) gives a convincing account of why demand and supply explanations are to be interpreted as causal in character, and how the explaining principles are to be read as including a *ceteris paribus* clause.

²⁰ But see Ledyard *et al.* (1997) for a somewhat different conclusion. John Ledyard is part of a minority of economists who favour the combinatorial over the simultaneous ascending mechanism.

²¹ See Schotter (1995) for this distinction and some useful discussion.

²² The questionnaire included questions like “What happens if you stay logged on after the initial withdrawal; What happens if you log in from multiple locations at the same time; what happens if you enter 0000 rather than 0; what happens if you are theoretically inactive but then you log on after various events; what happens if you log in at the last second of a session or have a power failure; what happens if you follow local software installation exactly to the ‘letter’ of the instructions; etc.” (Plott, 1997, p. 631, n. 4).

²³ In the context of the natural sciences, Nancy Cartwright calls it the ‘*problem of material abstraction*’: “A physicist may preach the principles by which a laser should operate; but only the engineers know how to extend, correct, modify, or sidestep those principles to suit the different materials they may weld together to produce an operating laser” (Cartwright, 1989, p. 211).

²⁴ Milgrom (1998) provides an example and a game-theoretic argument in support of such a hypothesis.

²⁵ “Access to Airwaves: Going, Going, Gone”, *Stanford Business School Magazine*, June 1994.

²⁶ For a detailed analysis and discussion of the early FCC auctions, cf. Cramton (1995; 1997, 1998), Ayres and Cramton (1996), McMillan (1996) and Milgrom (1995). The data of all the auctions run by the FCC is available on-line at www.fcc.gov/wtb/auctions.

²⁷ In experimental economics, the problem of external validity is often referred to as the problem of ‘parallelism’.

²⁸ The FCC, though, interrupted the final testbed experiments before their completion. One last round of bidding was unexpectedly called, and this probably accounts for the sudden rise in prices observable in figure 3 (see Plott, 1997, p. 636 footnote 6).

²⁹ See also Cramton (1995, 1997 and 1998).

³⁰ Al Gore claimed that “now we are using the auctions to put the licences in the hands of those who value them the most” (quoted in Milgrom, 1995, p. 1).

³¹ See for examples the auctions for the third generation of mobile telephony run in the UK in April 2000. Nothing has been said in this paper about the business aspects of mechanism design, but there is certainly an interesting story to be told about game

theory turning into a profitable enterprise. Several economists involved in the FCC project have later created consultancy firms.

³² Some companies, for example, devised sophisticated signalling systems using the last few digits of their bids, in order to communicate to other competitors their willingness to cooperate (Cramton and Schwartz, 2000).

³³ This is not meant to suggest that routines do not play a role in bidding behaviour, nor that real bidders constantly perform all the calculations required by rational choice models. A reasonable interpretation is perhaps that bidders must be able to get in and out of routine behaviour as required by the evolving strategic situation.

³⁴ A similar point on the importance of structure for rational choice models has been made by Satz and Ferejohn (1994), but it is my contention that the FCC case warrants a stronger interpretation than their mild instrumentalism (or ‘moderate externalism’).

³⁵ Peter Cramton (1995; 1997) has provided some interesting accounts of the early auctions, written from an insider’s perspective – Cramton was part of the PageNet bidding team.

³⁶ For a more detailed illustration of this point, see Barnes (1983).

³⁷ To be fair, these charges are partly justified by the fact that economists themselves have traditionally conceived the economy (the *whole* economy, if not society in general) as a gigantic machine. See Ingrao and Israel (1987) and Mirowski (1989).

³⁸ Herbert Simon’s ‘satisficing’ or ‘bounded rationality’ approach would be an obvious candidate for that.