How to Do Things with Experimental Economics[*]

Francesco Guala
University of Exeter
f.guala@ex.ac.uk

Published in *Do Economists Make Markets?*,
edited by Donald MacKenzie, Fabian Muniesa, and Lucia Siu.
Princeton University Press, 2007.

In July 2003 two US senators, Ron Wyden and Byron Dorgan, publicly denounced a Pentagon plan to create an on-line 'market for terror'. Anonymous buyers and sellers would exchange on a governmental website 'futures for terrorist attacks', effectively betting on how likely a certain site or prominent individual were to become a target of Bin Laden's kamikazes. The project, initially allocated an $8 million budget for two years, was defined 'morally repugnant' and 'grotesque' and was quickly withdrawn by the Bush administration following outrage in the media and the public opinion.[1]

The aborted plan wasn't merely an odd combination of army idiocy and market extremism. The market for terrorism was in many ways cutting-edge science. It was inspired by the repeated successful use of electronic markets to forecast uncertain events. The prototype and most famous example, the Iowa Electronic Markets, had been running since the late eighties and had successfully predicted the results of major political events, notably the US presidential elections, with better margin of errors than the latest polls.

The roots of electronic markets lie in experimental economics, a research programme recognised by the 2002 Nobel Prize.[2] The key idea – known as the 'Hayek Hypothesis' – is that markets can be extremely efficient mechanisms for the aggregation of information. Indeed they are able to perform the remarkable trick of transforming imperfect individual information into an efficient market outcome, by means of a signal (a price) that incorporates at once all the preferences and expectations of the individuals in the market.

[*] I would like to thank all participants to the 'Performativities of Economics' meeting for the lively and interesting discussion during and after the workshop. In particular, Vincent Lepinay, Edward Nik-Kah, Phil Mirowski, and Yuval Millo provided many comments which helped to improve the paper. As usual, I am responsible for all the remaining mistakes.

[1] See Hulse (2003).

[2] Two faculty members of the University of Iowa who ran the first electronic market in 1988 (Robert Forsythe and Forrest Nelson) had worked for many years at the California Institute of Technology, one of the pioneering centres for experimental economics. Forsythe is an experienced experimenter himself, as are other current directors of the Iowa project like Joyce Berg and Thomas Rietz. (I'd like to thank Joyce Berg for providing this information, in personal communication.)

The Hayek Hypothesis was a key weapon in the hands of free-market apologists during the so-called 'socialist calculation debate' of the thirties, but for a long time was little more than a speculation based on fairly abstract philosophical assumptions.[3] General equilibrium models in the Walrasian tradition moreover seemed to have cast doubt on the hypothesis, by relying heavily on perfect information and other unrealistic assumptions to prove the theoretical existence of efficient equilibria; until, very recently, the Hayek Hypothesis was corroborated in a series of ingenious laboratory experiments. The authors of such experiments were leading economists like Vernon Smith and Charles Plott, who since the sixties had devoted their career to constructing little 'flesh and blood' markets (with real human subjects) in their university labs.[4] Economists were also quick to exploit the opportunities provided by the internet revolution, and the Hayek Hypothesis soon took the very concrete form of future markets for events of all sorts – from the results of political elections to the Oscars, and indeed – were it not for the two senators and a hypersensitive post-9/11 public opinion – for terrorist attacks.[5]

But these are neither the only nor the most important applications of experimental economics to date. Since the eighties experimental economists have designed mechanisms for the allocation of airport slots (Grether, Isaac and Plott 1989), for the pricing of space stations (Plott and Porter 1996), for the regulation of inland water transportation (Hong and Plott 1982), of the gas industry (Grether and Plott 1984), and of gas transportation networks (Plott 1988); for the construction of the new Arizona Stock Exchange (Smith and Williams 1992), for the regulation of the market for new physicians and surgeons (Roth and Peranson 1999), and for the allocation of telecom licences (Plott 1997, Guala 2001). The list is incomplete and likely to grow in the next few years. The experimental game theorist Ken Binmore, who co-designed with Paul Klemperer the widely acclaimed 3G mobile phone auctions in the UK,[6] foresees applications in some key areas of the welfare state. How about a market for hospital beds?

> One way to do that would be to run a computerised market. Not the idiot internal market of Mrs Thatcher … I cannot imagine a more irresponsible experiment … a real computerised market so each morning someone from each hospital can update their screen and say what they are willing to buy and sell beds for because you have to have an exchange of real resources for this to work. (Binmore in Atkinson 2000, p. 22).

Given the current political climate, this is not an unlikely prospect. Experimental economics is a relatively rare instance of social science that *works*. I mean 'work' in a very broad way: it works sociologically, for after a low-key start it has been generally

[3] Which does *not* mean that it wasn't successful in winning the argument on the (im)possibility of a centrally planned efficient economy.
[4] Smith (1982a), Davis and Williams (1991), Plott (2000).
[5] See for instance the Austrian Political Stock Market, the Election Stock Market at the University of British Columbia, and the Hollywood Stock Exchange (all internet addresses are in the bibliography).
[6] The auctions raised £22.5 billion for the government, or approximately 2.5% of UK GDP; see Binmore and Klemperer (2002), Klemperer (2004).

accepted within the scientific profession, as recognised by the Nobel Prize.[7] It works also scientifically, in the sense that it generates replicable results, allows one to make fairly precise predictions, and seems to provide strong insights in the mechanisms that govern market behaviour. Experimental economists have even become so arrogant to name their society the *Economic Science Association*. (The subtext regarding the rest of economics is pretty obvious and astonishing - this is not a small club of cranks or heterodox social scientists.)

But most importantly, as we have seen, experimental economics can be (and has been) used effectively to *intervene*, to change the institutions that regulate and coordinate economic behaviour. Which takes me finally to the vaguely Austinian title of this chapter. Economists traditionally do things *with* models and field data; these are their basic tools, upon which their persuasive powers crucially rely. In order to do things 'with', they do things *to* models and data – they manipulate them, analyse them, and try to show 'what would happen if' such and such a policy were to be implemented in such and such circumstances. With experimental economics in contrast you can *do things to the economy*. You can manipulate and intervene in the micro-economies you have built in your laboratory, and this activity in turn is instrumental to intervening in real-world, full-size markets.

This view of laboratory experimentation as a tool for changing and building economic entities emerged slowly and with difficulty over the years, having to struggle against a tradition that sees experiments in quite a different light, as aimed mainly at theory-testing. Actually the struggle is not quite over yet, although the building tradition seems to gain more and more momentum. Part of this paper is devoted to drawing the contrast between these two traditions. I hope you will excuse my schematic attempt at reconstructing some very recent history of science, for it is eventually aimed at making a point about the core theme of this volume. The builders, I want to claim, are winning because they have understood performativity – or at least they have learnt how to use it constructively. The testers have chosen a weak model of social science, according to which performativity is primarily an impediment for scientific research. On the contrary, performativity is a *resource* for the social scientist, and a very powerful one too.

1. Testers and builders

That one could do things with experiments was by no means obvious in economics only a couple of decades ago. The relatively quick breakthrough of experimental economics within the current (broadly neoclassical) paradigm would be an interesting topic for a historian of science, but is a story that still remains to be told. What we have got instead is a series of recollections by the main protagonists – published partly for the sake of historical record, partly for propaganda, partly for blatantly self-serving purposes – which

[7] There is still some resistance, in the form of routine arguments about the absence of laws in the social sciences, the 'fact' that human beings are 'free to choose' etc. (see for instance Economics Focus 2001). But these are by now rear-guard skirmishes in a battle that has been largely won by experimental economists.

have crystallised in a sort of 'official' history of experimental economics.[8] What follows is by no means intended to fill the gap but rather to sketch the minimal historical background, without which it is difficult to appreciate the significance of the laboratory revolution in economics.

I'll skip the usual and useless attempts to trace the first prehistoric experiment back in time. Whether it was some betting experiment in the eighteenth century or a laboratory study of consumer theory in the nineteen-thirties has little importance, because experimental economics in its present form is entirely a post-WWII phenomenon. The mythology of the discipline customarily identifies three foundational moments. The first one is the Santa Monica conference held in 1952, where the newborn American community of experimental game theorists met shortly after the publication of von Neumann and Morgenstern's *Theory of Games and Economic Behavior* (1944). Many contributions described how real human beings (as opposed to perfectly rational agents) behave in simple social dilemma and bargaining situations. The second foundational event was almost simultaneous: the 1952 conference held in Paris where the earliest empirical counterexamples to von Neumann and Morgenstern's expected utility theory were presented to an audience of distinguished economists and statisticians. Like the Santa Monica conference, the Paris meeting was prompted by the publication of *The Theory of Games*, but focused more specifically on individual decision making. The third event is slightly posterior and significant almost only with hindsight: in 1962 Vernon Smith (later to become Nobel Laureate, in 2002) published his first experimental paper on the equilibrating and efficiency properties of a market governed by a double oral auction institution.[9]

A story based on the foundational myths leads pretty straightforwardly to identify three main currents within contemporary experimental economics: game-theory experiments, decision-theory experiments, and market-experiments. This is as accurate a taxonomy as many others, but turns out to be not very useful for my purposes. In order to understand what goes on in experimental economics, I believe, it is more useful to identify *two* distinct approaches: I shall call them respectively the *theory-testing* (or 'testing', for short) and the *institution-building* (or 'building') approach. This classification does not cut at the level of the theories that drive experimental research, but at the deeper level of the purposes of experimentation itself.

Roughly, the testers see experimentation through the spectacles of a philosophy of science textbook of the sixties; the builders, in contrast, are interested in 'doing things

[8] Examples can be found in Smith (1991a; 1992), Davis and Holt (1993, Ch.), Friedman and Sunder (1994, Ch.9), Kagel and Roth (eds. 1995, Ch.1). Leonard (1994) is the only study by a professional historian that I know of, but focuses on bargaining experiments only. Mirowski (2002) reconstructs the milieu of mid-twentieth century economics, where the conditions for the birth of experimental economics were created, and devotes a short section to Vernon Smith's experimental research programme (pp. 545-551). Two PhD dissertations at Notre Dame are beginning to investigate the origins of experimental economics and of the mechanism design tradition (Lee 2004, Nik-Kah 2004).

[9] The proceedings of the Santa Monica conference are published as Thrall, Coombs and Davis (1954), those of the Paris conference can be found in CNRS (1952). Vernon Smith's first experimental paper is Smith (1962).

with experiments'. (I shall explain what this means in more detail soon.) This dichotomy maps onto the traditional three-fold classification rather straightforwardly. Most of decision theory belongs to the testing approach, and most market experiments fall in the building camp. Game theory experiments are spread across the divide.

It is important to stress that the two groups are not neatly separated, and many people work in both traditions. Perhaps it is more accurate to say that testing and building refer to two logics of experimentation rather than two tribes. The building/testing taxonomy helps to understand, to begin with, why experimenters happen to have a mixed reputation in neoclassical economics. On the one hand, experimental results are frequently invoked by the enemies of neoclassicism, as providing the ultimate evidence (what is stronger than experimental evidence, after all?) that the received theory is deeply flawed. On the other, experimental economics is often cited as a source of most stunning confirmations of the standard theory; indeed experiments are found at the frontier of some impressive market reforms inspired by neoclassical economics.

Part of the explanation is that testers and builders tend to have very different agendas. Testers are often also dissenters; they look for refutations of the standard theory in the laboratory, and they find plenty. Builders have a more cautious attitude, they work inside the orthodoxy and tend not to make bold claims that may scare their fellow neoclassical economists. They also find lots of anomalies, but strategically highlight the discoveries that are broadly consistent with the neoclassical spirit.

The testing approach moreover tends to transgress disciplinary barriers. The testers of decision and game theory work in close contact with experimental and cognitive psychologists. They sometimes call themselves 'behavioural economists', by way of a contrast with the neoclassical habit of reasoning from models rather than from empirical data. Behavioural economists rely on various sources of empirical evidence, including laboratory experiments. They try to construct alternative models of human decision-making that usually depart from the standard assumptions of rational behaviour and are more firmly based on the data. In general, they don't get along very well with mainstream economists.[10]

Once upon a time the rhetoric of theory-testing was prevalent. In order to locate the shift in the balance of power between the testing and the building traditions it is useful to examine the official propaganda of the discipline in methodological articles, presidential addresses, and books (the sort of stuff economists don't normally write, unless there is a very good political reason to do so). It is difficult to say exactly when the shift took place, but my hunch is that it is fairly recent. Although pioneers like Charles Plott have been writing from a 'building' perspectives since the early eighties (see e.g. Plott 1981), the

[10] The history and current practice of experimental economics cannot be fully understood if one does not take the divide between psychology and economics seriously. The builders fought the battle on two fronts: at a purely rhetorical level, as mentioned in the main text, but also at a more subtle methodological level, by introducing standards of experimental validity that are at the same time more strict, more formalised, and more in line with the usual assumptions of economic theory, than those customarily adopted by experimental psychologists (I'll discuss these standards in section 4 below).

testing rhetoric is still prevalent in methodological overviews like Smith (1989) or Smith, McCabe and Rassenti (1991). In the nineties, slowly, titles like "The Economist as Engineer" (Roth 2002) or *Paving Wall Street: Experimental Economics and the Quest for the Perfect Market* (Miller 2002) have begun to appear more and more frequently.[11] In his post-Nobel writings, Vernon Smith engages in an overt apology of the instrument-building tradition in the natural sciences, from which economists have much to learn, in his view: 'I think all sciences are influenced far more by the machines builders than either the theorists or experimentalists' (Smith 2002a, p. 69); 'it's the machines that drive the new theories, hypotheses, and testing programs that take you from atoms, to protons, to quarks' (2002b, p. 105).

2. An example: social dilemma experiments

What kind of 'machines' can be built in an economic lab? Since most people are not familiar with economic experimentation, it is worth illustrating by means of a simple example. Social dilemma experiments are a good case because like most game theory experiments they cut across the divide between the testing and building traditions. They are also among the most replicated experiments, and happen to be widely popular outside of economics too. In a social dilemma situation an agent acts under the influence of two considerations pulling in opposite directions. Rational strategic considerations suggest that the individual payoffs are maximised by following one specific strategy (the 'free riding' strategy), but it is easy to see that the rational strategy leads to a socially inferior (Pareto-inefficient) outcome if universally followed. The simplest and most popular social dilemma situation is the one-shot prisoner's dilemma game (Table 1).

	Left	Right
Up	5, 5	0, 10
Down	10, 0	2, 2

Table 1

The first number in each cell represents the payoff of the row player, the second one of the column player. Here the free riding strategies are Down for Row and Right for Column, leading to a payoff of 2 units each. The reasoning behind this solution (the Nash-solution or Nash-equilibrium of the game, from the mathematician John Nash, recently celebrated in the Hollywood movie *A Beautiful Mind*) is simple: whatever Column may do, Row is better off by playing Down; and similarly Column is always better off by playing Right, regardless of what the opponent does. But, somehow paradoxically, both would be better off if they played Up-Left.

[11] See also Plott (1987, 1994), Smith and Williams (1992), Roth (1991), Schotter (1998), Milgrom (2004), Klemperer (2004).

The game in the laboratory is usually played simultaneously via computer networks and without the possibility of binding agreements; customarily players are also denied face-to-face interaction and the possibility of communication. I say 'customarily' because in three decades of experimentation almost every possible variation in the set-up has been explored, and it has been discovered that different arrangements have significant effects on the results.[12] I can't review these results in detail here, but as is well known in a 'standard' social dilemma experiment a considerable number of subjects play cooperatively (Up-Left, in the game above) contrary to the prediction of standard economic theory.

What does this mean? Most testers are pretty adamant that this is a falsification of the standard theory. What ought to be done, surely, is to reject the theory, and replace it by a better one that is able to account for this and other robust empirical anomalies.[13] Builders have a more sophisticated attitude: they begin by noticing that several subjects cooperate, but many others free ride. Then they ask what can be done to put them in line – whatever the 'line' is. Under what conditions does everybody's behaviour converge on the Nash equilibrium? And under what conditions does it converge to the Pareto optimum? How can we help people to achieve a desirable distribution of the payoffs (once 'desirable' has been defined precisely enough, of course)?

We shall examine some of the tools that builders use for this purpose later on. Here I would like to notice that the testing tradition, despite all the anomalous evidence it has accumulated, has been quite ineffective in defeating the standard theory. Why? A standard answer is that economists are simply not good scientists, that they are hopelessly influenced by their ideological commitments, or something along these lines. Another line, the one that I will follow here, is that in order to be successful you need to learn to *do things* with experimental economics, and the testers have not been very successful at that. The builders are way ahead from this respect, and builders are not interested in refuting the received theory. They rather want to *use* it, alongside with many other tools (such as, crucially, experiments) to perform the economic world.

3. Performativity as experimental bias

The debate on performativity in social science tends to focus on two interrelated issues. The first one is a general ontological thesis concerning the nature of social entities. Social properties are extrinsic properties of a special kind: they depend on the context, and in particular on what other human beings know, believe, or in a single word *intend* about the entity in question. The Queen (the individual Elisabeth Alexandra Mary Windsor) is the queen only if she is widely recognized as such. Her legitimate claim to the throne, moreover, depends on the original performance of a series of actions – including

[12] The most common format nowadays is probably not the prisoner's dilemma game but the so-called public goods game, where subjects play in groups of several players and have to decide how much money out of a given sum to contribute to a 'public project', knowing that the latter will produce some revenue that will be divided equally among the members of the group independently of their individual contributions. See Ledyard (1995) for an introduction and survey of results.

[13] For a representative example of this attitude cf. e.g. Kahneman, Tversky and Thaler (1986), Dawes and Thaler (1988).

linguistic utterances (think of a crowing ceremony for example) – which literally brought her social status into being. The contrast here is with natural properties (the structure of the molecule of a given substance, say) that are what they are (and constrain what you can do with them) independently of what people say or think about them. Of course this does not mean that certain natural properties are not created or brought about by human intervention, nor that human beliefs do not play a part in the genesis of these properties. But the idea is that natural properties are not 'made up' of such beliefs in the same way as social properties are.[14]

The second focal point in the debate concerns the process (or processes) by means of which social entities and phenomena begin, continue, and cease to exist. This being an empirical issue, no entirely general story can probably be told. But for the same reason, this is also where most of the interesting action is, from a sociological point of view. Most empirical studies tend to highlight the positive feedback effects (or 'looping' effects, to use Ian Hacking's (1999) expression) of social concepts: if by saying that X has the (social) property Y we induce people to treat X as if it had property Y, then property Y may well come into being. The social sciences, of course, can play an important role in such processes:

> The social sciences seek to refer to referring activities in general; the social sciences are particular instances of referring activities. Accordingly, full independence of knowledge and its referents cannot be hoped for in the social sciences. (Barnes 1983, p. 524)

In one of the best applications to economics so far, for example, Donald McKenzie has analysed how a 'looping effect' of this kind led to the self-referential verification of the Black-Scholes theory of efficient financial markets in the 1990s. The theory, to use Austin's terminology, 'performed' the market by helping creating and sustaining the entities it postulated. The markets were reformed and reshaped by regulators keeping the theory in mind; and the pricing model of the theory was widely used by stock-brokers when they operated in the market itself. This case study is particularly rich because it also provides an example of a destructive or 'counter-performative' effect, the bubble created by the LTCM investment fund which led to its eventual failure in 1998 (McKenzie 2001; McKenzie and Millo 2003; MacKenzie, this volume).

Michel Callon (1998) has recently proposed to extend the feedback story to account for the whole of economics. (Indeed, economics broadly construed so as to include accounting, marketing, management, and other disciplines customarily considered peripheral to the 'hard core' of neoclassical economic theory.) 'Economics, in the broad sense of the term, performs, shapes and formats the economy, rather than observing how it functions' (Callon 1998, p. 2). Callon's project is also based on a series of detailed empirical studies of how economics performs the economy; despite the boldness of the

[14] There are now various philosophical analyses of the ontology of the social world that follow this line of thought. Barnes (1983) provides one of the earliest and best discussions in my view. Searle (1995) is one of the most recent and popular ones. See also Gilbert (1989), Tuomela (2002).

thesis, again, the presumption is that at the micro-level there may be no general story to be told here.

A distinctive feature of contemporary work on performativity is its awareness of both its creative (positive) and its destructive (negative) aspects, whereas older accounts tended to focus on the negative side only. A classical example is Robert K. Merton's (1957) seminal discussion of the self-fulfilling prophecy. The rumour suddenly spreads that a bank is about to become insolvent. As a consequence, clients begin to withdraw their money from their accounts. Soon, the rumour turns into reality: the bank really *is* insolvent, 'merely' because people have become convinced that it is. Similarly, consider the much-discussed problem of predicting the results of an election. A prediction (based on a poll), once made public, may trigger a 'bandwagon' or an 'underdog' effect that will falsify the prediction itself.

The 'Mertonian' approach sees performativity as a threat, both for society and for social science. For society, because it may lead to disastrous results such as the failure of a financially sound bank. For social science, because it blurs the boundary between what scientists say about reality and reality itself, and in many cases seems to be an obstacle towards the use of social science for the prediction of future events.[15] Something like the Mertonian approach can be found in experimental economics too. Performativity worries are typically raised in the theory-testing tradition, and take the form of concerns about the representativeness of the sample of subjects. The standard 'laboratory rat' in experimental economics is the undergraduate student. And for obvious reasons of access the (self-selected) samples used in most experiments are largely made of *economics* students. The worry then is: do these individuals behave like everybody else? Isn't the fact that they are taught economic theory a source of bias in the experiments aimed at testing the theory itself?

This issue has been famously raised in the context of social dilemma experiments. Two experimental psychologists, Gerald Marwell and Ruth Ames (1981), first presented evidence that economics majors play the cooperative strategy less often than non-economics students in games of this kind. One tempting explanation is that they behave as free riders because economic theory tells them that that's the way in which people generally behave. Furthermore, the theory tells them that that's the *rational* way to behave. *Homo oeconomicus*, if this interpretation is correct, would turn out to be a straightforward effect of economic theory itself. But since not all people have a degree in economics, the Marwell and Ames result opens serious doubts about the generalisability of neoclassical economic models based on the assumption of rational selfish behaviour. It also raises the issue of the validity of laboratory experimentation itself: if people are so diverse that they behave in widely different ways depending on their cultural and

[15] It is interesting to see how this purely negative perspective was superseded in different areas of the social sciences. In economic theory, a simple fixed-point theorem can be used to demonstrate the very possibility of positively self-fulfilling predictions (Grunberg and Modigliani 1954, Simon 1957): a solution is logically possible, therefore the problem has been solved. In sociology, Krishna (1971) first argued that as a matter of fact social reality itself is the result of a massive and extraordinary complex series of performative acts or self-fulfilling prophecies. The interesting task, then, is to investigate the *robustness* of social entities (rules, norms, institutions) to changes in the beliefs and desires of individuals and groups.

educational background, how useful can these tests be? Performativity becomes a problem both for the theorist and the experimenter, from a theory-testing perspective.[16]

Is this interpretation of the Marwell and Ames result correct? Successive studies have replicated the significant difference between the behaviour of economics and non-economics students. But they have also cast doubt on the performativity interpretation of this phenomenon. The most plausible explanation of the Marwell and Ames result points towards a *selection* effect: the sort of people who tend to behave more individualistically are also those who tend to do economics degrees.[17] In a recent study Frey and Meier (2003) have found that freshers about to start an economics degree tend to be less cooperative even *before* they have attended their first economics class. This lower propensity to cooperate remains constant throughout their university career: teaching does not seem to make much difference to the way people behave (although interestingly the propensity to free ride tends to diminish slightly, but significantly, during PhD years).

The prominent game theorist Ariel Rubinstein also argues on the basis of pre- and post-class test results that teaching does not influence the way in which students behave – and fortunately so, he is keen to remark (Rubinstein 1999). Moreover experimenters' informal experience as well as the systematic analyses of subjects' debriefing interviews suggest that students often *believe* they are following the theory, whereas in fact they are not.[18] Indoctrination, again, seems to have less effect than one may initially have thought. This of course does not provide much relief to neoclassical theory – people after all still behave in various ways that differ widely from the theoretical prediction. But it suggests that those who behave as predicted by economic theory do not do it because they have been taught to do so. Which, in turn, shows two things: first, it's probably not so easy to make people behave as economists think they should. Secondly, if the performativity hypothesis is true – if economic theory helps shaping the economy, as Callon, McKenzie and others suggest – it must be in a more subtle way.

In order to figure out how, we have to look more carefully at the nature of economic models and theories. Economic theory does not merely describe how people behave; it describes how a specific kind of individual behaves *in some highly specific types of environment*. An economic model is a detailed description of the sort of circumstances that must occur for the interaction between agents of a certain kind (individualistic maximisers of their own utility) to produce outcomes of a certain kind (efficient market equilibria, typically). Among experimental economists, the builders have devoted more energies than anyone else to studying the institutional *structures* that govern market trading. The most innovative contribution of experimental economics lies in this area of

[16] This concern for problems of representativeness is quite typical of experimental psychology, whereas economists tend to worry about financial incentives and downplay representativeness (see Lowenstein 1999). We shall come to incentives shortly.

[17] They also tend to be male, and boys on average free ride more in social dilemma games.

[18] As an amateur experimenter, I remember several conversations with subjects who claimed confidently that they had followed the theory by cooperating in the initial rounds of a finitely repeated social dilemma game and by free riding in the last one. They usually appeared puzzled when I recalled what the theory actually says (that you should free ride right from the start). Apparently backward induction arguments are very difficult to digest.

research rather than in the study of individual decision making. And here performativity takes a rather different, more interesting and more complicated form.

4. How to do things with preferences

I have located the origins of the building tradition in the market experiments of Vernon Smith and his colleagues. The distinction between builders and testers however applies to the origins of experimental economics only with hindsight, for market experiments initially used to be presented as attempts to test the theory of competitive equilibrium. But how do you test such a theory? Economic theories are sets of models, and models are notoriously tricky entities. A literal reading of neoclassical models of competitive markets in the Walrasian tradition, for instance, leads to the rather paradoxical conclusion that they cannot describe any economy that does (or even can possibly ever) exist. There is nothing special with economic models, from this respect: classical mechanics describes the behaviour of dimensionless mass points and perfectly rigid objects, ignoring electromagnetic effects and the influence of other non-gravitational forces. Similarly, neoclassical equilibrium theory analyses the properties of frictionless markets populated by perfectly rational, perfectly informed agents trading homogeneous non-complementary goods.

The idealisation that has attracted most interest in experimental economics lies at the institutional level. Competitive markets in the real world are organised in various ways, in the sense that different systems of (explicit and implicit, formal and informal) rules regulate the interactions between buyers and sellers. If one is interested in issues of general equilibrium – as Walrasians are – it is obviously necessary to simplify and represent these different institutions by means of a single device. Walras introduced for this purpose an ideal auctioneer who collects from buyers and sellers the quantities they would be willing to trade at a given price. The auctioneer then adjusts the proposed price if the quantity offered falls short or exceeds the quantity demanded, until the two quantities coincide. The 'equilibrium point' is the price at which trading eventually takes place (the price that makes the market 'clear') and under various restrictive conditions can be shown to have the well-known efficiency properties formally defined by Pareto and his followers.

Although a few market institutions are vaguely similar to the Walrasian auctioneer (Walras himself was allegedly inspired by trading at the Paris stock exchange), the latter is largely a fictional entity, because no real market uses *tâtonnement* to determine prices. But the auctioneer has the advantage of being an entity of which the theory's equations are true: if such an institution existed, then Walrasian equilibrium theory would fit it perfectly. Indeed, Walras in the fourth edition of the *Elements of Pure Economics* seems to suggest that the term '*tâtonnement*' refers to the technique of solving a system of simultaneous equations by iteration.[19] The motivation behind the use of *tâtonnement* is probably more mathematical than empirical in character. But then either equilibrium theory is supposed to apply only to markets governed by (something very similar to) the Walrasian auctioneer; or the auctioneer really just 'stands for' a whole class of different

[19] On the Walrasian auctioneer and its various possible interpretations, see de Vroey (1998).

institutions which are supposed to deliver the same result (efficient equilibria, clearing markets) by means of different rules and procedures.

The first interpretation is pretty uninteresting for a theory aiming at policy relevance (remember that general equilibrium was used pretty unashamedly in very concrete political battles like the socialist calculation debate) because the institutions that are very similar to the Walrasian auctioneer are rare. But then is the second interpretation *true*? Are different market institutions equivalent? Economic theory was surprisingly silent on this issue until very recently, for a variety of reasons (see Mirowski, forthcoming) including the lack of analytical techniques to deal with it rigorously. In the sixties and seventies game theorists began to construct models of auction systems that seemed to provide some insight in the way different institutions work.[20] But then the same question could be raised again for these game-theoretic models: do they characterise correctly the functioning of *real* markets? Are they empirically adequate?

An obvious way to test this proposition is to observe different institutions at work. This sort of empirical testing however is difficult in non-laboratory circumstances. A major problem with field data is that some key variables of economic theory, like agents' preferences, are not directly observable. If you are interested in explaining, say, price variations in a market, in order to derive the demand and supply schedules (two crucial explanatory factors) from the observable data, you have to rely on auxiliary assumptions that are usually as difficult to test as the main research hypothesis itself.

Subjects' preferences and beliefs are directly unobservable in laboratory experiments too, of course, but can be more easily *controlled* therein. The way in which experimental economists try to do that is by *paying* their subjects. The idea of using monetary rewards often generates hilarity among non-economists ('Hey, these guys pay their subjects to behave like economists would like them to behave!'), whereas the absence of incentives is dismissed by economists equally bluntly ('What can you learn from cheap talk? Put your money where your mouth is!'). Indeed, the presence of 'adequate' monetary incentives (we shall see what 'adequate' means shortly) has become *de facto* a prerequisite for publication in economics journals – and, conversely, the lack of incentives is considered a sufficient condition for the rejection of an experimental paper.

Social, cognitive and economic psychologists tend to apply a less rigid policy. Many experiments in these areas are performed with incentive structures that would be considered inadequate in economics, and often lack monetary incentives altogether. [21] Early economic experiments (even 'paradigmatic' ones like Smith 1962, or Allais 1953) also lacked what contemporary experimental economists consider an 'adequate incentives structure'. The norms regulating financial incentives were codified later, in a series of papers written in the late seventies and early eighties by Vernon Smith (1976, 1982b) and

[20] See e.g. Vickrey (1961), Wilson (1977), Milgrom and Weber (1982).
[21] Unsurprisingly, then, the issue of incentives is often couched in terms of 'the economics-psychology methodological divide'. I don't want to review of this more general debate here, but see e.g. Cox and Isaac (1986), Hogarth and Reder (eds. 1986), Smith (1991b), Loewenstein (1999), Rabin (1998, 2002).

Luis Wilde (1980).[22] The use of incentives is regulated by four of the five so-called '*precepts* of experimental economics':

1. *Nonsatiation*: the medium of reward is such that of two otherwise equivalent alternatives subjects will always choose the one yielding more of the reward medium.
2. *Saliency*: the rewards are increasing in the good and decreasing in the bad outcomes of the experiment.
3. *Dominance*: the rewards dominate any subjective costs associated with participation in the experiment.
4. *Privacy*: each subject in an experiment receives information only about her own payoffs.

The fifth precept (*parallelism* or external validity) is mostly (although not entirely) independent from incentives issues, and I shall ignore it in this paper.[23] The precepts form the core of so-called 'Induced Value Theory' (Smith 1976), and are to be interpreted as 'a *proposed* set of *sufficient* conditions for a valid controlled microeconomic experiment' (Smith 1982b, p. 930, my emphasis). The precepts were proposed as hypothetical conditionals ('if you want to achieve control, you should do this and that'), and should emphatically not be taken as axioms valid a priori. 'The truth of these precepts can only be established empirically' (Smith, 1982b, p.930, n.10).

The precepts provide broad guidelines concerning the control of individual preferences, which may be implemented in various ways, and which may require ad hoc adjustments depending on the context and the particular experimental design one is using. In fact, money or financial incentives are never mentioned in the precepts. The principles only state in abstract terms what kind of properties an appropriate reward medium should have, but do not say what the medium should be. Money may be one way of implementing the precepts, but not necessarily the only one. In light of the fairly rigid interpretation that has become prevalent in experimental economics, the Smith-Wilde precepts appear distinctively liberal in their original formulation.

Even more importantly, the precepts were originally supposed to apply to market experiments only. In his seminal Induced Value Theory article Smith states explicitly that the principles apply 'to experiments designed to test price theory propositions conditional on known valuations. Separate experiments can be designed to test propositions in preference theory' (Smith 1976, p. 275). In order to explain their rationale, Smith couches the precepts in a conceptual framework borrowed largely from the mechanism design theory of the sixties and seventies. A so-called 'microeconomic system' is analysed into three major components: the environment, the institution, and the outcome. The outcome (the behaviour of the agents in the market) is modelled as a function of the environment and the institution. The institution is basically (I'm simplifying here) a set of rules governing behaviour by setting incentives, punishments, and their enforcement. The environment is a complex set of factors including the commodities to be exchanged, the

[22] The idea of using monetary rewards was borrowed, somehow ironically, from the work of two psychologists (Fouraker and Siegel 1963).
[23] But see Guala (1999, 2003, and 2005 Ch. 7).

agents in the market, their individual endowments, their utility functions, and the technology (costs).

In order to study empirically the effect of these factors on the outcome (the sort of prices that are generated in a market defined by a certain environment and a certain kind of institution, for example), the ability to control preferences is crucial. By controlling preferences, for example, one can try to systematically vary the supply/demand schedules in a given institution, and observe the results of such variations. Alternatively, one can keep the preferences fixed 'in the background' and observe the effect of using different institutions in a given environment (cf. Smith 1982b, p. 927).

A typical application works as follows. Suppose you want to induce in your experiment supply and demand schedules like those of Figure 1. (The 'swastikas' are the discrete, experimental counterparts of the perfectly smooth curves of textbook equilibrium theory.) The customary way of achieving this goal is by assigning your subjects some definite roles in the experiment, dividing them in groups of buyers and sellers with well-defined reservation prices. The reservation price of a seller can be interpreted as the cost of production for each unit of the exchange good. The reservation price of a buyer can instead be seen as the price the experimenter is willing to pay each buyer for a unit of the good once the experimental market is closed.

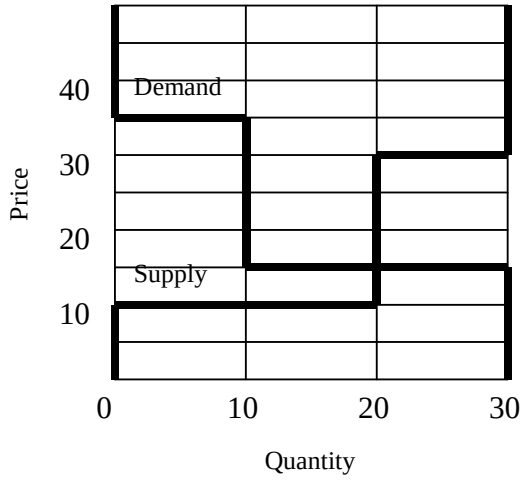

Figure 1

The supply/demand schedules of Figure 1 can be 'induced' by setting reservation prices as in the table below (assuming that each buyer can exchange only one unit of the good during the experiment).

No. of subjects	Reservation price
10 sellers	30 tokens
20 sellers	10 tokens
10 buyers	35 tokens
20 buyers	15 tokens

Table 2

Notice that the prices are expressed in experimental tokens. The key move, according to the precepts of induced value theory, is to make sure that the tokens will be exchanged (privately) at the end of the experiment for some other reward medium, at a rate that satisfies the criteria set out in the precepts themselves. Hence the habit of using real money, in quantities that are likely to dominate all other costs of participating in the experiment.

If this sort of control is effectively achieved, it is then possible to compare the effect (the outcome) of different institutions while keeping the preferences (the environment) fixed. To an observer this may seem a big 'if' but experience shows that it is quite easy to convince people to try to maximise the experimental payoffs. Whether monetary rewards play a crucial role or not (whether role-playing is a key factor, for instance) is obviously debatable.[24] But here the approach is what matters: market experiments work by creating *homini oeconomici* in the lab, not by questioning their existence.

5. Explaining anomalies away

I would like now to use a typical SSK trick, and examine a controversy raised by the (mis)application of Induced Value Theory. Induced Value Theory can be seen as a turning point in the history of experimental economics. A rigid implementation of the precepts makes no sense in the context of other (non-market) economic experiments, yet the precepts inform standard methodological practice in *all* areas of experimental economics. Of course this causes some friction, and scientific friction is very helpful to bring the tacit commitments of scientists into the open.

One obvious motivation behind the indiscriminate application of the precepts is economists' desire to mark a methodological distinction between what they do and psychologists' experimental practice. But there may be more to be said, here. Theory-testing experiments on social dilemma games, for example, do not fit the straightjacket of the precepts. When the assumptions of rational choice theory are themselves under investigation, the aim is to figure out whether individual preferences (and/or beliefs) have

[24] See for example the debate sparked recently by Hertwig and Ortmann (2001). On behalf of monetary incentives, it must be said that economic experiments have become a real business in some universities, used by students to top up their grants (an attractive alternative to a part-time job at MacDonald's, in other words).

the structure postulated by the standard models. The precepts lose much of their appeal in such a context, because clearly there is little point in trying to induce the behaviour one is supposed to be testing in the first place.

Yet, surprisingly, a strict implementation of the precepts is usually advocated in these contexts too. As we have seen a substantial portion of experimental subjects playing social dilemma games choose to play cooperatively, contrary to the prediction of standard game theory. The straightforward interpretation of these results is that many human beings (fortunately) do not behave as predicted by the theory. But a considerable number of economists reject this interpretation and argue that the problem must lie with incentives. If they don't conform to game theory predictions, people must be put in line. Experimental economics is then turned into the exploration of the conditions of applicability of an economic model.

The first step consists in arguing that the preference rankings of the subjects who play cooperatively in these experiments might be inadequately represented by the numbers in the classic prisoner's dilemma game matrix (see Table 1 above).[25] According to the orthodox interpretation of game theory, the numbers represent the (ordinal) structure of agents' preferences. The actual numbers do not even matter, as long as the payoffs are ordered 'Down-Left' > 'Up-Left' > 'Down-Right' > 'Up-Right' (from the perspective of the row player; modify accordingly for the column player). Thus, the argument goes, if we observe anomalous behaviour in the experiment, it is likely that the initial conditions postulated in the model were instantiated in the experiment. Subjects were not playing the prisoner's dilemma game, but another game of their choice.

I'm interested here in the general significance of arguments of this kind. For someone working in the testing tradition the standard reaction to the anomalies of cooperation is to conclude that individual agents do not behave as predicted by economic theory. For someone who believes in the control of individual preferences, in contrast, the immediate reaction is to try to make the anomalies disappear by means of a tighter design. The most obvious move is the scaling-up of monetary incentives: surely if one is playing for hundreds of dollars (rather than the relatively low payoffs commonly used in experiments with college students), he/she will have better reasons to behave as a proper *homo oeconomicus*. (Increasing the monetary rewards is an attempt to implement the dominance requirement, in other words.) Other similar devices are the strict enforcement of privacy (in order to neutralise other-regarding preferences), the introduction of training sessions at the beginning of the game (in order to make sure that subjects understand what is in their 'real' interest, what 'ought' to be done rationally, or what 'really' to expect from others), and so on.

[25] There is an interpretation of game theory according to which the preference structure of cooperative subjects is *necessarily* misrepresented by the prisoner's dilemma matrix, because the numbers represent *revealed* preferences or observed choices rather than psychological entities or dispositions. Ken Binmore has been the standard bearer of this view for a while (Binmore 1994), but since his position suffers from several problems and is probably inconsistent (see e.g. Hausman 2000) I shall ignore it here. The weaker and more reasonable position outlined in the text is defended for example by Weibull (2002).

The effect of these moves is mixed, for norms of fairness and reciprocation seem to be rather robust. But the interesting question is why is so much effort invested in preference control? Suppose we *did* manage to achieve control on subjects' preferences. What would be left to test in a trivial game like the one-shot prisoner's dilemma? The rationality hypothesis (that actions follow from preferences and beliefs) is not really in question in simple games like this. There is little to learn, from a theoretical viewpoint, by making sure that the 'right' preferences are instantiated in the experiment.

But imagine you ultimately intend to construct a little machine, a prisoner's dilemma in flesh and blood (and microchips, if the game is played on a PC network). Then these moves make much more sense. Why should one want to construct such a device, though? Social dilemmas epitomise the failure of uncoordinated strategic behaviour – a situation to be redressed rather than replicated. True, and in fact you don't do very much with a social dilemma machine.[26] Much higher stakes are placed on the applications of game theory to market design, especially in the area of auction theory. The example of social dilemmas is interesting because it shows how the official methodological apparatus of experimental economics is pretty incomprehensible from a theory-testing perspective. Experimental economics is successful not because it confirms or refutes neoclassical theory, but because it 'works'. Because you can do things with experiments.

6. Building economic machines

The trajectory of the institution-building tradition in experimental economics intersects with that of the new institutional economics and the theory of mechanism design.[27] Vernon Smith's methodological pronouncements in the late seventies draw explicit links with this theoretical literature. The main idea behind mechanism design theory (or the '$(New)^2$ Welfare Economics', or 'New Institutionalism' – I shall use these labels interchangeably) is to treat institutions as *variables* that affect the allocation of economic goods (see Hurwicz 1972, 1973). Normative (welfare) economics plays a role at the level of defining a set of criteria used to assess market allocations, or in other words the exogenously defined, presumably politically negotiated objectives to be achieved by means of the economic exchange. Then game theory enters the scene: the market institution is represented as a game that rational agents are trying to solve. The 'best' institution is the one that leads the agents to satisfy the welfare criteria 'as if guided by an invisible hand', by setting the right incentives and by giving them enough information to solve the problem they are facing.

The step from this abstract framework to the creation of an experimental branch of mechanism design is short: all you have to do is replace game theoretical agents with real human beings playing for real money, and abstract institutions with concrete systems of rules. The result, as Smith points out in several of his papers, is a dramatic increase in the realisticness of the result:

[26] Most recent research focuses on the (symmetric) problem of 'pushing' free riders towards the Pareto-optimal solution. See for instance Fehr and Gachter (2000), Burlando and Guala (2005).
[27] Nik-Kah (2004) investigates these connections.

> Laboratory microeconomies are real live economic systems, which are certainly richer, behaviourally, than the systems parametrized in our theories. (Smith 1982, pp. 254-255)

The success of a real-life (laboratory) market depends then on the successful matching of the appropriate kind of agents with the appropriate system of rules. The two are highly interdependent, because rules must be interpreted by agents, and the way in which the latter behave depends in part on the institutional rules.

Consider the assumption of rationality, to begin with. The economic engineer cannot just *assume* that market traders are rational selfish maximisers of the kind postulated in most economic models (including mechanism design theory). One has to *make sure* that this is the case, for the presence of a single 'crazy' player may have a devastating impact on the functioning of a market. The mobile phone auctions that have been run in many countries since the early 1990s provide a neat example of the challenges posed by market design. The auctioned goods are licences for frequencies, owned by the government and sold to private telecom companies. The exact value of each licence is unknown, but the general assumption is that potential users (the companies) can make a better estimate of their value than a bureaucrat or politician, because they have better knowledge of the market and the technology. Even the buyers, however, can only estimate – the market is dominated by *uncertainty* about the value of the goods. The value of one licence, moreover, is likely to depend in part on the ownership of other (neighbour) licences, complicating considerably the evaluation of the 'optimal' allocation. A popular design to deal with this kind of complexity and uncertainty is the 'simultaneous continuous ascending auction', where all the bidders can be active on different markets for different licences at the same time. The exact rules of the game can be rather complicated (see Milgrom 2004, Klemperer 2004, Guala 2001), but Ken Binmore, the experimental game theorist who co-designed the UK auction of 2000, simplifies them as follows:

> If a company wants to stay in the bidding it has to either hold the top bid for one of them or overbid a set amount. The price keeps going up and up until there are only five bidders left. […] The advantage of this design is that it allows the bidders to concentrate on what their valuations for the licences are. After each round what a bidder should do is to say what's my current value for each licence because the events of the last round might change their value for the licence. If you see someone withdrawing from the auction that you didn't expect to see withdraw that's valuable information to you and you might want to change your valuation on that basis, but once you know what your valuation is you then simply ask yourself well what minimum bid would I have to make to become top bidder on a licence. Subtract that bid from your valuation for that licence and that will give you your profit on that licence and then you simply bid to maximise your profit on the assumption that that bid will be the winner. (Binmore in Atkinson 2000, p. 22)

Underlying values, in other words, are not given but constantly updated in light of the moves made by other bidders. This transparency and exchange of information is the main advantage of the simultaneous continuous mechanism compared with other market institutions like sealed bid auctions. But then of course if other competitors behave irrationally, they may send misleading signals to the market. Game theory assumes common knowledge of rationality: I am rational, you are rational, and I know that you are rational, you know that I am rational, I know that you know that I'm rational, … and so on. At a more concrete level, the design of a market institution assumes behaviour with certain formal characteristics on the agents' part, but each agent must also be aware of these assumptions, and must be confident that the other agents are willing to and capable of fulfilling the mechanism's requirements. But how do you make sure that this is the case in a real market?

The answer is a neat example of performativity. Game theorists are keen to stress the simplicity of their preferred mechanisms and the small demands they impose on bidders – 'anybody can do that. That does not require any great skill and it's no secret' (Binmore in Atkinson 2000). But, just in case, 'all bidders have got a pet game theorist to give them their advice' (ibid.). Economists design the market *and* advise the companies that will compete in the market itself. The common knowledge problem becomes: I know that you know that I know … that you have a game theorist in your team. [28]

But that's not the end of the story. Economic rationality is not like Newton's laws, that are supposed to be at work everywhere in the universe. It is a fragile property that must be carefully preserved by creating a hospitable environment. It is a *capacity* or a *potentiality*, and the goal of experimental market design is to create the 'right' circumstances for it to be actualised.

> Designs are motivated by a mechanism (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 – a 'process' that mirrors the behavioral features of the mechanism. The theory suggests the existence of 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)

A good market must impose a certain amount of *discipline* in other words (a Foucauldian terminology is very appropriate here), and the precepts of Induced Value Theory help you to do that. The precepts define an artificial situation: it is simply not true that privacy, for example, is in general instantiated in non-laboratory economic situations. Such a requirement however is crucial in the process of applying highly abstract models to concrete cases: it helps to build the experimental counterpart of the theoretical restrictions that make demonstrations from economic models possible.

[28] The line-up of Market Design Inc., for example, a company created in 1995 on the wave of success of the first telecommunication auctions, is impressive: Peter Cramton, Lawrence Ausubel, John MacMillan, Preston McAfee, Paul Milgrom, Robert Wilson, Jeremy Bulow, Eric Maskin, and other among the finest US academic economists are among the principals.

One way to capture the process of market design is to imagine a hierarchical structure: 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 that define the conditions under which such concepts may be deductively demonstrated. But at this stage the description of the causal structure that brings about effects like efficient equilibria 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.[29]

Economists are guided by experimental, practical, as well as theoretical knowledge in designing their experiments so that these conditions are satisfied. But the circumstances in which an economic system maintains its own structural properties may be quite narrow and fragile. Consider for example how difficult it is to control information concerning the identity of bidders (and hence privacy) in a real auction; Richard Cramton, an economist who worked as a consultant for the PageNet team in the first US auctions for telecommunication licences, recalls for example that

> It was common for a bidder that did not need to bid, because it was the current high bidder, to pretend to place a bid, so as to conceal its identity. These pretend bids were not always successful before round 18, because a bidder could not ask for written confirmation of the pretend bid. Almost all bidders asked for written confirmation for their bids. To get a written confirmation, the bid assistant would have to walk across the room in public view. In round 18, the FCC announced, "Beginning with this round, you may go into the bidding booth and request from the bidding assistant a confirmation of your actions regardless of whether you bid, exercise a proactive weaver, or do not submit a bid." Even this met with limited success, since the sheet on which the written confirmation was printed was folded differently depending on whether it was a real bid or a fake bid. (Cramton 1995, p. 21, n.23)

Computerised auctions are used extensively to create 'appropriate' market conditions, precisely because they allow controlling tightly the quality, amount, and flow of information between buyers and sellers. But a computerised auction system obviously can be used only if we are absolutely sure that the institution will accomplish its goals – for it permits no tinkering with the rules and no adjustments like those described by Cramton. 'Black boxing' is appropriate only at an advanced stage of scientific engineering, when most problems and uncertainties have already been solved.[30]

Until then the fragility or sensitivity of a market mechanism to the details of the material (institutional) arrangements is of great concern to the economic engineer, whose

[29] This way of understanding the relation between models, experiments and engineering owes a lot to Nancy Cartwright's work. See in particular Cartwright (forthcoming).
[30] Muniesa (2000) describes the creation of a computerised stock exchange. On black-boxing and the study of markets (especially finance) see MacKenzie (2003).

machines are supposed to work for several years, in different contexts and without the constant supervision of their manufacturer. In order to build a successful auction then one has 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 a realistic view of individual cognitive capacities. 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. The economic engineer must design the market 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 make sure that the structure is 'right' (and experiments are invaluable for that).[31]

7. The philosophy, politics, and economics of market design

Part of the experimental economics and mechanism design revolution consisted in emancipating economics from its obsession with high theory and appreciating the complex relation between abstract and applied work. Paul Klemperer says provocatively that in practice mechanism design requires little more than undergraduate economic theory. The key lesson, in his view, is: 'pay more attention to elementary theory, to the wider context of the auctions, and to political pressures – and pay less attention to sophisticated mathematical theory' (Klemperer 2004, p. 125).

> The really bad mistake in running an auction is just to take an auction design off the shelf, as shown by a comparison of the British and subsequent European 3G auctions. Auction design is a matter of 'horses for courses', *not* one size fits all; each economic environment requires an auction design that is tailored to its special circumstances. (Binmore and Klemperer 2002, p. C94)

Again, this should come as no surprise to social students of science: several local factors determine the success or failure of a scientific application. Some of these factors are cognitive, some are physical or technological (e.g. the reliability of a piece of software), some are political. A market design, in order to be successful, must be attractive to its users, to the government as well as to the private firms who will compete in the newly designed arena. Consider the telecom auctions once again: the nightmare of the governments was to give away the licences for too little or even not sell them at all. The companies' executives in contrast had to justify the money spent in front of their bosses and shareholders; their nightmare was to pay a sum that looked unreasonably or unnecessarily high – by outbidding other firms by too great a margin for example.

These opposite interests affect the design immensely, for politicians, executives, shareholders, and the public in general do not necessarily see a market mechanism in the

[31] For a practitioner's view on so-called 'robustness' requirements, see Schotter (1998) and Klemperer (2004).

same way as an economist would. Consider a sealed-bid auction mechanism where the winners pay the price of their bid; executives are unlikely to love this mechanism, because it may be embarrassing to justify the difference between a successful bid and the second-highest bid – especially if the difference is in the area of hundreds of millions of euros or dollars. An alternative solution is to have a continuous ascending auction where the winner can always monitor the bids of other competitors. But this mechanism is more fragile to collusion or may lead to a collapse in the level of competition if potential buyers drop out too early from the market (if, for example, they are intimidated by a competitor's aggressive bidding at the beginning or even before the auction). A possible solution is to make sure that there are enough *serious* bidders right from the start by imposing high entrance fees that make it very costly to drop out with nothing in hand. But in order to be effective such fees must be very high, indeed quite close to the final price paid for the licences. And this is scary for the government officials, because setting the entrance fee too high could result in nobody participating in the auctions in the first place.

Solutions to all these problems must be negotiated (see Klemperer 2004, Ch. 3-4, for a general discussion); negotiation usually leads to small concessions, sometimes to concessions that seem *politically* small but may be economically high. (A small change like lowering the entrance fees can for instance lead to a loss of a few *billions* for the government.) And not all designs are equally robust to political pressure. Mechanism design has among other things taught that one has to be very careful about what happens outside the economic realm. This is big news in economics – a science that has tried to differentiate itself from the other social sciences most vigorously during the last half century or so.

Remember where it all started from: Walrasian general equilibrium theory does not (and presumably cannot) pay too much attention to the specific characteristics of single markets. However, at the price of some 'heroic' abstraction it delivers an entirely general proof of the invisible hand theorem, one that promises to establish once and for all the superiority of markets with respect to other systems of allocation. Unfortunately nobody has ever seen (or will ever see) a pure Walrasian market at work, and experimentation has demonstrated that 'impurities' matter enormously. Somehow paradoxically the really successful applications of neoclassical theory so far have revealed that markets work in subtly diverse ways, and that a general recipe for market design is a chimera. [32] Adam Smith's invisible hand requires a lot of fine-tuning and tinkering in order to transform individual greed into social benefits, reliably and consistently. But this important lesson simply makes economics look much more like 'proper' science – like physics or biochemistry, where general laws and theories are applied successfully to specific conditions only after a lot of effort and at the price of several adjustments and compromises.

An interesting question for the historians of the future is why this revolution is occurring now. I can imagine a plausible answer along the following lines: general equilibrium

[32] On the recent turn in neoclassical economics away from general equilibrium and towards the details of market mechanisms, see Mirowski (forthcoming).

theory – like much economics of the sixties and seventies – is 'cold war economics' (Mirowski 2002), science devoted first and foremost to winning an ideological game with extremely high stakes. The real limits of applicability of economic theory were too dangerous and tricky an issue to be properly discussed in such a climate. Market design and experimental economics in the building tradition is in contrast genuinely 'third way' economics. The market can do great things for you, if you learn to use it properly; the difficult task is to find out what 'properly' means.

On the philosophical side, experimental economics provides some excellent examples of performativity at work. A claim of this chapter has been that those who understand performativity as a resource (rather than a threat) ultimately win in the social-scientific game. I think this is true of experimental economics, but I suspect it applies as a maxim to several other areas of the social (and perhaps natural) sciences. By way of a conclusion, I would like to summarise some of the ways in which this interpretation of experimental economics relates to some of themes in the new (performativity-based) economic sociology of MacKenzie, Callon, and others.

Let us start from the ultimate ontological question: do the entities described by economic theory exist? As a philosopher, I don't think this is a question I should try to answer.[33] But let me try with a milder suggestion: *if* the entities described by economic theory exist, they are probably not very common. Economic theory seems to be still a long way off from providing an approximately accurate description of most of the economic world. I'm following here those philosophers and sociologists of science – like Bruno Latour, Nancy Cartwright, or John Dupré – who have insisted that science provides an accurate description of at best only *niches* of the real world.[34] Most of these niches, moreover, are artificially created so as to give the theory its 'best shot', so to speak, by eliminating all the disturbances and the imperfections that normally impede its application to 'naturally occurring' circumstances. The story that I have told here and elsewhere about experimental economics can be seen as just an extension of this overall philosophical outlook.

This 'localist' position, interestingly, was originally devised in the context of the natural sciences (biology for Latour and Dupré, physics for Cartwright). This suggests an important distinction to be introduced in the discussion of performativity in the social sciences. Economics helps shaping the economy in at least two different ways. The first one is indirectly by informing institutional design: economists identify the appropriate initial conditions (to use an old-fashioned philosophical concept) to bring about a certain effect or result. The policy-maker then implements the suggestion, for example by redesigning or by creating a new market that fulfils such requirements. This is not a very distinguished form of performativity however: natural science intervention often works in the same way, and performativity theories attempted, at least originally, to distinguish the peculiar nature of social entities from (an idealised version of) natural reality.[35] This is

[33] In Guala (2001) I was much less cautious, however. Callon (1998) also answers a bold 'yes' to this question.
[34] See e.g. Latour (1984), Cartwright (1999), Dupré (2001); see also Guala (2003).
[35] Barnes (1983) and Searle (1995) are typical in this respect.

not to deny that the initial conditions in social science are institutions, rules, informational constraints, etc. that usually need to be created and maintained by means of performative procedures. The point is rather that the relevant science plays only an indirect role in this process. Other institutions or actors do the main job (the SEC for financial markets, the FCC for telecommunications markets etc.) by setting the rules, the incentives and punishments that supposedly create the 'right' conditions for the result to be obtained.

Drawing again on some old-fashioned philosophy of science, notice that to set the 'right' initial conditions does not guarantee the success of a scientific application. One must also bet on the correctness of the relevant theories/models, or on the existence of the laws or causal mechanisms that supposedly connect the initial conditions with their effects. In economics this link is largely constituted by the actions or choices of the individuals in the market. Here is where the second role of economics in performing the economy becomes evident: economics not only identifies the 'right' conditions for the coordination of (given) individual action, but can shape (change) the behaviour of the individuals who will act in the designed environment. Here is where the performativity thesis has more bite, and where the natural vs. social science/world contrast becomes more striking – in the making of the *homo oeconomics.*

This second (more interesting) role of economic theory is perhaps the only one that is worthy of a new technical term - 'performativity'. Economics can shape behaviour because it works in part as a *norm* for the agents in the market, just like the priest's utterance 'you are now man and wife' creates powers and obligations for the individuals involved in a wedding ceremony. This special feature (normativity) distinguishes 'genuine' performativity from similar phenomena, such as 'bandwagon' and 'underdog' effects, that are often conflated with it.

This distinction between 'Type-1' (spurious) and 'Type-2' (genuine) performativity cuts at a different level from MacKenzie's 'generic' and 'Barnesian' performativities (in this volume). The latter refers to those (perhaps relatively rare) cases where a speech act (utterance, theory) brings about or perpetuates the very entities or phenomena it refers to. It denotes, in other words, a particularly tight self-referential loop triggered by the normative character of a speech act. Whether this special Barnesian (or Austinian) case is common in economics or not, is an interesting question but one which may prove to be particularly difficult to answer (it is, after all, a variant of the more general question of the truth of economic theory).[36] That's why I prefer to use performativity in a broader fashion, to include all those cases in which 'economics matters', but does so *in virtue of its normative character*.

This is what distinguishes my Type-1 from Type-2 (genuine) cases of performativity. It also allows to highlight the peculiar ontological role of social science discourse in changing the social world, by generating new entities and relationships. This take on performativity distinguishes sharply the tradition of ontological analysis that from Austin

[36] Mirowski and Nik-Kah (this volume) highlight several dangers lurking behind the use of performative language to describe the accomplishments of contemporary economic theory, not all of which I fear.

leads to Barnes and Searle, to the Actor-network tradition of Callon and his collaborators. Actor-network theorists, I suspect, find the first type of performativity more interesting *precisely* for the opposite reason: because it blurs the distinction between natural and social entities.

In this paper I have tried to give examples of both types of performativity, but one can find other fascinating instances in the existing literature.[37] In my discussion I have also tried to show how these two procedures are conceptually distinct but in practice tightly interdependent (which probably explains why the distinction is often overlooked). I think this interdependency is due to the simultaneous power and weakness of economics as a science. Economics is *powerful* because, unlike physics, it can in principle directly shape the economy (people's behaviour) along the lines suggested above, by virtue of its own authority, with or without the intermediate intervention and support of other institutions (the SEC or the FCC). But it is also *weak*, because it is not capable to do so always and everywhere. Indeed, one of my claims in this chapter is that the second (direct) form of performativity – without doubt the philosophically most interesting one – can rarely take place without the assistance of performative processes of the first (indirect) kind. Both forms thus go hand in hand, and can rarely, if ever, be observed independently from one another.

References

Allais, Maurice. 1953. "The Foundations of a Positive Theory of Choice Involving Risk and a Criticism of the Postulate and Axioms of the American School." In *Expected Utility Hypothesis and the Allais Paradox,* edited by Maurice Allais and Ole Hagen. Dordrecht: Reidel, 1979.

Atkinson, Mark. 2000. "Auctioneer with a Lot to Sell." *The Guardian* 22 April 2000: 22.

Austrian Political Stock Market. Technischen Universitat Wien: http://ebweb.tuwien.ac.at/apsm/index.html

Barnes, S. Barry. 1983. "Social Life as Bootstrapped Induction." *Sociology* 17: 524-545.

Binmore, Ken. 1994. *Game Theory and the Social Contract. Vol. 1: Playing Fair.* Cambridge, Mass: MIT Press.

- and Klemperer, Paul. 2002. "The Biggest Auction Ever: The Sale of the British 3G Telecom Licences." *Economic Journal* 112: C74-C96.

Burlando, Roberto M. and Guala, Francesco. 2005. "Heterogeneous Agents in Public Goods Experiments." *Experimental Economics* 8: 35-54.

[37] For example: the influence of the efficient markets theory in reforming (designing) financial markets vs. the use of the Black-Scholes model directly in calculating the prices of derivatives; see McKenzie (2001, this volume), McKenzie and Millo (2003).

Callon, Michel. 1998. “Introduction: The Embeddedness of Economic Markets in Economics.” In *The Laws of the Market*, edited by Michel Callon. London: Blackwell.
Cartwright, Nancy. 1999. *The Dappled World*. Cambridge: Cambridge University Press.
- Forthcoming. “Abstract and Concrete Knowledge.” In *The Rise and Fall of Historical Political Economy*, edited by Emma Rothschild and G. Stedman Jones.
CNRS. 1953. *Économétrie* (Paris 12-17 mai 1952). Paris: Colloques Internationaux du Centre National de la Recherche Scientifique.
Cox, James C. and Isaac, R. Mark. 1986. “Experimental Economics and Experimental Psychology: Ever the Twain Shall Meet?” In *Economic Psychology: Interactions in Theory and Application*, edited by A.J. MacFadyen e H.W. MacFadyen. New York: North Holland.
Cramton, Peter C. 1995. “Money Out of Thin Air: The Nationwide Narrowband PCS Auction.” *Journal of Economics and Management Strategy* 4: 267-343.
Davis, Douglas D. and Holt, Charles H. 1993. *Experimental Economics*. Princeton: Princeton University Press.
- and Williams, Arlington W. 1991. “The Hayek Hypothesis in Experimental Auctions: Institutional Effects and Market Power.” *Economic Inquiry* 29: 261-274.
Dawes, Robyn M. and Thaler, Richard H. 1988. “Anomalies: Cooperation”, *Journal of Economic Perspectives* 2: 187-197.
de Vroey, Michel. 1998. “Is the Tâtonnement Hypothesis a Good Caricature of Market Forces?” *Journal of Economic Methodology* 5: 201-222.
Dupré, John. 2001. “Economics without Mechanism.” In *The Economic World View*, edited by Uskali Mäki. Cambridge: Cambridge University Press.
Economics Focus. 1999. “News from the Lab.” *The Economist* May 8th: 96.
Election Stock Market. University of British Columbia: http://esm.ubc.ca/
Fehr, Ernst and Gachter, Simon. 2000. “Cooperation and Punishment in Public Goods Experiments.” *American Economic Review* 90: 980-994.
Fouraker, Lawrence E. and Siegel, Sidney. 1963. *Bargaining Behavior*. New York: McGraw-Hill.
Frey, Bruno and Meier, Stephen. 2003. “Are Political Economists Selfish and Indoctrinated? Evidence from a Natural Experiment”, *Economic Inquiry* 41: 448-462.
Friedman, Daniel and Sunder, Shyam. 1994. *Experimental Methods: A Primer for Economists*. Cambridge: Cambridge University Press.
Gilbert, Margaret. 1989. *On Social Facts*. London: Routledge.
Grether, David M., Isaac, R. Mark, and Plott, Charles L. 1989. *The Allocation of Scarce Resources: Experimental Economics and the Problem of Allocating Airport Slots*. San Francisco: Westview Press.
Grether, David M. and Plott, Charles L. 1984. “The Effect of Market Practices in Oligopolistic Markets: An Experimental Examination of the Ethyl Case.” *Economic Inquiry* 22: 479-507.
Grünberg, Emile and Modigliani, Franco. 1954. “The Predictability of Social Events.” *Journal of Political Economy* 62: 465–478.
Guala, Francesco. 1999. “The Problem of External Validity (or ‘Parallelism’) in Experimental Economics.” *Social Science Information* 38: 555-573.

- 2001. “Building Economic Machines: The FCC Auctions.” *Studies in History and Philosophy of Science* 32: 453-477.
- 2003. “Experimental Localism and External Validity.” *Philosophy of Science* 70: 1195-1205.
- 2005. *The Methodology of Experimental Economics*. New York: Cambridge University Press.
Hacking, Ian. 1999. *The Social Construction of What?* Cambridge, Mass: Harvard University Press.
Hausman, Daniel M. 2000. “Revealed Preference, Belief, and Game Theory.” *Economics and Philosophy* 16: 99-115.
Hertwig, R. and Ortmann, Andreas. 2001. “Experimental Practices in Economics: A Methodological Challenge for Psychologists?” *Behavioral and Brain Sciences* 24: 383-451.
Hogarth, Robin M. and Reder, Melvin W. (eds.) 1986. *Rational Choice: The Contrast Between Economics and Psychology*. Chicago: University of Chicago Press.
Hollywood Stock Exchange. http://www.hsx.com
Hong, James T. and Plott, Charles L. 1982. “Rate Filling Policies for Inland Water Transportation: An Experimental Approach.” *Bell Journal of Economics* 13: 1-19.
Hulse, Carl. 2003. “Pentagon Prepares a Futures Market On Terror Attacks.” *New York Times* July 29, Late Edition, Section A: 1.
Hurwicz, Leonid. 1972. “On Informationally Decentralized Systems.” In *Decision and Organization*, edited by R. Radner and C.B. McGuire. Amsterdam: North Holland.
- 1973. “The Design of Mechanisms for Resource Allocation.” *American Economic Review* (Proceedings) 63: 1-30.
Kagel, John H. and Roth, Alvin E. (eds.) 1995. *The Handbook of Experimental Economics*. Princeton: Princeton University Press.
Kahneman, Daniel, Tversky, Amos and Thaler, Richard. 1986. “Fairness and the Assumptions of Economics.“ In *Rational Choice: The Contrast between Economics and Psychology*, edited by Robin M. Hogarth and Melvin W. Reder. Chicago: University of Chicago Press.
Klemperer, Paul. 2002. “How Not to Run Auctions: The European 3G Telecom Auctions.” *European Economic Review* 46
- 2004. *Auctions: Theory and Practice*. Princeton: Princeton University Press.
Krishna, Daya. 1971. “The Self-fulfilling Prophecy and the Nature of Society.” *American Sociological Review* 36: 1104-1107.
Iowa Electronic Markets. University of Iowa, http://www.biz.uiowa.edu/iem/index.html
Latour, Bruno 1984. *Les microbes: guerre et paix.* Paris : Métailié; Engl. transl. *The Pasteurisation of France.* Cambridge, Mass.: Harvard University Press, 1988.
Ledyard, John O. 1995. “Public Goods: A Survey of Experimental Research.” In *The Handbook of Experimental Economics*, edited by John H. Kagel and Alvin E. Roth. Princeton: Princeton University Press.
Lee, Kyu S. 2004. “Rationality, Minds, and Machines in the Laboratory: A Thematic History of Vernon Smith’s Experimental Economics”, PhD dissertation, University of Notre Dame.

Leonard, Robert. 1994. “Laboratory Strife: Higgling as Experimental Science in Economics and Social Psychology.” In *Higgling*, edited by Neil B. De Marchi and Mary S. Morgan. HOPE supplement, Vol. 26. Durham: Duke University Press.
Loewenstein, George. 1999. “Experimental Economics from the Vantage-point of Behavioral Economics.” *Economic Journal* 109: F25-34.
MacKenzie, Donald. 2001. “Physics and Finance: S-Terms and Modern Finance as a Topic for Science Studies.” *Science, Technology, & Human Values* 26: 115-144.
- 2003. “Opening the Black Boxes of Global Finance.” Unpublished paper, University of Edinburgh.
- This volume. “Is Economics Performative? Option Theory and the Construction of Derivatives Markets”.
- and Millo, Yuval. 2003. “Constructing a Market, Performing Theory: The Historical Sociology of a Financial Derivatives Exchange.” *American Journal of Sociology* 109: 107-145.
Marwell, Gerald and Ames, Ruth E. 1981. “Economists Free Ride, Does Anyone Else?” *Journal of Public Economics* 15: 295-310.
Merton, Robert K. 1957. “The Self-fulfilling Prophecy.” in *Social Theory and Social Structure*. London: Free Press.
Milgrom, Paul R. 2004. *Putting Auction Theory to Work*. Cambridge: Cambridge University Press.
- and Weber, Robert J. 1982. “A Theory of Auctions and Competitive Bidding.” *Econometrica* 50: 1089-1122.
Miller, Ross M. 2002. *Paving Wall Street: Experimental Economics and the Quest for the Perfect Market*. New York: John Wiley & Sons.
Mirowski, Philip. (2002) *Machines Dream: Economics Becomes a Cyborg Science*, Cambridge, Cambridge University Press.
- Forthcoming. “Markets Come to Bits”, *Journal of Economic Behavior and Organization*.
- and Nik-Kah, Edward. This volume. “Markets Made Flesh: Callon, Performativity, and a Crisis in Science Studies, Augmented with Consideration of the FCC Auctions”.
Muniesa, Fabien. 2000. “Un robot walrasien: cotation electronique et justesse de la découverte des prix.” *Politix* 52: 121-154.
Nik-Kah, Edward. 2004. “Designs on the Mechanism”, PhD dissertation, University of Notre Dame.
Plott, Charles R. 1981. “Experimental Methods in Political Economy: A Tool for Regulatory Research.” Pp. 117-43 in *Attacking Regulatory Problems*, edited by A.R. Ferguson. Cambridge, Mass.: Ballinger,.
- 1987. “Dimensions of Parallelism: Some Policy Applications of Experimental Methods.” In *Laboratory Experimentation in Economics: Six Points of View*, edited by A.E. Roth. Cambridge: Cambridge University Press.
- 1988. “Research on Pricing in a Gas Transportation Network.” Office of Economic Policy Technical Report 88-2, Federal Energy Regulatory Commission.
- 1994. “Experimental Methods in Economics and Political Science: The Design and Testing of Policy Options.” *Human Dimensions Quarterly* 1: 5-8.

- 2000. “Markets as Information Gathering Tools.” *Southern Economic Journal* 67: 1-15.
- and Porter, D.P. 1996. “Market Architecture and Institutional Testbedding: An Experiment with Space Station Pricing Policies.” *Journal of Economic Behavior and Organization* 31: 237-272.
Rabin, Matthew. 1998. “Psychology and Economics.” *Journal of Economic Literature* 35: 11-46.
- 2002. “A Perspective on Psychology and Economics.” *European Economic Review* 46: 657–85.
Roth, Alvin E. 1991. “Game Theory as a Part of Empirical Economics.” *Economic Journal* 101: 107-114.
- 2002. “The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics.” *Econometrica* 70: 1341-1378.
- and Peranson, Elliott. 1999. “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design.” *American Economic Review* 89: 748-780
Rubinstein, Ariel. 1999. “Experience from a Course in Game Theory: Pre and Post-class Problem Sets as a Didactic Device.” *Games and Economic Behavior* 28: 155-170.
Schotter, Andrew. 1998. “A Practical Person’s Guide to Mechanism Selection: Some Lessons from Experimental Economics.” In *Organization with Incomplete Information*, edited by M. Majumdar. Cambridge: Cambridge University Press.
Searle, John R. 1995. *The Construction of Social Reality*. London: Penguin.
Simon, Herbert. 1957. “Bandwagon and Underdog Effects of Election Predictions.” In *Models of Man*. John Wiley.
Smith, Vernon L. 1962. “An Experimental Study of Competitive Market Behavior.” *Journal of Political Economy* 70: 111-137.
- 1982a. “Microeconomic Systems as an Experimental Science.” *American Economic Review* 72: 923-955.
- 1982b. “Markets as Economizers of Information: Experimental Examination of the ‘Hayek Hypothesis’.” *Economic Inquiry* 20: 165-179.
- 1989. “Theory, Experiment and Economics.” *Journal of Economic Perspectives* 3: 151-169.
- 1991a. *Papers in Experimental Economics*. Cambridge: Cambridge University Press.
- 1991b. “Rational Choice: The Contrast Between Economics and Psychology.” *Journal of Political Economy* 99: 877-897.
- 1992. “Game Theory and Experimental Economics: Beginnings and Early Influences.” In *Towards a History of Game Theory*, edited by E. Roy Weintraub. HOPE Supplement, Vol. 24. Durham: Duke University Press.
- 2002a. “Interviews: Vernon Smith.” In *Experimental Economics: Financial Markets, Auctions, and Decision Making*, edited by Fredrik Anderson and Hakan Holm. Dordrecht: Kluwer.
- 2002b. “Method in Experiment: Rhetoric and Reality.” *Experimental Economics* 5: 91-110.
- McCabe, Kevin A., and Rassenti, Stephen J. 1991. “Lakatos and Experimental Economics.” In *Appraising Economic Theories*, edited by Neil De Marchi and Mark Blaug. Aldershot: Elgar.

- and Williams, Arlington W. 1992. “Experimental Market Economics.” *Scientific American* December: 72-77.
Thrall, Robert M., Coombs, Clyde H. and Davis, Robert.L. (eds.) 1954. *Decision Processes*. New York: Wiley.
Tuomela, Raimo. 2002. *The Philosophy of Social Practices*. Cambridge: Cambridge University Press.
Vickrey, William. 1961. “Counterspeculation, Auctions, and Competitive Sealed Tenders.” *Journal of Finance* 16: 8-37.
von Neumann, John and Morgenstern, Oskar. 1944. *The Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
Weibull, Jorgen W. 2002. “Testing Game Theory.” SSE Discussion Paper 382, Stockholm School of Economics.
Wilde, Luis L. 1981. “On the Use of Laboratory Experiments in Economics.” In *Philosophy in Economics*, edited by J.C. Pitt. Dordrecht: Reidel.
Wilson, Robert B. 1977. “A Bidding Model of Perfect Competition.” *Review of Economic Studies* 44: 511-518.