Part I

What Is It? An Introduction to Experimental Economics

1 The Emergence of Experiments in Economics

> There is a property common to almost all the moral sciences, and by which they are distinguished from many of the physical; that is, that it is seldom in our power to make experiments in them.
>
> Mill (1836), cited in Guala (2005, p. 2).

This statement by John Stuart Mill, or similar remarks, introduces virtually all texts on the methodology of experiments in economics. At the time, and for a long time after that, controlled experiments in the social sciences, and especially in economics, were considered impossible to conduct; it appeared that experiments were reserved to the natural sciences, and that the testing of social and human behaviour in the framework of a controlled experiment would prove completely unworkable. Nowadays, experiments are a widely accepted means of generating knowledge in economics. Among many examples, it is shown by the fact that experimental or behavioural economics is part of the graduate programme of most universities, there are many books, handbooks and textbooks focusing on the field, and even a well-recognised academic journal ('*Experimental Economics*') is specialised on research using this method.

Before moving on to a detailed discussion of why and how laboratory experiments are performed in economics, we will explore this intriguing trend. What happened between the time experimental economics first came into existence and when it finally became an established member of the community? We will start by highlighting the progress of experimental methods in economics, from an area that was thought impracticable, meaningless or uninteresting, to an accepted and widely used process in economic research. In describing the reasons why there was such a sudden change of interest in and attitude towards experiments, we will examine some of the very first examples of experiments in economics. These examples are interesting not only from a historical point of view, but also because they underscore the main reasons for the change and how experimental economics has grown since – both in terms of the research questions that are addressed and in the type of answers it provides. These will be followed by three more recent examples which illustrate what the research programme has become today – a unified and also very diverse area of study.

The most obvious and powerful unifying factor of all works using laboratory experiments is, in fact, the methodology applied: a controlled environment allowing use of the observed behaviour of human beings to produce knowledge about economics. As the last section will show, a thorough study and presentation of this methodology requires

a wide-ranging knowledge of economic theory as a whole, and its relation to different application fields, analytical tools and approaches. It will soon become clear that no single textbook can possibly cover all these aspects: this chapter will offer a road map of everything this book is unable to cover, or can only cover in part. Perhaps more importantly, this chapter will try to convince you that in order to fully understand the rationale, contribution and practical lessons of the results generated by experiments in economics, the first step is to be aware of the choices of methodology and the reasoning behind them: this is what this book is all about.

1.1 The End of a Long-Standing Regretful Impossibility

Even if experiments in economics were considered impossible for a long time, they were nonetheless the object of considerable wishful thinking. If experiments could be implemented, they could be designed and put in place in order to provide empirical evidence and serve as a basis to enhance theory. This is implicitly acknowledged in a celebrated remark made by Friedman, 'We can seldom test particular predictions in the social sciences by experiments explicitly designed to eliminate what are judged to be the most important disturbing influences' (Friedman, 1953, p. 10). Experiments in the social science are seen as a very attractive, though impossible, way of testing theories. If feasible, experiments would allow researchers to neutralise all forces driving behaviour that are outside the scope of the theory. In that case, experiments would help elicit the empirical content of theory, and therefore identify the main driving forces of behaviour. This opinion was shared by many eminent economists long after 1953. In their groundbreaking principles textbook, Samuelson and Nordhaus noted that 'economists cannot perform the controlled experiments of chemists or biologists because they cannot easily control other important factors' (Samuelson and Nordhaus, 1985, p. 8). All of the remarks cited above show quite clearly how recent the appearance of experimental economics as a *bona fide* field of study is and also underline how desirable experiments are for research. Fortunately, the long-standing and powerful belief in the impossibility of experiments in the social sciences, however regretful, is now a thing of the past.

As a matter of fact, in a later edition of their textbook (which appeared less than ten years later) Samuelson and Nordhaus had already adopted a new and different mindset: 'Experimental economics is an exciting new development' (Samuelson and Nordhaus, 1992, p. 5). Between these two editions, economists had managed to set up experiments similar to the ones conducted in the natural sciences. But, even more importantly, the results generated by these experiments began to be considered by an increasing number of specialists to be sound empirical evidence.

From then on, the pace and scope of the changes taking place increased so rapidly that today the situation stands in sharp contrast with the earlier views expressed above. This phenomenon is illustrated, for instance, by the rise in the rate of academic publications related to experimental economics over the years. Figure 1.1 shows the results of a survey carried out by Noussair (2011) concerning the percentage of articles including experiments that have appeared in major academic economic journals. The survey

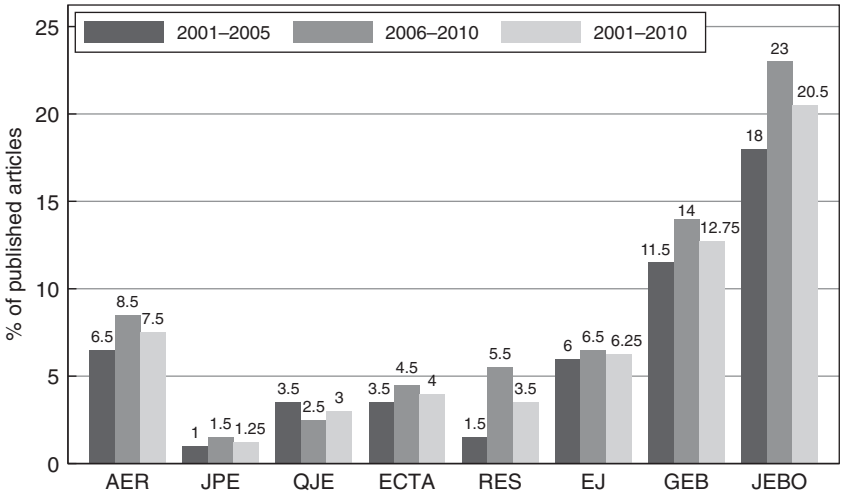

Figure 1.1 Trends in academic publishing in experimental economics

Note. Percentages of experimental articles from those appearing in the journals: *American Economic Review* (*AER*), *Journal of Political Economy* (*JPE*), *Quarterly Journal of Economics* (*QJE*), *Econometrica* (*Ecta*), *Review of Economic Studies* (*RES*), *Economic Journal* (*EJ*), *Games and Economic Behaviour* (*GEB*), *Journal of Economics, Behavior and Organization* (*JEBO*).
Source: Noussair (2011, p. 8).

covers the top five journals (*AER*, *JPE*, *QJE*, *ECTA*, *RES*) which experts acknowledge as the leading supports in the field; three other journals were added to the list: *EJ*, *GEB* and *JEBO*. These are more specialised and/or lower-ranked journals, but which are, nonetheless, highly influential and open to experimental works. The chart shows the change in the rates from 2001–2005 to 2006–2010. The first ten years of the new millennium saw a slight increase in the percentage of articles in the sample. More importantly, the share of experimental papers is very significant in most of these leading journals: from 2% to 7% in the top five journals, and from 5% to 20% in the more specialised ones. This a clear indication of the growing acceptance and recognition of this type of work by the academic community.

The four experimental economists who have been awarded the Nobel Prize in Economics in the first decades of the new millennium, who we will come across a number of times in this book, are another example of this recognition. In 2002, Vernon L. Smith and Daniel F. Kahneman were the joint recipients of the Nobel Prize in Economics. Smith was thus acknowledged as one of the founders of experimental economics and as someone who contributed to establishing it as a conclusive method. The main justification for the award was the introduction of the methodology per se (they received the prize 'for having established laboratory experiments as a tool in empirical economic analysis, especially in the study of alternative market mechanisms'). In terms of contributions, the field is seen as interdisciplinary in nature, with Kahneman receiving the prize 'for having integrated insights from psychological research into economic science, especially concerning human judgement and decision-making under uncertainty'. Ten years later, another renowned experimentalist, Alvin Roth, was also granted the

Nobel Prize. But this time, the co-winner was Lloyd Shapley, a pure theorist. Together they were recognised 'for the theory of stable allocations and the *practice* of market design'. It goes without saying that the Smith and Kahneman contributions are of major importance to the discipline, and that these three Nobel Prizes in themselves are convincing proof that experiments have been widely accepted as part of the field. But there is an interesting change in nature between the two prizes: while the first Nobel Prize was awarded for the methodological advance itself, the acknowledgement of Roth's contribution was based on actual laboratory results using the toolbox of experimental economics and applied to research issues that are at the core of economic theory. This is further evidence of the wide acceptance of experimental economics by the academic community. Last, Richard Thaler was awarded in 2017 for having incorporated 'psychologically realistic assumptions into analyses of economic decision-making'. Richard Thaler showed how experimental methods are particularly meaningful for uncovering deep psychological phenomena such as mental processes, self-control behaviour and social preferences. The award also underlines his contribution to public policies based on nudges (see Chapter 9). This is further evidence of the wide acceptance of experimental economics by the academic community, with results from the laboratory now being seen as useful in order to better design choice architectures.

In contrast with the quotes that opened this section, in which experiments were regarded with substantial scepticism, there is now substantial evidence that experimental economics has become a well-established and widely accepted empirical method. One may wonder how an entire new field has managed to surface in such a short period of time. As a first step towards a better understanding of how this change came about, we will show in the next section that this, in fact, was not the case at all: experiments in economics have existed for a long time, producing results that are much in line with the works that appear nowadays in leading publications. It appears that the reason for the lack of experiments in economics comes not so much from their practical impossibility, but rather from the main focus of academic research at the time. Since then, a change in focus occurred towards questions that are closer and closer to the kind of issue that experiments are well suited to investigate.

1.2 Why Such a Change: Two Early Examples

The two examples below are among the best known of the early experiments. They illustrate the state of infancy of experimental economics at the time, although they are now regarded as important and insightful contributions to economic knowledge.

1.2.1 How Do Competitive Markets Work?

In 1948, Harvard Professor Edward Chamberlin organised a game with his students. The aim was to replicate the functioning of a market in perfect competition with rational agents as closely as possible. Students were randomly assigned a card, which made each student either a seller or a buyer. In addition, the card displayed a price for a hypothetical

good to be sold or bought. For students playing as sellers, this price referred to the minimum price at which they were ready to sell. For the buyers, this price indicated the maximum price they were willing to pay to obtain the (hypothetical) good. Afterwards, the students walked freely in the classroom and bargained with their colleagues to either buy or sell the good. Once a deal had been made, the students came to Chamberlin's desk to report the price at which the good had been sold.

In this framework, economic theory predicts outcomes according to the two curves depicted in Figure 1.2, where the supply and demand curves were drawn based on the prices distributed to students – i.e. how many students were willing to buy or sell at each possible price that appears on their card: a 'induced values' design. The game is a textbook example of a market: the demand curve is decreasing in price, whereas the supply curve is increasing. The market equilibrium determines the actual price that should arise from strategic interactions, as well as the resulting quantities exchanged on the market; the unique stable price is the one that clears the market, in such a way that demand meets supply. This point is an equilibrium not only because the two sides happen to be equal, but more importantly because it is the only state of the market in which everyone agrees to stay – there is no possibility of doing better at the individual level by moving out of this situation. For any other price, there is either excess supply or excess demand, in which case either suppliers (sellers) or consumers (buyers) can be in a better situation by moving to another price level. There are thus strong reasons to believe that the equilibrium should result from real interactions in this particular environment.

Surprisingly enough, Chamberlin obtained the results reported in Figure 1.3 based on the actual behaviour of his students. The dashed line depicts the average price at which students traded their goods during the experiment: it is far below the straight line, or the competitive equilibrium price. There was also a huge variation in the actual prices,

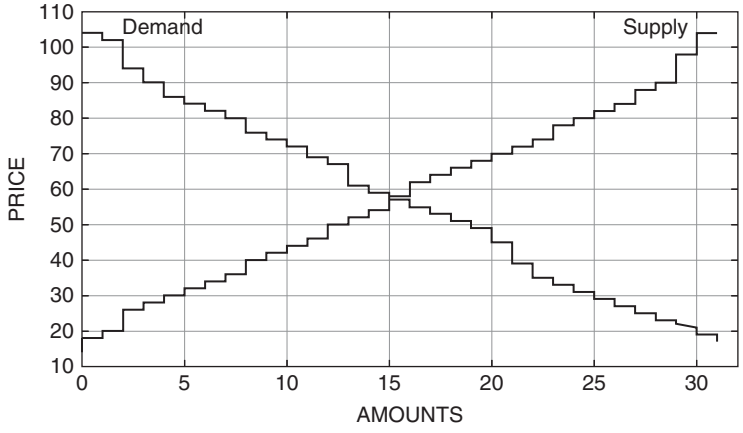

Figure 1.2 Market equilibrium in the Chamberlin (1948) experiment

Note. The figure shows the theoretical equilibrium of the market implemented in the laboratory – at the intersection of the (increasing) supply function and the (decreasing) demand function.
Source: Chamberlin (1948, p. 97, Figure 1).

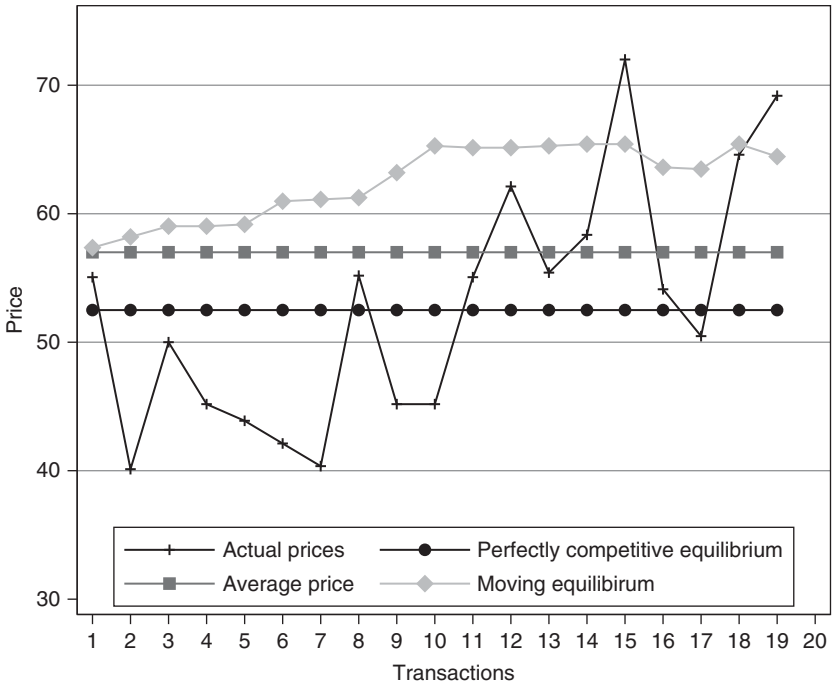

Figure 1.3 Observed behaviour in the Chamberlin (1948) experiment
Note. For each transaction in abscissa, the figure shows the actual price observed in the experiment as well as a recall of the theoretical equilibrium described in Figure 1.2.
Source: Chamberlin (1948, p. 101, Figure 3).

which are represented by the curving line. In addition, the equilibrium volume of trade is higher than what the theory would have predicted. Actual behaviour in this environment thus strongly departs from what economic theory expects, leading Chamberlin to conclude, 'Perhaps it is the assumption of a perfect market which is "strange" in the first place' (and interpret this as a support for his monopolistic competition model). This result is not, however, the end of the experimental story of markets.

Vernon Smith (who, as mentioned above, was subsequently awarded a Nobel Prize) was one of Chamberlin's students and participated in his classroom experiment. Around fifteen years later, in 1962, he decided to replicate Chamberlin's experiment, but with various changes in the environment – aimed at replicating what Smith thought were important actual driving forces of a competitive market. As in Chamberlin's experiment, each student received a card, making him either a buyer or a seller. This card also gave the student a reservation price: the price above which a buyer would not buy, and below which a seller would not sell. The changes implemented as compared to the seminal experiment are as follows. First of all, instead of having bilateral bargaining (or, at most, discussions in small groups) between students, the announcements of offers and demands become public, meaning that buyers and sellers could call out their offers in the room so that everybody could hear. This is aimed to make the information on prices public, so as to mimic what is achieved by an auctioneer receiving and distributing all

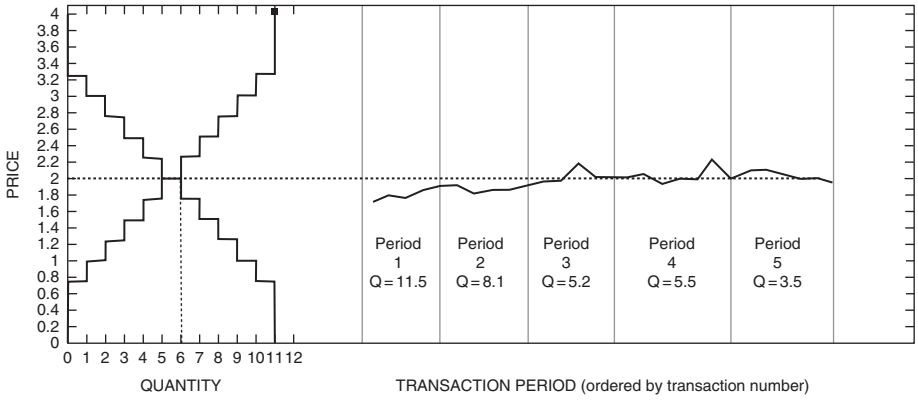

Figure 1.4 Predicted and observed behaviour in the Smith (1962) replication

Note. The left-hand side shows the theoretical market equilibrium – at the intersection of the (increasing) supply function and the (decreasing) demand function. The right-hand side shows the price and number of transactions in each market period.

Source: Smith (1962, p. 113, Figure 1).

offers. Second, the market experiment was repeated over several periods, and allowed the students to better understand the functioning of the market, hence getting closer to market behaviour of professional market traders.

Figure 1.4 reports the observed behaviour and theoretical predictions of the Smith experiment. The theoretical market plotted on the left-hand side shares the same features as the one implemented by Chamberlin. The curve on the right-hand side shows the prices at which market clears for five market periods. The contrast with the previous results is drastic: the observed prices smoothly converge towards the equilibrium price, and the number of transactions (reported on the bottom part of the graph) converges to the equilibrium quantity equal to 6.

Beyond the seminal insights about how the market works, these series of experiments help to describe the methodological issues behind experimental results. Both experiments aim to replicate competitive markets, but with different implementation choices. The best environment to describe markets is a matter of judgement, and the theoretical conclusion drawn will be entirely different whether one or the other experiment is believed to best capture the important features of the economic phenomenon. At the same time, the implementation differences between the two experiments also inform about the key features that explain behaviour in a market situation: the extent of information buyers and sellers receive, for instance, seems to be a critical driving force. Beyond rejection/support of the prediction, the experiment thus informs theory by highlighting the salient dimension to be taken into account. Lastly, as the Smith experiment clearly shows, it is not always the case that the theory is necessarily wrong or that experiments are designed expressly to reject the behavioural assumptions behind the theoretical results (as is sometimes taught, mainly by some academics who view experiment results with scepticism): in this case, experiments serve more to identify the circumstances under which these assumptions are actually accurate.

Table 1.1 The choice sequence of the Allais paradox

	Option A	Option B
A or B?	100% chance of winning 1 million	10% chance of winning 5 million 89% chance of winning 1 million 1% chance of winning nothing
	Option C	Option D
C or D?	11% chance of winning 1 million 89% chance of winning nothing	10% chance of winning 5 million 90% chance of winning nothing

Note. Each respondent was asked to make both choices in turn.
Source: Allais (1953, implemented in 1952).

1.2.2 Choice Consistency in Risky Decisions

The second example focuses on individual decision-making, rather than on strategic situations. During the annual conference of the American Economic Society held in New York City in 1953, Maurice Allais presented the economics professors attending the conference – especially those specialised in game theory and decision theory – with two binary choices. Respondents were shown Table 1.1 and asked to choose either A or B, and then either C or D.

Based on the axiomatic framework of decision theory, the first choice and the second choices are strongly related – although the choice between the two options per se is a matter of preferences that nobody can predict. To understand the link between the two decisions, let us first put aside the 89% probability of winning one million – in situations A and B – or nothing – in situations C and D. Apart from this 89% probability, both situations A and C have the same probability (11%) of winning one million. Similarly, situations B and D offer the same expected outcome: nothing with a probability equal to 1%, and five million with a probability of 10%. As a result, still disregarding this 89% probability, an individual who prefers A over B (B over A) should also prefer C over D (D over C). You can note that the outcome that results from the 89% probability is exactly the same for A and B on the one hand, and C and D on the other. Consequently, it only comes down to the addition of an identical outcome for each pair of situations: one million for A and B, nothing for C and D. It sounds reasonable to assume that this should not affect the preference ordering of consistent decision-makers.[1] Because of this very clever feature in the way situations are built, elicited choices provide a test of consistency: depending on individuals' unknown preferences, either A and C, or B and D, should be picked together; no other combination can be rationalised with classical decision theory. Using these choice situations, Allais was successful at tricking the economists at the conference. As he expected, 45% of the leading theorists (including Savage, one of the leading researchers in the field) to whom Allais submitted the choice

[1] This property of preferences is named the "independence axiom" in decision theory, which implies that if there are two different gambles and one is preferred to the other, then mixing them with another identical gamble should not alter the order of the preferences. This axiom is the one violated by the results of this experiment, which is now known as the common consequence or Allais paradox.