Experimental Economics in the Design of Environmental Markets

Andrew Reeson
CSIRO Sustainable Ecosystems
GPO Box 284, Canberra
Ph: 02 6242 1640
andrew.reeson@csiro.au

This is a working paper only. Please check with the author before quoting.
Please feel free to contact me with any comments, suggestions, criticisms, gifts etc.

Experimental economics is a relatively new and exciting economic discipline. It applies an empirical, scientific approach to the study of economics. It is useful from a theoretical perspective for helping to understand how and why markets work. Much economic theory involves applying complex mathematical techniques to model the behaviour of perfectly informed, completely rational economic agents. Unfortunately for much of this theory, people are seldom fully informed or completely rational, and they are notoriously bad at solving complex equations. Experimental economics puts people squarely at the centre of economics. It incorporates the psychological and social factors that can influence economic decision making.

Experimental economics also has great potential from an applied point of view, enabling economic policies and institutions to be tested under laboratory conditions before being inflicted on an unsuspecting world. Experiments can be applied by policy makers to test alternative institutions under controlled conditions, potentially revealing any flaws or unanticipated consequences of a proposed market design. This process is analogous to wind-tunnel testing used by engineers – the best available theory can be used to design a structure, but there remains some uncertainty as to exactly how it will perform. The wind tunnel process allows potential problems to be identified and the design to be fine tuned.

Experimental economics provides a safe and cost-effective means of testing a proposed market or other economic policy. Potential nasty surprises can be discovered in the confines of the laboratory, avoiding the social and economic consequences of a failed policy. Experiments can also be used to select among alternative institutional designs. There are many different ways in which a market can be set up, and apparently minor details often have a significant impact on performance. The experimental approach can be used to directly compare alternatives, facilitating informed policy choices.

Market-based instruments (MBIs) provide a flexible and cost-effective means of addressing environmental problems. They are increasingly being applied in Australia and overseas as an environmental policy tool. The CSIRO markets for ecosystem services research project has worked on the development of these markets for a range of issues including salinity, water quality and biodiversity conservation.[1] Often this involves creating a novel market institution where none has existed before. In these circumstances, experimental economics techniques is particularly valuable to examine how people are likely to respond to such an institution, and so inform the design of environmental markets.

[1] See www.ecosystemservicesproject.org

What is experimental economics?

Experimental economics involves real people facing real incentives, and therefore making real decisions, albeit within a simplified environment. The application of experiments to economics was pioneered by Vernon Smith, for which he was awarded the 2002 Nobel Prize. Smith's first experiment was a simple laboratory auction.[2] Carrying out an economic experiment requires creating a controlled environment. Smith did this by assigning his 22 participants into an equal number of buyers and sellers; a seller had a single 'unit' to sell, a buyer could purchase a single unit. Each person was given a card indicating the value of a unit to them. This creates supply and demand. For instance, if a unit is worth $4 to a seller, that person would be able to make a profit from selling for anything more than $4. Similarly, if a unit is worth $5 to a buyer, that person will profit from buying a unit for less than $5. The participants had different values, creating supply and demand functions.

The law of supply and demand states that, in a competitive market, equilibrium is reached at the point where supply and demand functions meet. At this equilibrium point, buyers and sellers have the same marginal cost (which becomes the market price), and efficiency is maximised. Smith's experiment was designed to test this law. In order to do this, he had to specify the market institution, that is how buyers and sellers would interact. He used a double auction – an institution in which buyers continually make bids to buy while sellers make offers to sell. Bids and offers were shouted out by participants, and written on a board in full view. A trade occurs when a buyer accepts an offer to sell, or a seller accepts a bid to buy; the trade occurs at the agreed price. Outstanding bids and offers can be cancelled by the originator – if it appears that no one will accept your offer, you can drop your price. The market ran for five minutes, after which any un-traded units were worthless. In each experiment the market was repeated several times.

As this was a controlled experimental scenario, the supply and demand functions were exactly known, and the number of trades and market price predicted by the law of supply and demand could be calculated. Smith's hypothesis[3] was that the market would not reach equilibrium as each trader only had a single piece of information – not nearly enough to work out the market price. However the experiments showed that the law held, and that an efficient market outcome was reached even with a relatively small number of traders with limited information. This early experiment therefore demonstrated that the double auction institution could deliver efficient outcomes where traders hold only limited private information. Many subsequent experiments have confirmed that the double auction performs well even under 'difficult' conditions.

Comparing institutions

Smith's experiment showed that a market institution, in this case a continuous double auction, could efficiently allocate resources among participants with limited information. It also showed how this happened, with prices gradually converging towards the equilibrium, and that the equilibrium does respond to changes in demand and supply. In terms of MBI design, Smith's approach shows how proposed market institutions can be tested in the laboratory, to ensure that they work as theory predicts. However it is possible to go further, applying experimental techniques to directly compare alternative institutions. There are numerous different ways in which a market can be structured. Experiments allow the performance of alternative structures to be compared under controlled conditions, so market design can be based on sound empirical data.

---

[2] Smith 1962

[3] as related by Miller 2002

The use of experiments to compare alternative institutions is exemplified by an early study by Charles Plott and Vernon Smith.[4] They compared an oral bid auction, in which buyers shout out bids and can change their bid if it is not accepted, with a posted bid auction in which buyers make a single bid which cannot be changed. In order to compare the two institutions they had to keep all other factors constant in their experiments. Buyers and sellers were given the same values, creating identical supply and demand functions in each experiment, and instructions which differed only in the details of the market organisation. They found that overall efficiency was high for both institutions, but consistently higher (nearly 100%) in the oral bid auction than in the posted bid market. These results suggest that, all things being equal, an open auction institution should be preferred over a posted offer institution.[5]

Experimental methods

These early examples show that control is the key to experimentation. An ideal experiment will have a single variable, while all other factors are held constant. This allows the effects of the variable of interest to be disentangled from everything else. For instance in their comparison of institutions, Plott and Smith used the same number of participants, the same payoff structures, etc. Therefore any differences in results can be related to the different institutions being tested.

Of course not all variables can be completely controlled. The most obvious of these is the human factor – the whole basis of experimental economics is that it uses real people, and since no two individuals are exactly the same, there will always be differences between experiments. One way around this might be to use the same individuals. For instance, Plott and Smith might have run an experimental session using the posted bid institution, and then used the same group of people in an oral auction. However they would still not be identical – people's behaviour in the second institution is likely to be influenced by their experience in the first. Therefore one could not be sure whether any observed differences were due to different institutions or different levels of experience.

Economic experiments ideally use different people in each session, to avoid the confounding effects of learning and experience. Each group is considered as a random sample of the population. Each group will be different, and may produce different results in the experiment (different individuals may respond differently to an institution; even the same individual can vary depending on mood, interactions with others, etc). To establish that observed differences are down to the institutions being tested, rather than just random differences between groups of people, requires replication – each experimental institution must be run several times, with a different group each time. Statistical analysis is then required to determine whether there are significant differences between the institutions, beyond what would be expected through chance alone.

Laboratory experiments achieve a high degree of control by creating a highly simplified decision making environment. For instance in Smith's early market experiments, participants were told only that they were trading 'units', each was told whether they were a buyer or a seller, and could only trade a single unit in each period. Therefore participants' actions are controlled to a considerable level. All they can choose is if, and at what price, to trade. This was all that was necessary to test Smith's hypothesis, and avoided the complications that might occur if people tried doing other things. However this control comes at the expense of realism. In most real markets, people's actions are far less constrained, and so behaviour may differ from that observed in the laboratory.

[4] Plott & Smith 1978

[5] However all things are usually not equal. There are also differences in transaction costs between these two institutions – see 'institutions matter' below for a discussion.

Experimental design typically involves a trade-off between realism and control. The choice depends on the objective of the experiment. Nowadays most experiments involve a group of people sitting at computer terminals. The use of computers allows a greater degree of control, as the instructions and interface can be exactly replicated in different sessions. They also facilitate anonymity, which may be necessary to observe 'real' behaviour, and the use of software allows complex, real-time trades and interactions. At the other extreme, there is growing interest in field experiments, in which behaviour is measured under far more realistic settings, but usually at the expense of some degree of control.

A laboratory experiment can be made more realistic by providing a greater degree of context. Smith's experiment, like most subsequent academic experiments, was largely de-contextualised. Participants were simply traders buying and selling 'units'. Some studies choose to provide a greater degree of context. For instance, participants could be told that they are farmers buying and selling salinity credits, or taking part in a conservation tender. This can provide greater realism, but it comes at the expense of control. People may express other motivations, such as environmental concerns, just as some may do in the real world. However it makes the experiment more complex, providing more opportunities for confusion and strategic bias (see below).

A key feature of economic experiments is that participants are paid real money based on the outcomes of their decisions. This means that, for them, the values they are assigned become real, and any transactions have a real value. This enables economists to observe real economic behaviour, rather than simply collecting hypothetical survey responses. Typically participants are paid a 'show-up fee' plus or minus the profit they make during the experiment. In our laboratory, participants get a $10 show-up fee, and usually leave with $20-$40. This does mean that running experiments can become costly, but it ensures that quality data are produced. And lessons about policy and market design are vastly cheaper when learned in the laboratory than through trial and error in the real world.

The participants in Smith's groundbreaking experiment were students from his economics class at Purdue University. The use of students in experiments has become standard practice, mainly for practical reasons. Most experimental economics research to date has been carried out within universities, so student populations have been the natural recruiting ground. Students are also more likely to be available during the day, and motivated by smaller amounts of money. However there is no reason for experiments to be limited to student participants, and some studies have deliberately targeted other groups such as CEOs[6] and farmers.[7] A key assumption in most experiments is that participants are representative of the broader population – can experiments with undergraduate students really tell us anything about how, for instance, landholders will respond to a particular institution? No two people are exactly alike, so within any group of people there will be variation. People may have different values, different psychological biases, different ways and abilities of learning, different experience and different social norms. If the group of people who take part in an experiment are 'typical' of the greater population (or the population of interest for a particular study, such as landholders in a certain catchment), then it may be possible to draw conclusions about the behaviour of that population from the behaviour of a few dozen individuals in the laboratory. However this could be more problematic if the participants are systematically different from the broader population.

There are three solutions to this problem. The first is if students are truly representative of the population of interest. There is some support for this from empirical studies. Secondly, even if the experimental participants are different from the population, many experimental findings will still hold. Experiments are typically intended to test whether a particular institution works

[6] Fehr & List 2004
[7] Ward 2005

under certain conditions, or compare the performance of alternative institutions. They are not intended to predict exactly what will happen if such an institution were introduced in the real world (as they are typically carried out in a highly simplified and controlled environment). Therefore if one institution outperforms another in a laboratory setting, designed to reflect the key aspects of the environment in question, then there is strong reason to believe that that institution would perform better in the real world. The third solution is simply to use 'real people' rather than students as experimental participants. It is likely that larger cash payments will be required to motivate people – if only trivial money is at stake, only trivial results can be obtained.

Carrying out experiments involving people requires a consideration of ethical issues. Informed consent must be obtained, which requires explaining the nature and purpose of the research. Unlike many well publicised psychological experiments, experimental economics as a discipline insists that there should be no deception of participants. The supply of volunteers is likely to run low if people feel they have been tricked or duped. Data collected during experiments is always confidential, and usually anonymous. Unlike most research, taking part in economic experiments has positive benefits for the participants as they have the opportunity to earn quite substantial amounts of money relatively easily!

Another issue of using people in experiments is the possibility of strategic bias. When Galileo was dropping his balls from the Leaning Tower of Pisa, it is safe to say that the balls had no interest in the outcome of the experiment. The same cannot always be said of people. Participants may attempt to bias the results of an experiment, perhaps in the hope of affecting subsequent policy choices. This could be particularly problematic if interested stakeholders are used as experimental participants. A further advantage of student participants is reducing strategic bias. However there is also a tendency for people to try and give the answer that they believe the researchers are looking for, which may be enhanced by using students in experiments run by university academics. Thorough experimental design, including anonymity for participants, can minimise this bias.

Institutions matter!

Vernon Smith's early experiments showed that markets work. Resources can be allocated efficiently through the market mechanism even with incompletely informed traders. It has often been assumed that decision makers must have complete information in order for a system to reach equilibrium[8]. In complex and dynamic economic systems it is most unlikely that any decision maker can know everything. However, provided people know their own values, and act reasonably rationally, the market extracts the relevant information and efficient outcomes can be reached.

A market institution is the set of rules which specifies how the market is implemented, and how agents interact with each other - particularly the nature and timing of messages between agents. In practical terms, it has become clear from experiments that details of the market institutions really matter in terms of market performance. Small changes in market institution can make for big changes in market outcomes. Governments around the world have recently begun using auctions to allocate goods such as licences for use of communication bandwidth, with the most notable example being 'third generation' mobile phone spectra. Each country adopted slightly different rules for auctioning 3G spectra, which resulted in massive variation in the amount of revenue raised[9].

[8] Smith 2003
[9] Klemperer 2004

Some kinds of policy problems suit some kinds of institution better than others. As described above, Plott and Smith shoed that posted offer markets converge more slowly and erratically, and are less efficient than, continuous double auctions. However there are also differences in transaction costs which mean that in many circumstances, posted offer is the preferred institution[10] – indeed it is by far the most common institution in daily economic life. Despite being slightly less efficient than auctions, it has much lower transaction costs – imagine if every item in the supermarket were sold by an open auction, going shopping would become a lengthy, complex process, and shops would need skilled traders instead of checkout assistants. In this case the increased transaction costs (your time and energy, and the shop's staff costs) would far outweigh any efficiency gains.[11] For higher value transactions, such as selling a house, or transactions that involve specialist traders such as on the stock market, the efficiency gains from auctions are more likely to outweigh the transaction costs.

Experimental economics and environmental markets

Experimental economics has been successfully applied to a wide range of problems, from designing auctions for mobile phone spectra[12] to allocating resources on space missions.[13] There are a growing number of examples of its application to inform the design of environmental markets.

Testing pollution permit markets

One such example is the design of an auction for tradeable pollution permits in Virginia, USA.[14] The auction, run by the state government, was to allocate permits for nitrous oxide emissions from power plants. The government sought an auction design that would maximise its revenue. A team of economists identified two alternative auction formats which were likely to be acceptable to stakeholders, a traditional sealed bid and an 'English clock' mechanism. The sealed bid requires participants to indicate the price they are willing to pay and the maximum number of permits they are prepared to buy at that price. Permits are sold to the highest bidders at the price they bid. In the English clock, the price increases incrementally, with bidders indicating how many permits they wish to purchase at the current price. The price continues to increase until the number of permits demanded equals the number being sold; all are then sold to the remaining bidders at that price. Under the English clock all successful buyers pay the same price. In the sealed bid format, each successful buyer is likely to pay a different price, which may cause some consternation among those who have paid more than others. This may lead some participants to bid less than they otherwise might, and so reduce overall revenues. On the other hand, the seller is able to get extra revenue from those who paid more than the lowest winning bid. So which works best?

A series of experiments were carried out to compare the alternative mechanisms. Experimental participants took on the roles of buyers in the auction. Each was given values for the units they purchased, as in Smith's original experiments. If they managed to purchase a unit for less than its value, they made a profit which they received in cash at the end of the session. Therefore, just like the energy companies, participants have a strong incentive to buy units as cheaply as possible. The experimental environment needs to reflect the key parameters of the actual market. In this case, energy companies are likely to demand multiple permits, for which they have a declining marginal value. Experimental participants were

[10] Smith Economics in the laboratory

[11] It is interesting, if unsurprising, to note that Plott and Smith found that income tends to be distributed in favour of the side which posts the offer, ie the supermarket wins, you lose.

[12] Klemperer 2004

[13] Ledyard *et al.* 2000

[14] Porter *et al*. 2005. The following description represents a considerable simplification of the actual experiments and auction mechanism.

therefore given multiple values for each unit they purchased. For instance the first unit may be worth $5, the second $4.50, the third $3.50 etc – if three units are purchased for $3 each, the buyer makes a profit of (5-3)+(4.50-3)+(3.50-3)=$4. Different participants have different values for units, just as different energy companies are likely to have different values for permits. Little was known about the value of permits to the various energy companies in the market, making it difficult to create an accurate experimental environment. Therefore each institution was tested in a series of environments with different value sets. If one performed best across all environments it would be a clear winner; if the best institution varied between environments, then it would be necessary to get a better estimate of the value set in the real world.

The experiments were carried out at George Mason University using student participants. While the value sets were designed to reflect the real world, the experiments were context-free, that is no mention was made of energy companies, pollution, etc. This was to avoid participants expressing other motivations such as a dislike of pollution or energy companies, which would be unlikely to be expressed by energy company executives. A computer interface was used to create a standardised and anonymous environment. To prevent collusion (which would be illegal in the real world auction) participants were not permitted to talk to each other. Each institution was tested at least five times to provide a statistically robust sample. The results indicated that the English clock mechanism generated significantly higher revenues. When the pollution permit auction went ahead, it adopted the English clock mechanism, and exceeded its revenue target by nearly 20% ($1.7 million). In this example, there can be no way of knowing the outcome had the experimental findings not been applied. However when the cost of running such experiments is small compared to the revenue generated by the market, there is every chance that it will be a worthwhile investment.

Designing salinity credit trading schemes

Experimental economics has been applied in Australia through a number of the National Action Plan for Salinity and Water Quality MBI Pilot Projects. The dryland salinity recharge credit MBI in the Bet Bet Catchment, Victoria, provides a good example of experimental economics in action.[15] The aim of this pilot was to implement a recharge credit trading system, to provide an efficient way of reducing the salt load entering the Murray River from the Bet Bet catchment. A series of experiments were carried out to address a number of questions, including i) that a cap and trade MBI could work in these circumstances; and ii) inform the detailed design of the market mechanism.

A simulated catchment was developed for the experiments, incorporating real biophysical data. This catchment consisted of 12 farms with different income and recharge characteristics. Possible management decisions for each farm were simplified to five options, representing different crop mixes, each associated with a different level of farm income and recharge. Thus the researchers created a controlled environment which was made realistic by incorporating key aspects of the real catchment, but with a simplified decision making process that could be handled by inexperienced participants (ie students rather than farmers). Experiments were carried out in the laboratories at Griffith University,[16] Brisbane, using student participants. There were 12 participants per session, each taking the role of a landholder on one of the simulated farms. Unlike the previous example, in this study participants were told about the context of the experiment, that they were farmers making crop management decisions with impacts on salinity levels. As discussed above, there are pros and cons of providing context in experiments. In this case the researchers judged that the environmental and social context was important to the decision making process. In this case, decisions are being made by individual landholders (or couples), with direct impacts on their neighbours. Therefore context is likely

[15] Connor *et al*. 2004
[16] Established and run by Dr John Tisdell, see http://www.economicexperiments.com/

to be much more relevant than in the Virginia auction example, in which decisions would be made by executives and professional traders in a large industry.

Each experimental session consisted of 10 periods. Participants made crop management and recharge trading decisions in each period, which represent one year. Operating profits (or losses) were calculated at the end of each period. These were converted to cash amounts – participants see a running total of the money they are earning during the experiment, which they receive in cash at the end. Participants were paid $20-$50 in these experiments. Profits depend on crop decisions and trading outcomes. As in the real world, some farms are inherently more profitable than others. All decisions are made anonymously, so participants cannot see the decisions others are making. The instructions and experimental interface are delivered via computer terminals, creating the standardised environment required to institutions under controlled conditions.

Two alternative mechanisms were tested – open call and closed call markets. In the closed call markets, buyers and sellers submit sealed bids and offers. At the end of the trading period, the market is 'called' – a clearing house assesses all the bids and offers and calculates a single market clearing price, at which supply is equal to demand. All bids to buy at or above this price are successful, as are offers to sell at or below the market clearing price. All trades occur at the market clearing price, so the price is the same for all successful bidders. Since the market is 'closed', only the market clearing price is announced, the actual bids and offers are not revealed. By contrast in the open call institution, all bids and offers are made public. The two institutions differ in the amount of information available to participants. Theoretically they should give the same outcome. But do they?

The results of the experiments showed that the market institutions delivered efficiency gains, compared to the alternative of regulation. In the early periods of the experimental sessions, corresponding to the first few years of any market, gains tended to be somewhat less than predicted by economic modelling. However as participants gained experience in the trading mechanism, overall gains from trade increased. These experiments clearly demonstrated that the closed call market performs better than the open call in these circumstances. With the closed call, the market tended to converge towards the equilibrium, where overall efficiency is maximised. However in the open call, prices were more volatile, and did not converge so well to the equilibrium. The authors suggest that the larger amounts of information in the open call market actually make it harder for people to work out a successful strategy. These experiments therefore show that a cap and trade MBI does have the potential to deliver efficiency gains in reducing recharge, and that a closed call market is preferred to an open call format in this instance.

Conclusion

There is growing interest around Australia in the use of market-based instruments to achieve environmental policy goals. A number of researchers are applying experimental methods to inform the design of these markets. The markets for ecosystem services research team in CSIRO Sustainable Ecosystems are currently applying experiments to assist in the design of auctions for biodiversity conservation and salinity mitigation. There is a role for experimental research both to adapt existing market institutions to fit the context in question, and to design novel 'smart' institutions capable of applying markets to more complex issues.

Experiments can be used to show if, and how, an economic institution works. They can also be used to compare alternative institutions under various conditions. A key role of experimental economics in market design is to inform the choice between alternative institutions. It is particularly useful for testing novel mechanisms, or mechanisms which are

being proposed for novel applications, as relatively little will be known about how they will perform.

Experiments can also be useful in policy design more generally. They can reveal much about how people make decisions, and how they respond to various incentives. The laboratory can provide a safe and relatively inexpensive environment to potential policies and markets. Used wisely, it can bring human behaviour into formal economics and take the guesswork out of institutional design.

References

Connor, J.D., Ward, J., Thomson, D. & Clifton, C. (2004) Design and implementation of a landholder agreement for recharge credit trade in the Upper Bet Bet Creek Catchment, Victoria.

Fehr, E. & List, J.A. (2004) The hidden costs and returns of incentives - trust and trustworthiness among CEOs. *Journal of the European Economic Association* **2**, 743-771.

Klemperer, P. (2004) *Auctions: Theory and Practice*. Princeton University Press.

Ledyard, J., Porter, D. & Wessen, R. (2000) A Market-Based Mechanism for Allocating Space Shuttle Secondary Payload Priority. *Experimental Economics* **2**, 173-195.

Miller, R.M. (2002) *Experimental Economics: How we can build better financial markets*. Wiley.

Plott, C.R. & Smith, V.L. (1978) An experimental examination of two exchange institutions. *Review of Economic Studies* **45**, 133-153.

Porter, D., Rassenti, S., Shobe, W., Smith, V. & Winn, A. (2005) The design, testing and implementation of Virginia's NOx allowance auction.

Reeson, A.F. & Tisdell, J. (2006) When good incentives go bad: An experimental study of institutions, motivations and crowding out. www.ecosystemservicesproject.org/html/publications/markets.html

Smith, V.L. (1962) An experimental study of competitive market behavior. *Journal of Political Economy* **70**, 111-137.

Smith, V.L. (2003) What is experimental economics? www.ices-gmu.org/article.php/368.html

Smith, V.L. Economics in the laboratory. www.ices-gmu.org/article.php/369.html

Ward, J. (2005) Tradeable recharge credits in Coleambally Irrigation Area: Laboratory tests of alternative institutional frameworks. www.ecosystemservicesproject.org/html/publications/markets.html

Further reading

Carpenter, J.P., Harrison, G.W. and List, J.a.E. (2005) Field experiments in economics. Research in experimental economics, volume 10. San Diego, Elsevier.

Davis, D.D. and Holt, C.A. (1993) Experimental economics. Princeton University Press.

Friedman, D. and Cassar, A. (2004) Economics lab: An intensive course in experimental economics. New York: Routledge.

Kagel, J.H. and Roth, A.E.E. (1995) The handbook of experimental economics. Princeton University Press.