

An Introduction to Experimental Economics

Eugenio Proto

Overview

- 1 The introduction: Introducing experimental economics as a methodology, though I will also cover a little about the history of the method and look at the strengths and weaknesses that it brings to the discipline.
- 2 Individual Decision Experiments: Theory and Methods: How to measure individual preferences under risk and ambiguity/uncertainty using laboratory experiments. Explore major behavioural paradoxes: situations, where conventional game theory predictions fail to explain behaviour observed in the laboratory/field (such as Allais Paradox, Ellsberg Paradox, etc.).

Overview, Cont'd

- 3 Individual Decision Experiments: Risk Attitudes and Stochastic Choice: Methods of measuring individual risk attitudes
- 4 Cooperation and Altruism: Games of cooperation like Prisoner Dilemma or Public Good Game
- 5 Bargaining: whether and to what extent players in experimental games behave according to the predictions of the conventional game theory (Dictator Game, Ultimatum Game, coordination games)
- 6 Markets: experimental market and explore whether and how experimental markets reach equilibrium.
- 7 Experimental Herding: Analyze the herding behaviour in the laboratory.

Overview, Cont'd

- 8 Neuro-Economics: Neuro-Biological basis of human behaviour, based on laboratory experiments.
- 9 Field Experiments: You will find out about how field experiments differ from other types of experiments (laboratory experiments, natural experiments and thought experiments).

Reading

There is no course textbook, but the best single resource is: Kagel and Roth (1995): *The Handbook of Experimental Economics*, Princeton University Press.

Lecture notes and course materials can be found on my website for this course.

The ultimate aim of this course is to give you a clearer idea of what experimental economics can do that other methods in economics cannot and how experiments can provide special insights into behaviour.

A Brief History

Some of the oldest experiments within economics are older than you might think!

- Market experiments: Decentralized markets - Chamberlain (1948) induced demand and cost structure; Double auction - Vernon Smith (1962, 1964).
- Game experiments: Prisoners dilemma 1950s - Originally by psychologists and sociologists; Oligopoly games Reinhard Selten (1959).
- Individual choice experiments and choice under uncertainty: Allais paradox (1953).


Experimental Economics in the Discipline

Running experiments is now an established method to explain and/or describe economic activity which bring economics into alignment with many of the natural sciences which rely on experimental methods (e.g. physics and biology).

This is backed up by publications, citations and even a Nobel prize (Vernon Smith, 2002).

For example note that from the last 12 years 11% of the most-cited papers are experimental which is roughly the same number as theoretical papers.

A Quick Look at the AER



When to Use Experiments?

Experiments allow us to do several things but we have to be careful not to use experiments when existing methods may be better (in many senses).

We will look at a few accepted reasons why experiments add to the discipline by providing new ways to do things and new insights.

Let's start with the most obvious rationale and the one that our colleagues in the natural sciences would immediately recognize as extremely important...

Testing Theory

Experiments can test theories under precisely controlled and/or measured conditions that are typically unavailable in field data. The key thing that experimental economics provides is **control** of various forms:

- Institutions (e.g. voting rules, communication, etc.) though not perfect (e.g. social norms).
- Incentives (payoffs) again not perfect (e.g. altruism).
- Measure and checking confounding or unobservable (in the field) variables (e.g. beliefs).
- Randomization (avoids some self-selection problems).

Game Theory

The development of **Game Theory** gave particular impetus to experimental economics in the 1950s, as game theory offered testable theories of economic behavior that depended on the fine structure of both the strategic environment and the preferences of the players.

It is therefore no surprise that many experimental economists are also (game) theorists often seeking to test their own theories or those of their peers.

Empirical Regularities

Experimental data can allow us to search for **regularities**, and exploring and documenting unanticipated regularities has given experimental economics some of its biggest hits (from Allais onwards).

These are often all about violations of the predictions of existing theories which explains the links with behavioural economics.

Challenging theory: do people play Nash strategies, do they really never play dominated strategies, can they optimize decisions, use Bayes rule, calculate risks properly, solve the sorts of problems implicit in economic theory, etc.?

Helping theory: For example, if there are multiple equilibria which is selected?

Informing Theory

As well as helping to refine theoretical ideas experiments can also assist in the formulation of new theories, to explain newly observed regularities, and devising new experiments to help distinguish among such theories.

Examples: behavioural economics, generalized and alternative models of expected utility theory, learning in games, bounded rationality, etc.

We will look at a specific example later (informational herding), but for now consider a famous experiment that started the ball rolling on much of behavioural economics...

The Allais Paradox

the **Allais paradox**.

- Choice 1: A : 4000 with $p = 0.20$ or B : 3000 with $p = 0.25$.
- Choice 2: C : 4000 with $p = 0.80$ or D : 3000 with certainty.
- EUT predicts people choose $A-C$ or $B-D$, but in experiments $A-D$ or $B-C$ are commonplace.

Evaluating Assumptions

Theory often contains **assumptions** (as N increases such and such happens - or when N is infinite firms behave as competitive) - but what does this mean in practice (how big a number is treated as infinite?).

Where theory is questioned there is also the issue of when and why and this relates to so-called “stress-testing”: a theory may fail for a certain set of parameters - will it do better with others? (eg how do contributions to a public good game change as N rises?).

In this way experimental economics may help to be clearer on when theory works and when it does not.

Policy

There are also policy-oriented experiments. Most commonly to do with **market design**.

There are many examples, but think about something as huge as conducting a spectrum auction or as complex as the effects on worker motivation of various compensation schemes.

Would it make sense to just go ahead without first testing the mechanism you have invented? (UK spectrum auctions example: Binmore and Klemperer).

Replicate Reality?

So we now have an idea of what experiments **can** do, but what makes a **good** experiment?

Should an experiment replicate reality? Should an experiment replicate a formal model? Typically no to both!

The real goal is a design that offers the best opportunity to learn something useful and to answer the questions that motivate your research.

An experiment is judged by its **impact** on our understanding of behaviour.

A Good Experiment

So, if an experiment is not just a replication of reality or a model, what really makes an experiment good? The answer depends on what you are testing or exploring, and who you are talking to. A good design:

- Is **simple** compared to reality and even simpler than relevant models (remember that models are themselves a simplification of reality).
- Is designed to test specific hypothesis or set of hypotheses.
- Tests or controls for alternative hypotheses (this is the bit that is often missed!).

Potential alternative hypotheses may again depend whom you talk to, which is why psychology experiments often look so different to economics experiments.

Confounding

A good design controls for the most plausible **alternative** hypotheses that might explain what is being observed:

- Avoid confounding theories that give an equally plausible rationale for behaviour under the experimental design.
- Protect ourselves from fooling ourselves into believing what we want to believe.
- Science done by people who are following up on their intuitions, and (often) investigate hypotheses that they believe to be true.
- The same intuition that causes you to believe the hypothesis might give you a good idea of situations in which the hypothesis will hold.
- But if there are other reasons that those conclusions might hold, you have to make sure that you haven't just created a situation that gives you the results you expect, but not for the reason that you believe.

Alternative Hypotheses

What are the most plausible alternative hypotheses we should be controlling for?

We have already noted that this typically depends on who you are talking to (economists/psychologists?).

This may also depend on recent developments in theory, in the laboratory or in the field. Like any other branch of economics current fashions are important!

Some alternative hypotheses become obsolete with time because they have been rejected or are simply out of fashion.

Testing Alternatives

Design by subtraction: if you think X is the reason for the result, design an experiment in which X is impossible, but all other explanations are possible. This can isolate the effect you want. Example: this is the basis of the most famous of natural experiments - twin studies -in which genetics as a rationale is removed by using genetically identical individuals. In economics an example might be making all incentives the same to see if incentives really change behaviour in different settings.

Design by manipulation: change a parameter that you think should make X a more likely explanation, or make it easier for X to be expressed in the game. Example: making incentives more or less like a tournament to see if different genders respond to different types of incentives (Niederle and Vesterlund, 2007). This tends to be the more common approach in economics.

Treatments

Next up let's look at what we can do to heighten our control.

First off we have different **treatments** to compare with a **control** (e.g. different payment methods, different information provided to subjects). Good practice involves:

- Testing hypothesis by changing one variable at a time.
- Only changing variables which are directly relevant to the hypothesis being tested, otherwise holding the environment fixed.
- Avoiding confounds (don't change more than one thing at a time).

Uncontrolled Factors

We can deal with many uncontrolled factors via **randomization**.

For example, experiments designed to test how subjects' attitudes towards fairness are affected by some treatment variable. Subjects enter the lab with differing attitudes about fairness so a true controlled experiment can't be run:

- By randomly assigning subjects to treatments, we can eliminate subjects differing attitudes as a cause of differences between treatments. This relies on the **law of large numbers**, implying that a large sample may be necessary.
- Or, we can measure variables which you think may affect fairness directly: gender or age for example. This explains why collecting demographic information (via a questionnaire) is so standard in experiments.

Within vs Between

Within vs. between (panel vs. cross section) designs allow some indirect experimental control.

Within-subject design: participants make decisions in all treatments.

Between-subject design: different participants make decisions in each treatment.

Under a within-subject design each subject is its own control. This is great as we need not worry about having different characteristics of participants in each treatment (often easier to get significance). However on the flip side there is the disadvantage of order effects or even fatigue.

More Design Choices

One round versus many rounds? (With implications for learning!).

Pay one (randomized) round or all rounds? (With implications for behaviour?).

Use language that is neutral? (**priming** or **framing** might be problematic and we will come back to this later).

Train participants, or test them before you use them as participants in your experiment? (Again learning is an issue).

The Bottom Line

Good experiment identifies an interesting question or questions (issues that are better addressed through a controlled experiment than through gathering field data).

It should determine a precise set of hypotheses.

The design should involve a simple environment that allows you to test the hypotheses that matters. The more complicated the environment the more likely you are to lose control and be unable to draw inference.

It should deal with confounding alternatives.

Realism

So far it seems like experiments are a fantastic method but there are many objections we need to consider. We will look at a few of the major ones in sequence.

First off is the potential lack of realism (the **external validity** issue). How can we respond?

- Experiments involve real subjects making real money.
- Less relevant for testing theories (again remember that models are themselves a simplification of reality).
- Realism can be added in controlled steps (much as with theory).

Representativeness

Next is the issue of the representativeness of subjects. Again, how can we respond:

- Less important for comparative statics and when we know demographics and can control for them.
- Can be tested and may not be important in some situations (e.g. when a theory is supposed to apply to everyone).
- We can go into the **field** for greater realism and representativeness (we will come back to this later).

Incentives

An old criticism of experimental work was the lack of clear incentives and experiments may even rely only on **intrinsic motivation**. This is less of an issue for experimental economics where the norm is to provide clear performance-related pay (which provides **extrinsic motivation**).

The advantages are clear as subjects are likely to make more effort (if economics is at all correct!) and there is likely more control over incentives if they are clear. It also makes clear that the experimenter wants subjects to follow their own interests (and not the interests of the experimenter or “society”).

But there are also disadvantages: it is expensive and in some sense limits the stakes (may make things seem trivial). It may also crowd-out intrinsic motivation (Gneezy and Rustichini, 2000).

Clarity

Might experiments also involve hard to understand instructions which fail to capture the key features of the problem at hand? Could this explain why subjects failed to act optimally? Savage's famous response to the Allais experiment - supposedly after at first contradicting EUT with his answers he "realised" his error and so would not do so again.

Concrete clear wording or neutral language (fully detailed in any published work to allow replication) seems a good reply to this issue but it is a valid point that needs to be dealt with and poor instructions may lead to an experiments downfall.

Concrete Wording?

- Advantages of concrete wording: can help with understanding of the experiment and can bring the experiment closer to research question.
- Disadvantages of concrete wording: loss of control as you don't know how subjects perceive their role e.g. subjects might "role-play".
- Norm in experimental economics: "Neutral" language describing actions and their payoff consequences.

Priming and Framing

Closely related to the issue of clarity is the issue of **framing** (how you describe something may change behaviour) and **priming** (the order in which you explain things may have an effect).

Again in general the answer is to provide a clear statement of what wording you use and try to avoid priming or use “neutral” frames as much as possible though to some extent it may be impossible to eliminate them entirely. Checking a result under different frames or making the frame as close to what would be seen in reality are also possibilities.

Priming or framing can actually be studied directly during the experiment (and may actually be the subject of the experiment!) in which case you can try to test different frames or different levels of priming to evaluate the impact.

Time-line: Outline

With a better idea of the objections it we are now in a position to think about “good practice”.

The process of starting an experimental project is long and complex. This is an idea of what you would need to think about (roughly in order):

- Formulate a research question.
- Choose design to address the research question: treatment variable(s), within vs. between, required number of independent observations, number of sessions/subjects.
- Prepare an experimental outline.
- Seek funding.
- Ethical approval.

Time-line: Details

So we have an outline and funding plus ethical approval. Next come the detailed instructions and design plus a pilot.

- Write instructions for each treatment and an overall “script”.
- Prepare a questionnaire (useful for controls).
- If computerized, make sure the software is capable.
- Organize money (when/how to pay).
- Recruit for the pilot.
- Run pilot experiment.

Time-line: The Experiment

Finally we can run the experiment and produce a paper.

- Improve the design/instructions based on the pilot experience.
- Recruit subjects.
- Run the experiment.
- Analyze data and write the paper!

Experiments usually yield something (even if it is not what you expected) but can occasionally produce very unclear results so experimental economists (like any other economists) have to be ready to abandon a paper if it provides no clear insights.

Practical matters

Next let's think a little more about some important practical issues.

Chemists used to use metal beakers for mixing ingredients, until they learned that the metal sometimes reacts with the chemicals. Then they switched to glass. Lesson: make sure you limit the things that can react with your treatment. You want to be able to replicate environment!

Provide a precise description of the experimental protocol, so make sure instructions are easy to understand.

No Distractions!

Good practice also involves eliminated distractions and enabling privacy:

- No talking (participants, experimenters) and no distractions.
- If a subject needs help then answer questions privately and quietly.
- In smaller groups reduce the risk that one participant says something out loud and affects the whole group.
- Use appropriate screens if privacy is important (it usually is).

Deception

In psychology experiments deception is commonplace and experimenters may even use stooges. Advantages: can lower costs, easier to study rare situations, easier to design experiment.

However, experimental economics do not use deception. Why? With deception comes a fear that your subjects will no longer believe what you are telling them (loss of control) and may even start to outguess you (changes behaviour). This has an externality for other experiments in the locality.

Close to impossible to publish economic research with even minor deception though it is “allowed” to not reveal the whole truth.

Anonymity

Just as experimental economists don't use deception, we also try to preserve anonymity. But there are many types of anonymity:

- Single blind: subjects guaranteed that no other subjects (or indeed, no one other than the researchers) will be able to ever identify their actions or payoffs.
- Double blind: subjects guaranteed that no one can link their decisions to their identity, even the researchers. Generally this is better than single-blind.
- Abandoning anonymity: if abandoning anonymity is an important part of the design then subjects should know what information about them is to be revealed publicly, and have the option of withdrawing if they do not want information about them revealed publicly.

Ethics

What is ethical depends on the discipline: medical experimenters do things psychologists might worry about and psychologists in turn do things economists might worry about. (Example: trauma healing experiments in medicine).

Deception is a good example: fine in psychology, not in economics. General ethical boards don't like to see deception even in psychology.

Ethical approval normally requires high levels of anonymity and also the ability for subjects to withdraw when they wish (in case they have concerns about privacy).

Matching

For some game theoretic experiments there is also the issue of how to match up subjects:

- Partners: always play with same group.
- Strangers: randomly re-matched before playing each game.
- Perfect strangers: subjects do not play with the same subjects more than once.

Replication

Finally, of huge importance is the ability to replicate experimental work. Replication checks for robustness, experimenter effects, etc. and being able to replicate work is typically a strict requirement for publication.

To need to enable replication also provides strong incentives to get it right.

It also necessitates the complete availability of data, instructions, software and procedures. Notice that it does not conflict with anonymity.

Going into the Field?

More external validity built-in (based in the real world?). Real people vs undergraduates (but students **are** real people!).

The subject pool is spot-on: use market traders to study trading strategies, use politicians to study legislative bargaining, etc.

You may want really large samples (thousands of people) and in that case the laboratory may actually be more expensive (you can go to developing nations to run large-scale field experiments).

You may wish to test if a change would have a sizeable effect when many other things happen as well in a “real-world” environment.

Sometimes the field is just the right place: e.g. examining the impact of a major real-world catastrophe may be hard to motivate in the laboratory.

Staying in the Laboratory?

The laboratory provides more control which is after all a key feature of experimental economics:

- Typically it is easier to get strict instructions followed when experiments are run in the laboratory.
- Students may follow difficult instructions more easily.
- Precise control over the incentives.

Transparent and Replicable?

The laboratory offers more transparency: the subject pool (undergraduates) is well understood. In the field you may worry you use a subject pool prone to some bias, that is then attributed to the experiment.

Laboratory experiments are more replicable: Lab experiments are very easy (and cheap) to replicate. This may make us more comfortable with surprising results.

Ethical issues may be easier to overcome in the laboratory.

Field experiments may simply be infeasible in terms of design, cost or opportunity.

Sometimes it is possible to compromise and either bring the laboratory to the field or field-elements to the laboratory to gain some of the positive aspects of each.

References

Here are the references mentioned in the lectures.

- Gneezy, U., Rustichini, A. (2000). “Pay enough or don’t pay at all.” *Quarterly Journal of Economics*, 115: 791-810.
- Kagel, J. H. and Roth, A. E. (Eds.) (1995). *The Handbook of Experimental Economics*, Princeton University Press.
- Niederle, M., Vesterlund, L. (2007). “Do women shy away from competition? Do men compete too much?” *Quarterly Journal of Economics*, 122, 1067-1101.