How should researchers interested in social and economic policy untangle cause and effect? A new book by **Joshua Angrist** and **Jörn-Steffen Pischke** shows how the five core econometric tools – randomised trials, regression, instrumental variables, regression discontinuity designs and differences-in-differences – accomplish this. These tools lie at the heart of CEP research.

The path from cause to effect: mastering ‘metrics

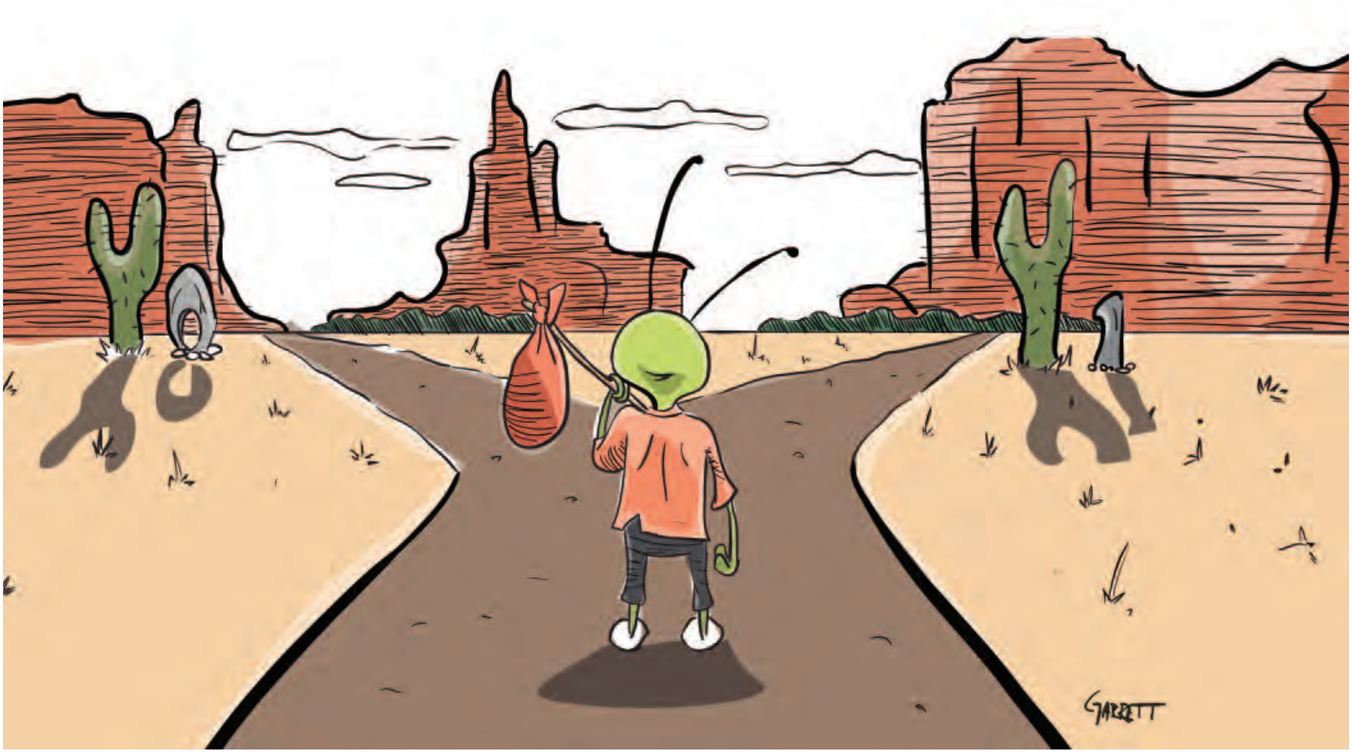

The most interesting and useful economic and social research asks big questions about cause and effect. Does access to free health insurance (as with the UK's NHS) make people healthier? Does going to a school or college with high achieving peers really make the kids who go there smarter? Should abusive domestic partners be referred to social services or simply arrested? Can loose monetary policy save shaky banks in a financial crisis?

Many obstacles litter the path from cause to effect, and the raw data often refuse to reveal the way to causal enlightenment. In a new book written primarily for undergraduate economics students (but also, we hope, for policy-makers and an economically literate citizenry), we explain how masters of the 'metrics trade uncover reliable evidence of causal connections.

We explain by example, with applications and case studies ripped from the headlines, and, in some cases, from our students' lives. We first consider the causal effects of health insurance. Obamacare extended subsidised health insurance coverage to many low-income workers who would otherwise have been uninsured. This is costly but seems justified by a health dividend: a simple comparison of the insured and the uninsured reveals the insured to be much healthier than the uninsured.

Does the relative health of the insured indeed mean that policies like Obamacare improve health? Not necessarily. The case for causality gets weaker when we notice that Americans who have health insurance are richer and more educated than the uninsured. Maybe it's those attributes, and not insurance itself, that are responsible for better health among the insured. Comparisons between the health of the insured and uninsured are not *ceteris paribus* – Latin for 'other things equal'. Rather, such simple comparisons are contaminated by other differences, a problem known to social scientists as 'selection bias'.

Scientists can engineer *ceteris paribus* conditions by running an experiment – called a randomised trial – where they vary only one thing at a time, like giving health insurance to some individuals but not to others. These experiments are much like the clinical trials that doctors have used to evaluate drugs and medical interventions since the middle of the twentieth century.

Although randomised trials are expensive and time-consuming, they have become an increasingly important tool in social science research. The power of an experiment comes from the fact that it separates the variable whose effects we're interested in (say, insurance status) from the selection bias that plagues naïve comparisons of insured and uninsured (the fact that the insured are richer, more educated, etc.).

In this spirit, we explain and interpret results from two remarkable social experiments that randomised access to healthcare coverage in the United States: the RAND Health Insurance Experiment from the 1970s; and the recent Oregon Health Insurance Lottery, which extended state sponsored healthcare coverage to a random subset of low-income applicants. Both experiments reveal that those covered by more generous insurance use more costly healthcare. Yet the extra healthcare consumed by those randomly assigned to the insured group generates few dividends in terms of better health! Insurance helps the insured avoid financial catastrophe when they fall sick – but it doesn't appear to make them healthier.

Experiments like these are expensive and slow to bear their research fruit. We therefore teach our readers to look to the experimental method as a benchmark while also explaining how masters of 'metrics extract causal evidence from the data generated in the course of everyday life. Our book demonstrates the four most important 'metrics tools employed in this effort: regression analysis; instrumental variables; regression discontinuity designs; and differences-in-differences.

Regression analysis attempts to eliminate selection bias by making like-for-like comparisons. Our regression example asks whether there's an earnings payoff to spending upwards of $50,000 a year on private university tuition, as many young Americans do, rather than going to a cheaper state subsidised university (a choice that surely will become relevant for more and more Europeans in due course).

Students who attend relatively selective private universities are likely to have higher earnings for many reasons – they come from richer families, for example. This is the selection bias that plagues a simple comparison of students attending cheap and expensive universities. We use regression to show that the most important sources of selection bias here are the universities to which applicants are admitted. Conditional on where you *could* have gone, where you actually go (U Penn versus Penn State, say) matters little, at least as far as your subsequent wages go.

The most interesting economic and social research asks questions about cause and effect

We next explore the use of instrumental variables, a remarkably flexible and powerful tool that is closely connected with randomised trials, but cheaper and more accessible! In research on questions where the variable of interest can't be manipulated directly, we can instead randomise incentives to choose a particular treatment. In the Oregon Health Insurance Lottery, for example, the lucky winners were only 25 percentage points more likely to receive state sponsored insurance. Instrumental variables readily fix the problems in analysing data from such an experiment.

The real power of the instrumental variables method lies in its ability to harness many useful sources of naturally

> Modern econometrics offers powerful tools for analysing the relationships hidden in large and complex data sets

occurring variation. We use instrumental variables, for example, to ask whether kids who grow up in a larger family get less education as a result. Instrumental variables for family size can be constructed from randomly occurring twin births (sometimes the birth lottery generates a bonus!) and sibling sex composition (mothers of two boys or two girls are substantially more likely to have a third child). Because twinning and sibling sex composition are essentially randomly assigned, they're unrelated to family background and other sources of selection bias.

Our fourth tool – the regression discontinuity design – compares people who are narrowly on opposite sides of a fateful policy cut-off. To illustrate, children who took the entrance exam for a selective school (like the 11 plus in English grammar schools) but just missed being accepted should be a good control group for those who obtained the minimum mark for admission. Applying this to prestigious 'exam schools' in New York and Boston reveals that those who missed the cut-off for these selective schools seem to learn no less than those who just scraped into the exam schools. This is an example of the *ceteris paribus* principle in action: find two groups of people who are distinguished by one key feature – in this case the type of school they attended, with 'other things equal'.

Our last tool – differences-in-differences – compares trajectories over time instead of contrasting differences in levels at a point in time. We apply this to explore the topical question of whether it's worth saving a teetering bank. Walter Bagehot, the editorial father of *The Economist* magazine, famously commented: 'The cardinal maxim is, that any aid to a present bad Bank is the surest mode of preventing the establishment of a future good Bank.' Was he right? This question lay at the heart of macroeconomic policy responses to the financial crisis of 2008.

Our 'diffs-in-diffs' chapter recounts research showing how during the Great Depression, the Atlanta-based district of the US Federal Reserve instituted a policy of lending to troubled banks, while the Fed's St. Louis-based district restricted credit. These districts shared a border that split the state of Mississippi, creating a natural experiment, since other economic and policy conditions across this arbitrary boundary were similar. The diffs-in-diffs analysis reveals that the Atlanta Fed's liquidity injections saved banks and improved its district's economic trajectory, while the St. Louis Fed's district sank more deeply into depression.

Our five econometric tools – which we call the Furious Five, inspired by the Kung Fu theme woven through the book – are central to causal analysis. We reveal their awesome power through interesting and relevant examples. We hope our readers will learn to wield these tools skilfully – first by reading the book, but mostly, as with all sophisticated tools, by a regimen of personal experimentation and practice. Diligent 'metrics apprentices will reap rewards not just in scholarly work but also through a wide range of applications in business and public policy.

Mastering 'Metrics: The Path from Cause to Effect by Joshua Angrist and Jörn-Steffen Pischke is published by Princeton University Press (http://masteringmetrics.com).

Joshua Angrist is the Ford Professor of Economics at MIT. **Jörn-Steffen Pischke** is professor of economics at LSE and a research associate in CEP's programmes on labour markets and wellbeing.