

Natural experiments help answer important questions

*This year's Laureates – **David Card, Joshua Angrist and Guido Imbens** – have shown that natural experiments can be used to answer central questions for society, such as how minimum wages and immigration affect the labour market. They have also clarified exactly which conclusions about cause and effect can be drawn using this research approach. Together, they have revolutionised empirical research in the economic sciences.*

If we are to make good decisions, we must understand the consequences of our choices. This applies to individuals as well as public policy makers: young people who are making educational choices want to know how these affect their future income; politicians considering a range of reforms want to know how these affect employment and income distribution, for example. However, answering broad questions about cause and effect is not easy, because we will never know what would have happened if we had made a different choice.

One way of establishing causality is to use randomised experiments, where researchers allocate individuals to treatment groups by a random draw. This method is used to investigate the efficacy of new medicines, among other things, but is not suitable for investigating many societal issues – for example, we cannot have a randomised experiment determining who gets to attend upper-secondary school and who does not.

Despite these challenges, the Laureates have demonstrated that many of society's big questions can be answered. Their solution is to use natural experiments – situations arising in real life that resemble randomised experiments. These natural experiments may be due to natural random variations, institutional rules or policy changes. In pioneering work from the early 1990s, David Card analysed some central questions in labour economics – such as the effects of a minimum wage, immigration and education – using this approach. The results of these studies challenged conventional wisdom and led to new research, to which Card has continued to make important contributions. Overall, we now have a considerably better understanding of how the labour market operates than we did 30 years ago.


Natural experiments differ from clinical trials in one important way – in a clinical trial, the researcher has complete control over who is offered a treatment and eventually receives it (the treatment group) and who is not offered the treatment and therefore does not receive it (the control group). In a natural experiment, the researcher also has access to data from treatment and control groups but, unlike a clinical trial, the individuals may themselves have chosen whether they want to participate in the intervention being offered. This makes it much more difficult to interpret the results of a natural experiment. In an innovative study from 1994, Joshua Angrist and Guido Imbens showed what conclusions about causation can be drawn from natural experiments in which people cannot be forced to participate in the programme being studied (nor forbidden from doing so). The framework they created has radically changed how researchers approach empirical questions using data from natural experiments or randomised field experiments.

One example of a natural experiment

Let us use a concrete example to illustrate how a natural experiment works. One question that is relevant to both society and to young people considering their future, is how much more you would earn if you chose to study longer. An initial attempt to answer this question could involve looking at data on how people's earnings relate to their education. In every conceivable context, people with more years of education have higher incomes. For example, for men born in the US during the 1930s, earnings were, on average, seven per cent higher for those with one additional year of education.

The association between education and income

There is a clear correlation between years of education and income. But does this mean there is also causation?



The figure uses data from Angrist and Krueger (1991). People with 12 years of education have incomes that are 12 per cent higher than those of people with 11 years of education. People with 16 years of education have 65 per cent higher incomes than people with 11 years of education.

So, can we conclude that an extra year of education adds an extra seven per cent on your income? The answer to this question is no – people who choose a long education differ in many ways from those who choose a short education. For example, some people may be talented at studying and at working. These people are likely to continue studying, but they would still probably have had a high income even if they hadn't. It may also be the case that only people who expect education to pay off choose to study longer.


Similar problems occur if you want to investigate how income affects longevity. The data show that people with higher incomes live longer – but is this really due to their higher incomes, or do these people have other attributes that mean they both live longer and earn more? It is easy to think of more examples where there is reason to question whether correlation actually implies a true causal relationship.

So, how can we use a natural experiment to examine whether additional years of education affect future income? Joshua Angrist and his colleague Alan Krueger (now deceased) showed how this could be done in a landmark article. In the US, children can leave school when they turn 16 or 17, depending on the state where they go to school. Because all children who are born in a particular calendar year start school on the same date, children who are born early in the year can leave school sooner than children born later in the year. When Angrist and Krueger compared people born in the first and fourth quarters of the year, they saw that the first group had, on average, spent less time in education. People born in the first quarter also had lower incomes than those born in the fourth quarter. As adults they thus had both less education and lower incomes than those born late in the year.

Because chance decides exactly when a person is born, Angrist and Krueger were able to use this natural experiment to establish a causal relationship showing that more education leads to higher earnings: the effect of an additional year of education on income was nine per cent. It was surprising that this effect was stronger than the association between education and income, which amounted to seven per cent. If ambitious and intelligent people have both high levels of education and high incomes (regardless of education) the result should have been the opposite; the correlation should have been stronger than the causal relationship. This observation raised new questions about how to interpret the results of natural experiments – questions that were later answered by Joshua Angrist and Guido Imbens.

People born late in the year have more years of education and higher incomes

Additional years of education have a positive effect on income. The figure uses data from Angrist and Krueger (1991).



It would be easy to believe that situations which enable natural experiments are very unusual, especially those that can be used to answer important questions. Research conducted over the past 30 years has shown that this is not the case: natural experiments occur frequently. For example, they may arise due to policy changes in some regions of a country, admission cut-offs in higher education, or income thresholds in tax and benefit systems, which mean that some individuals are exposed to an intervention while other, similar, individuals are not. There is thus unintended randomness that divides people into control and treatment groups, providing researchers with opportunities for uncovering causal relationships.

Understanding labour markets

The effects of a minimum wage

In the early 1990s, the conventional wisdom among economists was that higher minimum wages lead to lower employment because they increase wage costs for businesses. However, the evidence supporting this conclusion was not fully convincing; there were indeed many studies that indicated a negative correlation between minimum wages and employment, but did this really mean that higher minimum wages led to higher unemployment? Reverse causation could even be the issue: when unemployment rises, employers can set lower wages which, in turn, may lead to demands to increase the minimum wage.



To investigate how increased minimum wages affect employment, Card and Krueger used a natural experiment. In the early 1990s, the minimum hourly wage in New Jersey was raised from 4.25 dollars to 5.05 dollars. Just studying what happened in New Jersey after this increase does not give a reliable answer to the question, as numerous other factors can influence how employment levels change over time. As with randomised experiments, a control group was needed, i.e., a group where wages didn't change but all the other factors were the same.

The effect of increasing the minimum wage

Card and Krueger used a natural experiment to study how increasing the minimum wage affects employment.

The researchers identified a treatment group (restaurants in New Jersey) and a control group (restaurants in eastern Pennsylvania) to measure the effect of increasing the minimum wage.

● CONTROL GROUP ● TREATMENT GROUP



1 April 1992: The hourly minimum wage in New Jersey was increased from 4.25 dollars to 5.05 dollars. Despite this, employment in New Jersey was not affected.

Card and Krueger noted that there was no increase in neighbouring Pennsylvania. Of course, there were differences between the two states, but it is likely that the labour markets would evolve similarly close to the border. So, they studied the effects on employment in two neighbouring areas – New Jersey

and eastern Pennsylvania – which have a similar labour market, but where the minimum wage was increased on one side of the border but not the other. There was no apparent reason to believe that any factor (such as the economic situation) apart from the increase in the minimum wage would affect employment trends differently on either side of the border. Thus, if a change in the number of employees was observed in New Jersey, and this differed from any change on the other side of the border, there was good reason to interpret this as an effect of the increase in the minimum wage.

Card and Krueger focused on employment in fast-food restaurants, an industry where pay is low and minimum wages matter. Contrary to previous research, they found that an increase in the minimum wage had no effect on the number of employees. David Card arrived at the same conclusion in a couple of studies in the early 1990s. This pioneering research has led to a large number of follow-up studies. The overall conclusion is that the negative effects of increasing the minimum wage are small, and significantly smaller than was believed 30 years ago.

The work conducted by Card at the beginning of the 1990s also led to new research, which attempted to explain the absence of negative effects on employment. One possible explanation is that companies can transfer increased costs to consumers in the form of higher prices, without significant reductions in demand. Another explanation is that companies that dominate their local labour market can keep wages low; an increased minimum wage thus means that more people want to work, leading to increased employment. When companies have such power over the market, we cannot determine in advance how employment will be affected by changes to the minimum wage. The many studies inspired by Card and Krueger's work have considerably improved our understanding of the labour market.

Research on immigration and education

Another important issue is how the labour market is affected by immigration. To answer this question, we need to know what would have happened if there had not been any immigration. Because immigrants are likely to settle in regions with a growing labour market, simply comparing regions with and without many migrants is not enough to establish a causal relationship. A unique event in the history of the US gave rise to a natural experiment, which David Card used to investigate how immigration affects the labour market. In April 1980, Fidel Castro unexpectedly allowed all Cubans who wished to leave the country to do so. Between May and September, 125,000 Cubans emigrated to the US. Many of them settled in Miami, which entailed an increase in the Miami labour force of around seven per cent. To examine how this huge influx of workers affected the labour market in Miami, David Card compared the wage and employment trends in Miami with the evolution of wages and employment in four comparison cities.

Despite the enormous increase in labour supply, Card found no negative effects for Miami residents with low levels of education. Wages did not fall and unemployment did not increase relative to the other cities. This study generated large amounts of new empirical work, and we now have a better understanding of the effects of immigration. For example, follow-up studies have shown that increased immigration has a positive effect on income for many groups who were born in the country, while people who immigrated at an earlier time are negatively affected. One explanation for this is that the natives switch to jobs that require good native language skills, and where they do not have to compete with immigrants for jobs.

Card has also made important contributions regarding the impact of school resources on students' future success on the labour market. Once again, his results questioned the received wisdom – previous research suggested that the relationship between increased resources and school performance, as

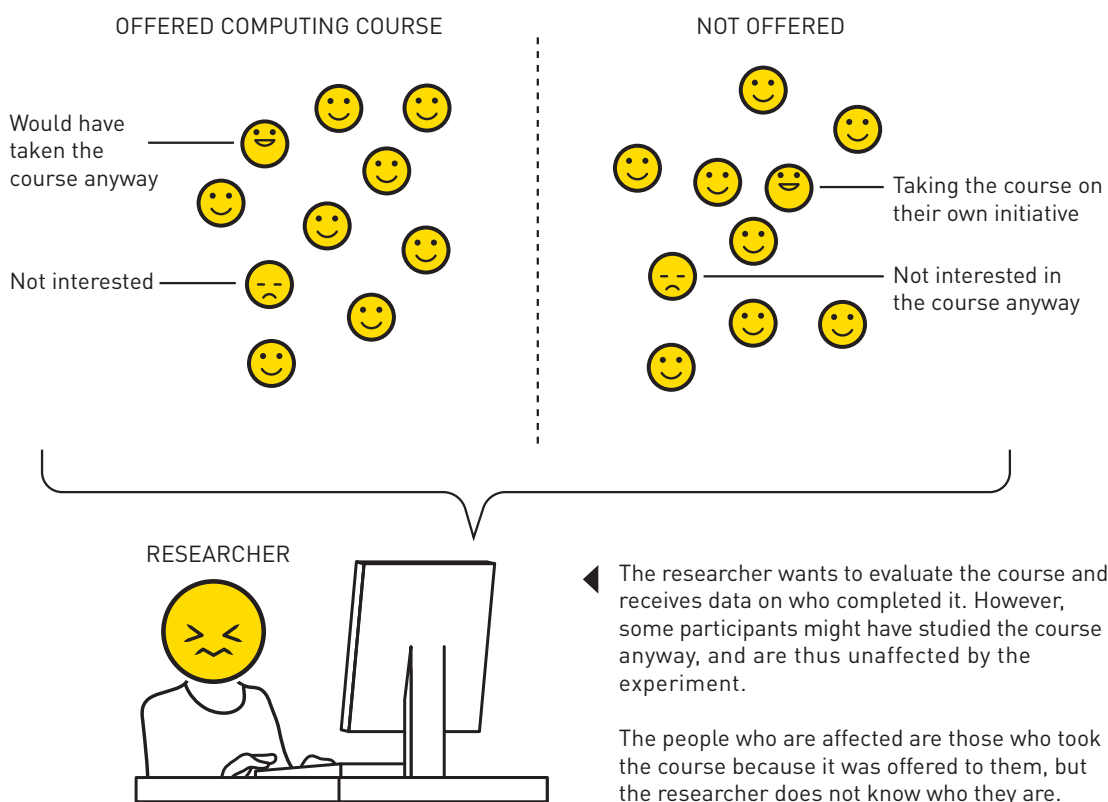
well as labour market opportunities later in life, was weak. However, one problem was that previous work had not considered the possibility of compensatory resource allocation. For example, it is likely that decision-makers invest more in educational quality in schools where student achievement is low.

To examine whether school resources have an impact on students' future labour market success, David Card and Alan Krueger compared the returns on education for people who lived in the same state in the US, but who had grown up in different states – for example, people who had grown up in Alabama or Iowa, but now lived in California. The idea is that people who have moved to California and have the same level of education are comparable. If the returns on education differ, this is probably due to Alabama and Iowa having invested differently in their education systems. Card and Krueger found that resources are important: returns on education increased with teacher density in the state in which the individuals had grown up.

This research also inspired many new studies. There is now relatively strong empirical support to show that investments in education influence students' later success on the labour market. This effect is particularly strong for students from disadvantaged backgrounds.

Local average treatment effect

Joshua Angrist and Guido Imbens showed how natural experiments can be used to arrive at precise conclusions about cause and effect. Natural experiments differ from clinical trials as the researcher does not have complete control over who receives the treatment.



A new framework for studies of causal relationships

In all realistic scenarios, the effect of an intervention – for example, the effect of additional schooling on earnings – varies between people. Moreover, individuals are affected differently by a natural experiment. The opportunity to leave school at 16 will hardly affect those who had already planned

to go to university. Similar problems arise in studies based upon actual experiments, because we cannot typically force individuals to participate in an intervention. The subgroup who eventually chooses to participate probably consists of individuals who believe that they will benefit from the interventions. However, a researcher who analyses the data only knows who has participated, not why – there is no information about who participated solely because they were offered the opportunity, thanks to the natural experiment (or the randomised experiment), and who would have done so anyway. How can a causal relationship between education and income be established?

Joshua Angrist and Guido Imbens tackled this problem in an influential study from the mid-1990s. More specifically, they asked the following question: Under which conditions can we use a natural experiment to estimate the effects of a particular intervention, such as a computing course, when the effects vary between individuals and we do not have complete control of who participates? How can we estimate this effect and how should it be interpreted?

Somewhat simplified, we can imagine a natural experiment as if it randomly divides individuals into a treatment group and a control group. The treatment group is entitled to participate in a programme while the control group is not. Angrist and Imbens showed that it is possible to estimate the effect of the programme by applying a two-step process (known as the instrumental variables method). The first step investigates how the natural experiment affects the probability of programme participation. The second step then considers this probability when evaluating the effect of the actual programme. Given a few assumptions, which Imbens and Angrist formulated and discussed in detail, the researchers can thus estimate the impact of the programme, even when there is no information about who was actually affected by the natural experiment. One important conclusion is that it is only possible to estimate the effect among the people who changed their behaviour as a result of the natural experiment. This implies that Angrist and Krueger's conclusion about the effect on income of an additional year of education – which they estimated to be nine per cent – only applies to those people who actually chose to leave school when given the chance. It is not possible to determine which individuals are included in this group, but we can determine its size. The effect for this group has been named the local average treatment effect, LATE.

Joshua Angrist and Guido Imbens thus showed exactly what conclusions about cause and effect can be drawn from natural experiments. Their analysis is also relevant for randomised experiments where we do not have complete control over who participates in the intervention, which is the case in almost all field experiments. The framework developed by Angrist and Imbens has been widely adopted by researchers who work with observational data. By clarifying the assumptions necessary to establish a causal relationship, their framework has also increased the transparency – and thus credibility – of empirical research.

A revolution in empirical research

The Laureates' contributions from the early 1990s demonstrate that it is possible to answer important questions about cause and effect using natural experiments. Their contributions complement and strengthen one another: Angrist and Imbens' methodological insights about natural experiments and Card's applications of this approach to important questions led the way for other researchers. We now have a coherent framework which, among other things, means that we know how the results of such studies should be interpreted. The work of the Laureates has revolutionised empirical research in the social sciences and significantly improved the ability of the research community to answer questions of great importance to us all.

FURTHER READING

Additional information on this year's prizes, including a scientific background in English, is available on the website of the Royal Swedish Academy of Sciences, www.kva.se, and at www.nobelprize.org, where you can watch video from the press conferences, the Nobel Lectures and more. Information on exhibitions and activities related to the Nobel Prizes and the Prize in Economic Sciences is available at www.nobelprizemuseum.se

The Royal Swedish Academy of Sciences has decided to award the Sveriges Riksbanks Prize in Economic Sciences in Memory of Alfred Nobel 2021

with one half to

DAVID CARD

Born 1956 in Guelph, Canada. Ph.D. 1983 from Princeton University, USA. Class of 1950 Professor of Economics, University of California, Berkeley, USA.

and the other half jointly to

JOSHUA D. ANGRIST

Born 1960 in Columbus, Ohio, USA. Ph.D. 1989 from Princeton University, USA. Ford Professor of Economics, Massachusetts Institute of Technology, Cambridge, USA.

GUIDO W. IMBENS

Born 1963 in the Netherlands. Ph.D. 1991 from Brown University, Providence, USA. The Applied Econometrics Professor and Professor of Economics, Stanford University, USA.

“for his empirical contributions to labour economics”

“for their methodological contributions to the analysis of causal relationships”