International Journal of Economics and Finance; Vol. 14, No. 10; 2022
ISSN 1916-971X E-ISSN 1916-9728
Published by Canadian Center of Science and Education

Reflections on Learning from Observational Data

Caleb Piche-Larocque[1], Joseph Findlay[1] & Akhter Faroque[1]

[1] Department of Economics, Laurentian University, Ontario, Canada

Correspondence: Akhter Faroque, Department of Economics, Laurentian University, Ontario, P3E 2C6, Canada.

Received: July 4, 2022 Accepted: September 14, 2022 Online Published: September 18, 2022

doi:10.5539/ijef.v14n10p56 URL: https://doi.org/10.5539/ijef.v14n10p56

Abstract

The social sciences study various aspects of human behaviour – social, economic and political – based on observational data. Observational data are inaccurate and subject to simultaneity, seasonality, structural breaks, random variation and too many interlocking variables masking the underlying causal patterns. During the past two decades or so, the use experimental data (RCTs) has become widely popular across the social sciences, creating a tension between the supporters and critics of the new and the old methodologies. In this paper, we first review these methodologies, both observational and experimental, focusing on how economists and other social scientists try to learn about the underlying causal relationships from the correlations contained in the data. We then reflect on whether the new or the old methodologies should be the way forward from a purely statistical and a broader policy and development perspectives.

Keywords: correlation, causality, observational data, Random Cotrolled Trials (RCTs)

1. Introduction

The social sciences study various aspects of human behaviour - economic, social and political -based on observational data. The natural sciences study the behaviour of the cosmos, plants and the animal kingdom based on experimental data. The purpose of this paper is to outline the difficulties of learning and how social scientists, especially, economists try to learn from observational data.

The natural laws that govern the behaviour of the cosmos, plants and the animal kingdom are fixed and predictable. Experimental data used in the natural sciences are collected under controlled environments, often using sophisticated machines. Such data are devoid of measurement errors, structural breaks, and have clear causal patterns.

In contrast, the social and economic principles that govern human behaviour are irregular and unpredictable. The decisions and choices people make are influenced by many uncontrolled influences, both past, present and future. National statistical agencies collect data based on various types of surveys; as such, they contain significant human and measurement errors. Thus, in comparison to experimental data, observational data are inaccurate and messy, being subject to structural breaks, simultaneity, seasonality, random variation and too many interlocking variables masking the underlying causal patterns.

The remainder of the paper has three sections. Section 2 describes the three essential aspects to learning from observational data. Section 3 highlights how economists and other social scientists combine theory/prior knowledge and various statistical techniques to estimate the true causal relationships from observational data. This section also describes the wide popularity of random controlled trials (CRTs) as a way of conducting causal analysis in the social sciences. Section 4 reflects on the advantages and disadvantages of reliance on experimental vs. non-experimental methods from both the purely statistical and the broader policy perspectives. Section 5 concludes the paper.

2. Three Dimensions to Learning from Observational Data

The attempt to understand human behaviour from observational data has, of necessity, three essential dimensions: data visualization, data transformation and retrieving causation from correlation contained in the data.

Data visualization is an important preliminary step, since numbers often fool us. Human brains are naturally wired to process information visually with graphs and charts, less so with numbers.

To see the power of visual over numerical methods, one only needs to consider the four pairs of artificially generated data - known as Anscombe's Quartet (Anscombe, 1973). Each of the four pairs of samples has the same mean, same variance, same correlation, and the same linear regression. However, it is only when they are graphed that we immediately see the very different patterns contained in them that different interpretations.

The second is that data transformation is often necessary prior to analysis, since we tend to see pattern even where there is none. Granger and Newbold (1974) demonstrated that when two time series variables X and Y are generated by two completely independent non-stationary (unit root) processes, then the OLS regression of one variable on the other gives "spurious regression" effects, *even though there is no actual relationship between X and Y.* Data transformation can convert non-stationary into stationary data prior to analysis and avoid spurious effects.

The third aspect of learning from observational data is the most challenging, as it necessitates the use of sophisticated econometric techniques in order to tease out the underlying causal patterns. Knowledge of the true causal relationship is important both for the development of theory and as a guide for effective policy intervention.

Consider the case of national unemployment, which may be caused by deficient aggregate demand for goods and services or by structural problems specific to the labour market. In the absence of accurate advance knowledge of the underlying cause of the unemployment, policy intervention may be completely ineffective. For instance, the use of expansionary fiscal or monetary policy will be completely misguided, if the true cause of unemployment is structural, such as rigid high real wages or high minimum wages.

3. Difficulties of Drawing Causal Inference in the Social Sciences

We have already noted above that observational data are messy because they are subject to structural breaks, seasonality, simultaneity and too many interlocking variables masking the underlying causal patterns. Presence of structural breaks (Hansen, 2001) and seasonality (Barnett & Dobson, 2010) in the data can be testing and dealt with in straightforward ways. However, simultaneity and learning about causality poses a much bigger challenge to social scientists. To illustrate the nature of the difficulties, here we consider a simple example from economics.

Every first-year economics student is taught that market demand for any product depends on the market price of the product and on many other variables, some of which are observable while others are unobservable or costly to observe. For example, the demand for orange juice depends on the price of orange juice, consumers' average income, prices of substitutes, prices of complements all of which are observables for which we have data. But market demand for Orange juice may also depend on consumers' taste and preferences and on their expectation of the future price of orange juice, which are unobservable. Our objective is to recover the unbiased (causal) effect of price (X) on quantity (Y) of orange juice from market data. To do so, we specify the following regression equation

$$Y = \alpha + \beta X + \theta_1 W_1 + \theta_2 W_2 + \ldots + \theta_k W_k + \varepsilon \tag{1}$$

where, the W's are control variables dictated by economic theory, which include all observable determinants of the demand for orange juice besides price (X) for which we have data. A fundamental problem is that the coefficient β attached to price (X) will not measure the true causal effect of price (X) on quantity (Y) even in this model (1) that controls for all observable determinants Y. This is because X is in fact an endogenous variable since price (X) and quantity (Y) are simultaneously determined. This means that X is correlated to the error term ε, so that the least square estimate of β will include not just the effect of a change in X but also of ε. This same problem also arises when X is correlated to an omitted variable (e.g., an unobserved determinant of Y (e.g., determinants of consumers' taste and preferences) or when any of the covariates contain non-random measurements errors. Thus, even if our estimate of β is statistically significant, this simply means that price (X) and quantity (Y) are strongly correlated, but it will not measure the unbiased (causal) effect of price on quantity.

It is at this juncture that applied economists and econometricians have made a major contribution to advancing social sciences research: they have proposed methods for breaking the correlation between X and ε and recovering the true value of β. More generally, they have proposed a set of statistical procedures that are suitable for determining the causal effect of any intervention/policy of interest on an outcome variable, in a variety of contexts across the social science. Below we review these procedures, including instrumental variable (IV), regression discontinuity, natural experiments, differences-in-differences (DID) and random controlled trials (RCTs) estimation.

3.1 Instrumental Variable Estimation

If all Ws are orthogonal to X and ε, then we can rewrite (1) simply as

$$Y = \alpha + \beta X + \varepsilon \tag{2}$$

without affecting the least squares estimate of our parameter of interest β,

$$\hat{\beta} = \frac{Cov\ (X,Y)}{Var(X)}$$

which after substitution for Y may be written as:

$$\hat{\beta} = \beta + \frac{Cov(X,\varepsilon)}{Var(X)} \tag{3}$$

It is easier to see from equation (3) why the least squares estimate $\hat{\beta}$ will not measure an unbiased (causal) effect of X on Y as long as X is correlated to ε for any of the reasons stated above. Back in 1915, Philip G. Wright proposed a solution to this problem (Wright, 2015). He suggested that if we can find a variable Z that is strongly correlated with X, but uncorrelated with ε, we could obtain an unbiased estimate of β by running the following two-step regressions:

Step 1: Regress: $Z = \alpha + \beta X + \theta_1 W_1 + \theta_2 W_2 + \ldots + \theta_k W_k + \varepsilon$ (get predicted series Z)

Step 2: Regress: $Y = \alpha + \beta \hat{Z} + \theta_1 W_1 + \theta_2 W_2 + \ldots + \theta_k W_k + \varepsilon$ ($\hat{\beta}$ in step 2 is unbiased)

3.2 Regression Discontinuity

In classroom discussion, instructors may ask students to perform a ceteris-paribus thought experiment to illustrate how, in theory, we could determine the true causal effect of price (X) on demand for orange juice (Y) in the context of equation (1) above. Suppose, we raise only the price of orange juice (X) by one unit, holding all the other determinants of demand (Y), both observable and unobservable, constant. Then the response of Y due to the one unit rise in X, that is, $\frac{\delta Y}{\delta X}$, will denote the true causal effect of price on quantity. Note that in this thought experiment we are measuring the response of Y by comparing the levels of Y after and before the change in price, holding all other determinants of Y constant. In the real world, however, with observational data, it is very difficult to hold all other things constant to estimate of the causal effect we seek without introducing some sort of (selection) bias in our estimate.

Thistlethwaite and Campbell (1960) - a psychologist and a statistician – who first introduced the regression discontinuity design (RDD) in social sciences research showed that, under some conditions, it is possible to recover the unbiased (causal) effect even with observational data. This is true in the context of program evaluations or interventions that have a cut-off point determining who is eligible to participate and who is not. The most distinctive characteristic of RDD is that participants are (randomly) assigned to program or comparison groups solely based on a cut-off score on a pre-program measure. Researchers can then estimate the causal effect of the cut-off/threshold by comparing subjects on each side of the cut-off/threshold.

For example, senior high school students are awarded merit scholarships based on whether or not they score above a threshold value on the PSAT test. Suppose we want to know if students who receive scholarships score a higher GPA in college/university years, compared to students who do not receive scholarships. We can answer this question by comparing the college/university performance of students who are just above the threshold with the performance of those are just below the threshold. This is because students who are just above or just below the threshold are likely to be similar in characteristics such as IQ and study habits, but only students above the threshold receive scholarships. The regression discontinuity design exploits this fact to estimate causal effect of the threshold. Two excellent references on RDD are the review article by Lee and Lemieux (2010) and the text by Angrist and Pischke (2009).

3.3 Quasi Experiment or Natural Experiments

In some situations, extreme natural or man-made events may cut through the problems of simultaneous causation and reveal the true causal direction. Rare events such as a major war, earthquake, a pandemic, or the sudden bursting of a housing bubble may act as a quasi-experiment, sometimes called "natural" experiment, that leave little doubt as to what is the cause and what is the effect.

For example, GDP (Y) and government budget (D) are typically simultaneously related, making the estimation of 'multiplier effects' of fiscal policy very difficult. The Keynesian economists estimate this multiplier to be around 1.5; while conservatives believe it is zero. In 2008, many countries suffered a 'deleveraging shock' - a sudden realization by households that debt levels are excessive and they cut back in consumption expenditure in order to pay off debt. The evidence shows that countries with higher debt levels have experienced a bigger deleveraging shock (cutbacks in private consumption expenditure) and have suffered deeper recessions. In this case, there is little doubt that drop in aggregate demand caused the drop in output Y, and not the other way around. For more details and applications of natural experiments see Dunning (2012) and Rosenzweig and Wolpin (2000).

3.4 The Differences-in-Differences Estimator (Panel Data)

In a quasi-experiment, the researcher does not have control over the randomization, so that some differences might remain between the treatment and control groups even after controlling for W in equation 1 above. One way to adjust for those remaining differences between the two groups is to compare not the outcomes Y but the change in the outcomes pre- and post treatment, thereby adjusting for differences in pre-treatment values of Y in the two groups. Because this estimator is the difference across groups in changes, or difference over time, it is called the differences-in-differences estimator. For example, in his study of the effect of immigration on low-skilled workers' wages, Card (1990) used a differences-in-differences estimator to compare the change in wages in Miami with the change in wages in other U.S. cities. For details, see Stock and Watson (2020).

3.5 Retrieving the Causal Effect Using Experimental Data

The best way to isolate cause and effect and make sure that we are studying the effect of changing only explanatory variable while keeping all other covariates constant, is to perform a randomized controlled trial, as advocated by Angrist and Pischke (2009; 2010). In fact, the impact of RCTs has been so large during the past two decades that many writers have come to refer to RCTs as the 'gold standard' for doing causal analysis in the social sciences (Note 1).

RCTs are a study design that draws a random sample from a population of interest and then randomly assigns the sample participants into a treatment group or a control group. This procedure allows researchers to identify and measure the causal effect of a treatment or policy intervention in a variety of fields. For example:

- The causal effect of a job training program on employment (labour economics)
- The causal effect of a new drug on health (health economics)
- The causal effect of micro-credit programs on poverty (development economics; Abhijit Banerjee)

Suppose, a pharmaceutical company wants to know the causal effect of a new drug (X) on some health outcome (Y). There is no past data, how do we proceed?

- Randomly select a sample of, say, 500 individuals from a population of adults so that the sample is representative of the population.
- Then randomly (by flipping a coin) assign the individuals into two groups: Treatment Group and a Control Group, so that the two groups are identical in every respect, except for the fact that the Treatment Group is given a fixed dose of the new drug but not the control group.

Then, the causal effect of the drug can be measured by running a simple regression.

$$Y_i = \beta_0 + \beta_1 X_i + e_i \tag{4}$$

where X_i is the treatment level and e_i is the error term that contains all of the omitted determinants of outcome Y_i. If treatment is the same for all members in the treatment group, then X_i is binary: $X_i = 1$ indicates that the ith individual received the treatment and $X_i = 0$ indicates that he/she did not receive the treatment. Then OLS will give an unbiased estimate of β_1 - the causal effect of the treatment on the population. This method of estimation of causal effect is called the 'Differences Estimator'.

4. Reflections on the Use of Experimental and Non-Experimental Methods in the Social Sciences

We devote this section to an analysis of the relative merits and demerits of experimental vs. observational procedures for causal analysis in the social sciences. Our goal here is not to try to rank the various methods, simply because no one method can claim to be systematically superior to all other methods in all circumstances. All procedures, including experimental ones, require assumptions and when those assumptions are/are not valid in a particular context, the causal conclusions that follow from those methods are likely to be valid/invalid. Thus, the issue of which of method should to used really depends on the context, what question being investigated, the assumptions that can be acceptably employed and on the costs are of different kinds of mistakes (Deaton & Cartwright, 2016).

The fundamental difficulty of conducting causal analysis in the social silences is one of replicating, in practice, the ceteris paribus principle of changing only one explanatory variable while holding other covariates constant.

This difficulty pervades both observational and experimental procedures. Observational studies try to overcome this difficulty by combining theory or prior knowledge to identity and statistical techniques to estimate the causal hypothesis of interest and other control covariates. By contrast, experimental studies (CRTs) do not require any specialized knowledge of the subject matter or any prior knowledge. Instead, such studies replace theory/prior knowledge with randomization of the selection process, in the belief that randomization makes the treatment and the control groups the same in all respects, except for one (the treatment), thus eliminating any selection bias. However, as it happens, even randomization may not equalize everything else but the treatment, so that biases can confound causal estimates even in experimental studies.

Below we examine the potential sources bias in observational and experimental studies first from a purely statistical and then a broader policy perspective to justify our viewpoint.

First, estimation biases in observational studies may arise because theory and/or prior expert knowledge may fail to impose the restrictions necessary to identity the true causal relationship. This is especially likely to be true in areas of application where theory is not well developed or where expert knowledge is lacking. For example, in our own discipline of economics, static theories are usually fairly well developed, but such theories have little to say about the dynamic interactions between the response and the explanatory variables. Consequently, causal estimates may be contaminated by biases due to failure to specify sufficient dynamics. Biases can also arise from reliance on weak instrumental variables due to lack of knowledge of strong instrumentals and natural experiments that may not satisfy the ceteris paribus principle of keeping other variables constant, except for the causal variable of interest.

Experimental studies (CRTs) conducted under ideal conditions where randomization creates a perfect balance of all other covariates between the treatment and control groups, except for the treatment, can result in unbiased causal estimates. However, in practice, even randomization is susceptible to several sources of potential biases. Randomization, in fact, must occur at two separate stages. Stage one involves drawing a random sample from a population whose properties the investigator wishes to study. In practice, however, stage one selection may be based on politics or convenience that introduces bias. In stage two, the individuals in the sample are assigned randomly into a treatment and a control group, but can be several post-randomization sources of biases. For example, individuals in the treatment group may not comply or forget to comply with the experimenter’s instructions, especially if they dislike the instructions or the instructions violate their beliefs. Similarly, individuals in the control group may fail to follow their instructions. These reasons – and others – may contaminate the causal effect of interest.

Second, RCTs are based on the assumption that the difference between the means of the treatment and the control groups measures the true average treatment effect (ATF) for the population being studied. This, however, is true only when the RCT is repeated many times on the same population, which is not done in practice. In a single trial, the difference in means between the treatment and the control groups will be equal to the ATF plus a term that reflects the imbalance in the net effects of all other causes. The net average balance of other causes (error term) will not be eliminated in any single trial of the RCT and nothing in the randomization limits its size of this source of bias.

Third, sample size matters for both observational and experimental studies. In general, sample sizes in observational studies range from medium to large, and in some cases, they can be very large, as in the case of unemployment and income household surveys. By contrast, sample size in CRTs are generally small: less than a hundred in medical interventions; less than a thousand plots in interventions of studies of economic development. Thus causal estimates from CRTs, even if they are unbiased, will tend to be less imprecise (have larger variance) compared to the estimates derived from observational studies, which may be biased. This trade off between bias and precision in observational vs. experimental studies makes the choice between them complicated. In some cases, the greater precision may dominate the size of unbiasedness and so an observational study may be preferable; in other cases, the opposite may be true.

Finally, like many non-experimental studies, RCTs are sensitive to outliers in the data and to asymmetrical distributions of the treatment effect. The effects of many treatments are asymmetric. For example, in a micro-financing scheme, a few talented, but credit-constrained entrepreneurs may experience a large and positive effect, while there is no effect for the majority of borrowers. Similarly, a health intervention may have no effect on the majority, but a large effect on a small group of people. Experimentations conducted by Deaton and Cartwright (2016) shows that an RCT can yield completely different results depending on whether an outlier falls in the treatment or the control group.

Turning to the broader policy and development front, here too we think that no one method, observational or experimental, can claim to have a special status in the tool-kit of statistical methods available to social scientists. The simplicity of RCTs identifying the cause and effect through randomization may seem like an important advantage over observational studies dependent on theory or prior knowledge. But, even if a RCT is able to produce an unbiased causal estimate in a specific context, trying to transport that result to other contexts, to other place, populations, and countries, is likely to run against the grains of heterogeneous cultural, behavioural, religious, social, and economic practices across populations. Simple extrapolations of the results of a particular CRT to other (broader) populations is likely to require additional assumptions or more general theories of human behaviour to interpret the results of the CRT.

The current debate between the proponents and critics of experimental (CRT) methods in social science research is reminiscent of a similar debate that took place in the 1940s between the proponents (Burns & Mitchel, 1946) and critics (Koopmans, 1947) when leading indicators was introduced as a new method for business cycle research in economics. At the heart of the debate, then as mow, is the role theory or prior knowledge in research methodology. In the current debate, our views on the subject is shaped by the many writings of Deaton and Cartwright (1916; 2018) and Cartwright and Deaton (2016). We think it is appropriate to use a direct quote from these authors that expresses our own views regarding the way forward for social science research from here:

"Economists and other social scientists know a great deal, and there are many areas of theory and prior knowledge that are jointly endorsed by large numbers of knowledgeable researchers. Such information needs to be built on and incorporated into new knowledge, not discarded in the face of aggressive know-nothing ignorance. The results of RCTs must be integrated with other knowledge, including the practical wisdom of policy makers if they are to be usable outside the context in which they were constructed." (Deaton & Cartwright, 2016).

5. Conclusions

This paper outlines the difficulties of learning about human behaviour based on observational data. These difficulties arise from two facts: observational data contain significant human and measurement errors and the social and economic principles that govern the generation of observational data are variable and unpredictable. The decisions and choices of people are subject to many uncontrolled influences past, present and future, so that observational data are subject to simultanety, seasonality, structural breaks, random variation and too many interlocking variables mask the underlying causal patterns.

Econometricians have devised alternative methodologies for dealing with the confounding features of observational data. In particular, they have developed alternative techniques for extracting the causal effects from the correlations contained in observational data, including instrumental variables, regression discontinuity, quasi experiments, and differences-in-differences estimation and, in more recent years, conducting random controlled trials. Each of these estimation methods is based on its own set of assumptions and is applicable only in specific circumstances that satisfy those assumptions. None is universally applicable to all situations.

If the goal of the researcher is to estimate the unbiased average causal effect of a treatment only for subjects included within a sample, then there no better way to achieve this than to conduct a random controlled trial (RCT). However, if the goal is extended to draw inferences about the population from which the sample is drawn, then RCTs compared to large-sample non-experimental data, involve a trade-off between bias and efficiency or precision. Any attempt to further generalize the results from a RCT to other populations is even more problematic, because of heterogeneity across populations. We conclude this paper by suggesting that the choice between experimental and observational data is not one or the other, rather both methodologies can make valuable contributions to our understanding of the social, economic and political behaviour of human beings.

References

Angrist, J. D., & Pischke, J. S. (2009). *Mostly Harmless Econometrics*. Princeton: Princeton University Press. https://doi.org/10.1515/9781400829828

Angrist, J. D., & Jörn-Steffen, P. (2010). The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics. *Journal of Economic Perspectives, 24*(2), 3-30. https://doi.org/10.1257/jep.24.2.3

Anscombe, F. J. (1973). Graphs in Statistical Analysis. *American Statistician, 27*(1), 17-21. https://doi.org/10.2307/2682899

Barnett, A. G., & Dobson, A. J. (2010). *Analysing Seasonal Health Data*. Springer. https://doi.org/10.1007/978-3-642-10748-1

Burns, A. F., & Mitchell, W. C. (1946). Measuring Business Cycles. *Studies in Business Cycles*, No. 2. National Bureau of Economic Research.

Card, D. (1990). The Impact of the Mariel Boatlift on the Miami Labor Market. *Industrial and Labor Relations Review*, *43*(2), 245-257. https://doi.org/10.2307/2523702

Deaton, A., & Cartwright, N. (2016). Understanding and Misunderstanding Randomized Controlled Trials. *NBER Working Paper No. 22595*. https://doi.org/10.3386/w22595

Deaton, A., & Cartwright, N. (2018). *Reflections on CRTs.* Retrieved from https://par.nsf.gov/servlets/purl/10067322

Dunning, T. (2012). *Natural Experiments in the Social Sciences: A Design-Based Approach*. Cambridge, UK: Cambridge University Press. https://doi.org/10.1017/CBO9781139084444

Granger, C. W. J., & Newbold, P. (1974). Spurious regressions in econometrics. *Journal of Econometrics*, *2*(2), 111-120. https://doi.org/10.1016/0304-4076(74)90034-7

Hansen, B. E (2001). The New Econometrics of Structural Change: Dating Breaks in U.S. Labor Productivity. *Journal of Economic Perspectives*, *15*(4), 117-128. https://doi.org/10.1257/jep.15.4.117

Imbens, G., & Lemieux, T. (2008). Regression Discontinuity Designs: A Guide to Practice. *Journal of Econometrics*, *142*(2), 615-635. https://doi.org/10.1016/j.jeconom.2007.05.001

Koopmans, T. C. (1947). Measurement Without Theory. *The Review of Economics and Statistics*, *29*(3), 161-172. The MIT Press. https://doi.org/10.2307/1928627

Lee, D. S., & Lemieux, T. (2010). Regression discontinuity design in economics. *Journal of Economic Literature*, *48*, 281-355. https://doi.org/10.1257/jel.48.2.281

Online debate. (2015). Between Abhijit Banerjee and Angus Deaton on the merits of RCTs. Retrieved from https://nyudri.wordpress.com/initiatives/deaton-v-banerjee/.

Nancy, C., & Angus, D. (2016). *The limitations of randomised controlled trial.* Retrieved from https://cepr.org/voxeu/columns/limitations-randomised-controlled-trials.

Rosenzweig, M. R., & Kenneth, I. W. (2000). Natural Experiments in Economics. *Journal of Economic Literature*, *38*(4), 827-874. https://doi.org/10.1257/jel.38.4.827

Stock and Watson. (2020). *Introduction to econometrics* (4th ed.). New York: Pearson.

Thistlethwaite, D., & Campbell, D. (1960). Regression-Discontinuity Analysis: An alternative to the ex post facto experiment. *Journal of Educational Psychology*, 51(6), 309-317. https://doi.org/10.1037/h0044319

Wright, P. G. (1915). Moore's Economic Cycles (a review). *The Quarterly Journal of Economics*, *29*(3), 631-641. https://doi.org/10.2307/1885466

Note

Note 1. Joshua Angrist, David Card, and Guido Imbens were awarded the 2021 Nobel Prize in Economic Sciences for their contributions to labor economics and the analysis of natural experiments.