

Joshua Angrist

On charter schools, the elite illusion, and the “Stones Age” of econometrics

As a teenager growing up in Pittsburgh, Joshua Angrist became fed up with high school and said his goodbyes to it after his junior year. Today, at the Massachusetts Institute of Technology, he’s a top researcher in labor economics and the economics of education — with work that includes a series of famed studies of policy choices for K-12 schooling.

Much of his work has been based on ingenious “natural experiments,” that is, episodes in which two or more groups of people were randomly exposed to different policies or different experiences. Such occurrences are an opportunity for Angrist and his co-authors to use the tools of econometrics to assess the effects of those differences — whether that’s a large classroom versus a small classroom or education at a charter school versus education at a conventional public school.

Angrist’s first natural experiment looked at labor market outcomes for men who were drafted during the Vietnam War era compared with those of men who weren’t drafted. The idea came to him from his labor economics teacher and Ph.D. adviser at Princeton, Orley Ashenfelter, who mentioned in class one day that he had seen a news article about a study in the *New England Journal of Medicine* in which epidemiologists investigated the long-term health effects of being drafted.

“They had done this very clever thing where they used the fact that draft lottery numbers were randomly assigned,” Angrist remembers, “and they compared people who had high and low numbers to test the causal effects.”

Ashenfelter remarked to the class that this use of the draft lottery was a great idea and that somebody should use it to look at the effects of the draft on the men’s earnings. Angrist agreed; immediately after class, he went to the library to start the research that became his doctoral thesis.

Angrist found that in the early 1980s, well after the end of the war, veterans — whether they served in Vietnam or elsewhere — took an earnings hit of around 15 percent compared with non-veterans in the same period. (Angrist himself served as a paratrooper in the Israeli army before he went to grad school.)

In addition to his research and teaching responsibilities at MIT, Angrist is a co-founder and co-director of MIT’s School Effectiveness and Inequality Initiative. He is the author, with Jörn-Steffen Pischke, of the



econometrics textbooks *Mostly Harmless Econometrics: An Empiricist’s Companion* (2009) and, for undergraduates, *Mastering ‘Metrics: The Path from Cause to Effect* (2015). He also teaches econometrics in a series of free videos offered through the nonprofit Marginal Revolution University.

David A. Price interviewed Angrist via videoconference in March 2020.

EF: You left high school early and then took a while to decide to go to college. Why?

Angrist: I wasn’t a very good high school student. I didn’t pay much attention to school, and I didn’t go very often. I also liked to work because I wanted money. So I started working as a busboy, and I liked having money because money is good when you’re a teenager. Then I thought, well, I’ll just leave school if I can and go work full time.

I figured out that the Pennsylvania high school graduation requirements were pretty minimal. So I was able to graduate with a bare bones diploma after my junior year. Then I worked for a while, mostly in institutions for the intellectually disabled because I had experience with this kind of work at summer camp. But from my work experience, I realized I should probably go to college.

EF: Was there anything in particular that brought home to you the idea that you should go to college?

Angrist: I saw that the work was not that interesting. And it wasn't clear where it would lead. So I got bored.

“Empiricists should work on stuff that’s nearby. Then you can have some visibility into what’s unique.”

to just mimic whatever the Americans and British are doing. I think a better strategy is to say, “Well, what’s special and interesting about where I am?”

EF: At some point, you became interested in economics. How did that happen?

Angrist: Because of the work I’d been doing, I thought I would go into special education, so I took a psych course. But I also took Econ 101, and I had a wonderful economics teacher, a man named Bob Piron. At least he was very much to my taste. He was funny. He was provocative. He would call on people; he would tease people. Today, I suppose he would get fired for this, but as a teacher, I try to do a version of what he did that’s a little more in tune with the times.

Piron was clear as a bell, and he was challenging, and I just enjoyed the whole thing. So I thought, I want more of this. I guess there was also the fact that I did like the material and I had an affinity for it. I started taking all the econ I could get and all the math I could stomach.

I also experienced another example of the power of having good teachers. Oberlin, where I went, is a small liberal arts college. It has a good economics department and they invite their best seniors to write a thesis, so I did it. What Oberlin does that’s very nice is they invite an outside examiner, somebody who is a well-known researcher. They invited Orley Ashenfelter, a labor economist at Princeton.

Not everybody wants to come to Oberlin for a few days, but Orley did that; he’s that kind of person. I met him and got to know him a little bit and he took a shine to me, I was lucky, and he said, “Why don’t you come to Princeton?” So after I graduated from Oberlin and after my Israeli army service, I went to graduate school at Princeton because of him and to work with him.

CLASS SIZE

EF: A lot of your recent research has focused on evaluating K-12 schools and education policies. One of your many well-known articles in this area found that limits on class size were associated with higher test scores for fourth and fifth graders in Israel. How did that work come about?

Angrist: It came out of the fact that I was living in Israel at the time and was working with Victor Lavy. He was my main collaborator while I was on the faculty at Hebrew University. Victor and I started writing papers about Israeli schools and have continued to do so ever since.

One thing I learned is that empiricists should work on stuff that’s nearby. Then you can have some visibility into what’s unique and try to get on to projects that other people can’t do. This is particularly true for empiricists who are working outside the United States. There’s a temptation

And it turned out that the Israeli school system had a lot of interesting things going on. One was that they had a rule about class size that can actually be dated back to the Talmud. Even though the details of the rule have changed, we call it Maimonides’ Rule, because the biblical sage and scholar Moses Maimonides had said in the 13th century that that’s what you’re supposed to do.

If you’re in a grade cohort of 41, they’ll split your class because you’re over the cap of 40; if you’re in a cohort of 39, you’ll stay lumped. So you get a nice natural experiment there. I think that paper is more methodologically significant than substantive in that it was one of the first of a wave of regression discontinuity studies. But we did find that larger class sizes reduced test scores. It’s a story of selection bias: Larger classes are in the areas that are more densely populated, and in Israel in particular, that’s urban areas and richer people.

A couple of years ago, Victor and I went back with much more data, and we replicated the Maimonides study again on Israel. In the original study, we had two years of data, from ’91 and ’92. But then in a study published earlier this year in *American Economic Review: Insights*, Victor and I and two of our graduate students — one from Israel, Adi Shany, and one from MIT, Jetson Leder-Luis — collected a lot more data and we reestimated the whole thing. We did not get the original finding, actually. In the newer, much larger sample, there’s not much relationship, basically none, between class size and achievement.

I can’t say that we actually figured out why it changed. Overall, classes have gotten smaller in Israel; maybe we’re into a zone where it doesn’t matter much anymore. It used to be Israeli classes were quite large, in the high 30s. Now they’re more like in America. So that’s one possibility. It could also be that the earlier estimates were kind of noisy and we got lucky, and when we had more precision, we didn’t see anything.

So the original Maimonides finding doesn’t hold up. I think there are other, more robust effects in education, like the effects of no-excuses charter schools, that have been replicated over and over by me and others. That seems more robust than the class size effect.

CHARTER SCHOOLS AND “NO EXCUSES”

EF: As you point out, you found benefits to students from attending charter schools in Boston or from attending a KIPP charter school. What were the benefits, and where do they seem to have come from?

Angrist: We’ve studied lots of charter schools. We have a research organization at MIT called SEII, the

School Effectiveness and Inequality Initiative, that does research on human capital. It's run by Parag Pathak, David Autor, and me. We have a K-12 division and a higher education division. The K-12 division has looked at a lot of school reform ideas: We've looked at charters of various types; we've looked at takeovers, where a low-performing public school is given over lock, stock, and barrel to an outside manager; we've looked at vouchers.

Charters of a particular type known as "no-excuses charters" are very effective. They're prevalent in large and urban districts, like Boston, New York, Washington, New Orleans, and Chicago. They serve low-income students, mostly minority, and many of them are organized in what are called charter management organizations, which are networks that are like franchises. KIPP is a big one.

And they have a model that seems to work. It's partly that they have a lot of inputs, so they have a long day and a long year. Actually, my daughter teaches at one, so we always have to plan our vacation around the fact that she starts teaching Aug. 1, a full month ahead of the traditional schools. She has very long workdays: Her day starts early, ends late, and in the evenings, she's on the phone with parents. Her experience is representative.

Some other things that set no-excuses charters apart are that they emphasize traditional reading and math, they have an emphasis on discipline and comportment, they have low-stakes rewards, they use data intensively, they observe the teachers quite often, and they give the teachers a lot of feedback. They also tend to target standardized tests because that's part of their accountability system.

But there are other charters that aren't nearly as good. Some of them are bad, at least as far as achievement goes. We have a paper on Massachusetts charters that shows that the suburban charters that are not in this no-excuses paradigm tend to take middle-class and upper middle-class children and reduce their achievement. It's not visible to the parents, because basically everybody does OK, but if you use the charter admissions lotteries to estimate causal effects, you get a negative effect.

EF: Does your work and the work of your colleagues tell us anything about where the gains are coming from?

Angrist: It seems to be a mix of things. The practices of the no-excuses schools are highly correlated, so you don't really get one-at-a-time randomization, even when

Joshua Angrist

► Present Positions

Ford Professor of Economics,
Massachusetts Institute of Technology
Co-Director, MIT School Effectiveness
and Inequality Initiative

► Selected Additional Affiliations

Research Associate, National Bureau of
Economic Research
Instructor, Marginal Revolution
University

► Selected Past Positions

Associate Professor of Economics,
Massachusetts Institute of Technology
(1996-1998); Associate Professor of
Economics, Hebrew University (1995-
1996); Senior Lecturer in Economics,
Hebrew University (1991-1995);
Assistant Professor of Economics,
Harvard University (1989-1991)

► Education

Ph.D. (1989), Princeton University;
B.A. (1982), Oberlin College

admissions are randomized. We do see that it's not enough just to have more instruction time. It looks like other things are important too. It isn't the level of traditional resources, though; for instance, their classes are not smaller. And it isn't peer effects. Mostly the kids in these schools come in as very low achieving.

EF: Have you looked at college completion?

Angrist: We've looked at longer-term outcomes related to college completion, but it's a work in progress because you have to wait a long time. Many of the students we study are going to middle school, then they have to go to high school, then they have to get into college. We have a paper on Boston charter effects on college enrollment and persistence. We're working on extending that.

EF: Are there problems with assessing schools' success on the basis of achievement scores?

Angrist: The main problem might be teaching to the test and that's at the expense of something else. I'm not too worried about that, because what we show in our paper is that schools that boost test scores tend to boost college. And we think college is unambiguously good.

We're not the only ones to link achievement value-added with longer-term outcomes. Raj Chetty has some work on that where it's pretty convincing that if you go to a school that boosts achievement, you're going to have higher earnings, for example. But our holy grail here is to get data on longer-term outcomes, like earnings. And I think we will.

EF: You've also found, with co-authors, that takeovers of public schools in New Orleans and Boston led to substantial gains for students. How did you determine this, and why were the schools more effective after the takeovers?

Angrist: The charter world has many variations. The most common charter model is what we call a startup — somebody decides they want to start a charter school and admits kids by lottery. But an alternative model is the takeover. Every state has an accountability system with standards that require schools to meet certain criteria. When they fail to meet these standards, they're at risk of intervention by the state. Some states, including Massachusetts, have an intervention that involves the public school essentially being taken over by an outside operator.

Boston had takeovers. And New Orleans is actually an all-charter district now, but it moved to that as individual schools were being taken over by charter operators.

That's good for research, because you can look at schools that are struggling just as much but are not taken over or are not yet taken over and use them as a counterfactual. The reason that's important is that people say kids who apply to the startups are self-selected and so they're sort of primed to gain from the charter treatment. But the way the takeover model works in Boston and New Orleans is that the outside operator inherits not only the building, but also the existing enrollment. So they can't cherry-pick applicants. What we show is that successful charter management organizations that run successful startups also succeed in takeover scenarios.

THE ELITE ILLUSION

EF: You've looked at the question of how much peers matter. Many parents obviously seek schools where they believe their children will have higher-quality peers, whatever they may mean by that term. You and your co-authors have looked at Boston and New York City selective public schools, and you concluded that peer effects don't seem to matter much. Why is that?

Angrist: I'm always beating that drum. I think people are easily fooled by peer effects. Parag, Atila Abdulkadiroglu, and I call it "the elite illusion." We made that the title of a paper. I think it's a pervasive phenomenon. You look at the Boston Latin School, or if you live in Northern Virginia, there's Thomas Jefferson High School for Science and Technology. And in New York, you have Brooklyn Tech and Bronx Science and Stuyvesant.

And so people say, "Look at those awesome children, look how well they did." Well, they wouldn't get into the selective school if they weren't awesome, but that's distinct from the question of whether there's a causal effect. When you actually drill down and do a credible comparison of students who are just above and just below the cutoff, you find out that elite performance is indeed illusory, an artifact of selection. The kids who go to those schools do well because they were already doing well when they got in, but there's no peer effect from being exposed to higher-achieving peers.

We also have papers where we show that the elite illusion is not just a phenomenon relevant for marginal kids. This is in response to an objection that goes, "If you're the last kid admitted to Stuyvesant, it's not good for you because you're not strong enough." We can refute that with some of our research designs.

EF: Are there other school situations where you think peer effects might turn out to be more important?

Angrist: There might be, but a lot of the evidence that I

find compelling goes the other way. I teach undergrad and grad econometrics, and one of my favorite examples for teaching regression is a paper by Alan Krueger and Stacy Dale that looks at the effects of going to a more selective college. It turns out that if you got into MIT or Harvard, it actually doesn't matter where you go. Alan and Stacy showed that in two very clever, well-controlled studies. And Jack Mountjoy, in a paper with Brent Hickman, just replicated that for a much larger sample. There isn't any earnings advantage from going to a more selective school once you control for the selection bias. So there's also an elite illusion at the college level, which I think is more important to upper-income families, because they're desperate for their kids to go to the top schools. So desperate, in fact, that a few commit criminal fraud to get their kids into more selective schools.

A theme of my work is that there's a lot of selection bias in simple comparisons. It's good to be skeptical of all strong empirical relationships and ask yourself and others, "Is that statistical connection really a causal effect, or is it just telling me something about the process that generates access to this school or whatever it is?" And the process that generates access to selective colleges and universities is kind of designed to mislead you. People who get into selective colleges and universities are picked because they're people who are likely to be smart and successful.

TECHNOLOGY IN THE CLASSROOM

EF: Going back to K-12, you did research with Victor Lavy on the effect of computer-aided instruction in Israel.

Angrist: Now computer-aided instruction is all the rage — personalized learning, using a lot of technology in the classroom. Victor and I had an early paper on the effects of that, taking advantage of something that happened in Israel.

EF: You found that the expected benefits didn't come to pass and that for some grades the effects were even negative. Any thoughts as to what happened?

Angrist: I'm skeptical of computer-aided instruction. I don't allow any electronics in my classroom. No laptops, no phones of course. They're a huge distraction.

That's actually been shown in a randomized trial by one of my former students, Kyle Greenberg, who's now a professor at West Point. West Point is a college like any other; it has some special features, but they teach college courses and they teach a lot of economics. Kyle and two of his colleagues did a randomized trial on allowing laptops and iPads in the classroom, and the treatment effect of that is a big negative effect.

Another example of that is laptops. The One Laptop per Child program was promoted by people in MIT's Media Lab, and they raised a lot of money for it. Eventually, the

Inter-American Development Bank figured out that they ought to make them show that it's worth doing. And the eventual randomized trial on One Laptop showed little in the way of learning gains.

EF: What started you down the road of looking so much at K-12 education?

Angrist: Well, K-12 education is economically important, but SEII also looks at higher education, which is equally important. We have a big randomized evaluation going on of financial aid, and I've done two randomized trials of financial incentives with Phil Oreopoulos.

Still, K-12 has been of particular interest to me. I think maybe in my personal case, there is a kind of peer effect in the sense that when I was in Israel, I was working with Victor Lavy and he was interested in that. And at MIT, my research direction was influenced by the arrival of Parag Pathak, who does a lot of work on market design. And Parag's thesis work was on school choice.

THE FUTURE OF ECONOMETRICS

EF: You've co-authored two econometrics texts and you teach the subject online through Marginal Revolution University. You've written that you "hope to bring undergraduate econometrics instruction out of the Stones Age." What did you mean by that?

Angrist: That comes from an article Steve Pischke and I wrote. Steve is my co-author on the books and an old friend of mine. Steve and I noticed that the way econometrics is taught is divorced from the way modern empirical work is carried out. The conventional econometrics course devotes a lot of time to things that aren't very important, like heteroskedasticity and generalized least squares, and little or even no time to questions of research design, how to figure out whether something affects something — in other words, causality — and how you should interpret regression estimates. Steve and I tackled that in our graduate book, *Mostly Harmless Econometrics*, and we followed up in *Mastering Metrics* for undergrads.

The Stones Age idea is that the Stones are an awesome band, but their heyday was the '70s. And even though modern empirical econometric research looks nothing like 1970s econometrics, econometrics instruction for the most part looks just like it looked in 1975.

EF: What do you think the relationship will be between econometrics and tools such as neural networks?

Angrist: I just wrote a paper about machine learning applications in labor economics with my former student Brigham Frandsen. Machine learning is a good example of a kind of empiricism that's running way ahead of theory.

We have a fairly negative take on it. We show that a lot of machine learning tools that are very popular now, both in economics and in the wider world of data science, don't translate well to econometric applications and that some of our stalwarts — regression and two-stage least squares — are better.

But that's an area of ongoing research, and it's rapidly evolving. There are plenty of questions there. Some of them are theoretical, and I won't be answering those questions, but some are practical: whether there's any value added from this new toolkit. So far, I'm skeptical.

EF: What are you working on now?

Angrist: We're writing up the results from our randomized evaluation of financial aid. That's been six years in the making, maybe more. I'm also doing a lot of work on value-added models for schools with a team from SEII.

Older problems where I think there's still some work to be done include the question of how to deal with weak instruments. An instrumental variable is something that's randomly assigned, or as good as randomly assigned, that provides or induces variation in something that you're interested in studying the effect of.

My interest in this originated as a result of work inspired by a 1991 paper with Alan Krueger using quarter of birth to construct instruments for schooling. In some models, we had lots of interactions and some of the interactions don't contribute much. John Bound, David Jaeger, and Regina Baker famously showed in a 1995 article that some of the resulting estimates are biased. That's the weak instruments problem. It produced a lot of work trying to address that problem, and an amazing thing is that there are still open questions there and these issues have practical consequences.

One area that I think is ripe for the picking is application of the microeconomic tools for the discovery and estimation of causal relationships to questions in macro — for instance, to the effects of monetary policy, which is a very well-defined causal question and in principle could be answered by a randomized trial.

In fact, Bob Lucas has an essay where he described how he would study monetary policy using an experiment at an amusement park near Pittsburgh. The park is called Kennywood Park, and he wrote about Kennywood because at the time he was in Pittsburgh at Carnegie Mellon. What he found interesting about Kennywood is that it had at the time, I don't know if it still does, its own currency. You didn't spend dollars in Kennywood, you spent Kennywood tickets. And so they could experiment if they wanted with their currency, they could print more of it and so on.

Lucas explained how you could use the Kennywood world to discover causal effects. I'd like to see that reasoning applied more often in real empirical work in macro. That is happening, but I think the causal empirical macro agenda should be pursued more aggressively. **EF**