Journal of Economic Perspectives—Volume 36, Number 1—Winter 2022—Pages 177–190

Isaiah Andrews, 2021 John Bates Clark Medalist

Anna Mikusheva and Jesse M. Shapiro

Isaiah Andrews is an exceptionally warm and caring person who adds to the knowledge and joy of his students, teachers, and colleagues. He is also a brilliant econometrician, and it is through his research in econometrics that he has come to be recognized with the 2021 John Bates Clark Medal. The Clark Medal is the most recent in a list of impressive accomplishments, which include a MacArthur Fellowship, a Fellowship in the Econometric Society, a Sloan Research Fellowship, a Junior Fellowship at the Harvard Society of Fellows, and others.

These and other dazzling intellectual achievements accompany a record of exceptional service to the profession. Isaiah currently serves as a co-editor of the *American Economic Review* and previously served as an associate editor at four different journals. He serves on the American Economic Association Committee on the Status of Minority Groups in the Economics Profession. He has co-organized a recent meeting of the NBER Working Group on Race and Stratification and has served on program committees for both the American Economic Association and the Econometric Society.

In addition to serving the wider profession, Isaiah has made important service contributions to his home department. After moving to Harvard in 2018, Isaiah helped to launch an Econometrics Clinic to which students in all fields can come for research advice. The clinic fits with Isaiah's strongly held view that econometric theory derives value from its ability to influence and improve empirical research on

■ *Anna Mikusheva is Associate Professor of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts. Jesse M. Shapiro is the George Gund Professor of Economics and Business Administration, Harvard University, Cambridge, Massachusetts.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.36.1.177.

Isaiah Andrews

important real-world problems. The clinic also fits with Isaiah's otherworldly talents as a teacher and communicator, able to distill a complex idea in a way that conveys its meaning to nonspecialists while still preserving many of its subtleties. It is difficult to convey how talented and special Isaiah is in this dimension; we hope that all readers have the chance someday to experience it for themselves.

Isaiah's paternal grandmother, Viola Andrews, graduated from Cheyney State Teachers College, considered the oldest of the nation's historically black colleges and universities (Cheyney University of Pennsylvania 2021). Isaiah's maternal grandparents both completed dissertations involving empirical research (Smith 1966; Smith 1971). Isaiah's mother and father received their PhDs in economics from Yale in 1984 and 1986, respectively (Smith 1984; Andrews 1986). As the child of two Yale-educated economists, Isaiah planned to avoid both Yale and economics, but eventually chose to pursue both. Isaiah received his BA in math and economics from Yale in 2009 and began his PhD studies in economics at MIT in that same year.

It was not inevitable that Isaiah would become a theoretical econometrician. For example, he might have become a theoretical macroeconomist. But at a fateful dinner at a graduate recruiting event in his first year at MIT, Isaiah sat next to Mikusheva. The seat led to a summer research assistantship. The research assistantship led to a coauthorship. And the rest is the intellectual story we are here to tell.

In the remainder of this article, we discuss Isaiah's contributions to econometric theory and empirical practice. A large portion of Isaiah's work concerns instrumental variables and related methods. Accordingly, we begin with a short review of these methods, motivated by the classic case of estimating the elasticity of demand for an agricultural commodity.

To wit, say that we want to estimate an elasticity of demand for wheat using historical data on wheat prices and quantities sold. Tempting though it may be, it is usually a bad idea to estimate the elasticity of demand via ordinary least squares regression of the log of quantity sold on the log of price (Working 1927). Because any observed price-quantity combination is an equilibrium outcome determined both by supply and demand, both price and quantity are influenced by unobserved determinants of demand, and so the ordinary least squares estimate of the demand elasticity is not generally reliable.

A possible solution is to find an *instrumental variable*, say, weather, that induces variation in prices that is unrelated to the unobserved determinants of demand, and to trace the influence of the instrumental variable, through prices, on quantity sold (Wright 1928). Intuitively, if changes in the weather shift the supply curve and are not related to unobserved determinants of demand, then changes in prices and quantity induced by changes in weather will trace out the demand curve.

In this example, the instrumental variables estimator of the elasticity of demand can be written as a ratio of the ordinary least squares estimator of the effect of the weather on the log of quantity sold to the ordinary least squares estimator of the effect of the weather on the log of price. There are two important conditions needed for the validity of this type of estimator. The first is called the *relevance condition*. It states that the instrumental variable is correlated with the variable we wish to instrument for. In our example, this means that the weather is correlated with the log of price. The second is called the *exclusion restriction*. It states that the instrumental variable is uncorrelated with unobserved determinants of the outcome. In our example, this means that the weather is uncorrelated with unobserved determinants of the log of quantity demanded.

Much of Isaiah's research concerns situations in which the relevance condition or the exclusion restriction are not guaranteed to hold convincingly, or even to hold at all, in the economic setting at hand. We discuss the two conditions in turn. In Table 1, we list a selection of Isaiah's published papers, and we refer to those papers by number in the discussion that follows.

The Relevance Condition

If the relevance condition were to fail (say, if the weather were not correlated with the log of price), then the instrumental variable would not help us recover the true value of the elasticity of demand because variation in the weather would tell us nothing about the influence of the price on the quantity demanded. Indeed, if in a particular sample the weather were totally uncorrelated with the log of the price, then the denominator of the instrumental variables estimator would be zero, and the estimator itself would be undefined!

Consider, though, what would happen if the effect of weather on the log of the price were very small but not exactly zero. In this case, the elasticity could, in principle, be estimated, but the instrumental variables estimator would be very sensitive

Table 1
Selected Publications of Isaiah Andrews

1	"Weak Identification in Maximum Likelihood: A Question of Information" (with Anna Mikusheva). 2014 *American Economic Reviews: Papers and Proceedings* 104 (5): 195–99.
2	"Maximum Likelihood Inference in Weakly Identified DSGE Models" (with Anna Mikusheva). 2015. *Quantitative Economics* 6 (1):123–52.
3	"A Geometric Approach to Nonlinear Econometric Models" (with Anna Mikusheva). 2016. *Econometrica* 84 (3): 1249–64.
4	"Conditional Inference with a Functional Nuisance Parameter" (with Anna Mikusheva). 2016. *Econometrica* 84 (4): 1571–1612.
5	"The Allocation of Future Business: Dynamic Relational Contracts with Multiple Agents" (with Daniel Barron). 2016. *American Economic Review* 106 (9): 2742–59.
6	"Conditional Linear Combination Tests for Weakly Identified Models." 2016. *Econometrica* 84 (6): 2155–82.
7	"Unbiased Instrumental Variables Estimation Under Known First-Stage Sign" (with Timothy B. Armstrong). 2017. *Quantitative Economics* 8 (2): 479–503.
8	"Measuring the Sensitivity of Parameter Estimates to Estimation Moments" (with Matthew Gentzkow and Jesse M. Shapiro). 2017. *Quarterly Journal of Economics* 132 (4): 1553–92.
9	"Valid Two-Step Identification-Robust Confidence Sets for GMM." 2018. *Review of Economics and Statistics* 100 (2): 337–48.
10	"On the Structure of IV Estimands." 2019. *Journal of Econometrics* 211 (1): 294–307.
11	"Identification of and Correction for Publication Bias" (with Maximilian Kasy). 2019. *American Economic Review* 109 (8): 2766–94.
12	"Weak Instruments in IV Regression: Theory and Practice" (with James Stock and Liyang Sun). 2019. *Annual Review of Economics* 11: 727–53.
13	"A Simple Approximation for Evaluating External Validity Bias" (with Emily Oster). 2019. *Economics Letters* 178: 58–62.
14	"On the Informativeness of Descriptive Statistics for Structural Estimates" (with Matthew Gentzkow and Jesse M. Shapiro, Matthew Gentzkow's Fisher-Schultz Lecture). 2020. *Econometrica* 88 (6): 2231–58.
15	"Transparency in Structural Research" (with Matthew Gentzkow and Jesse M. Shapiro, invited discussion paper). 2020. *Journal of Business and Economic Statistics* 38(4): 711–22.
16	"Inference After Estimation of Breaks" (with Toru Kitagawa and Adam McCloskey). Forthcoming. *Journal of Econometrics.*
17	"A Model of Scientific Communication" (with Jesse M. Shapiro). Forthcoming. *Econometrica.*
18	"Inference for Linear Conditional Moment Inequalities" (with Jonathan Roth and Ariel Pakes).
19	"Inference on Winners" (with Toru Kitagawa and Adam McCloskey).
20	"Optimal Decision Rules for Weak GMM" (with Anna Mikusheva). Forthcoming. *Econometrica*

to small variations in its denominator, because as the denominator approaches zero, the estimator's value would "explode" toward negative or positive infinity.

Figure 1 illustrates this situation in a case where more favorable weather is correlated with lower prices and greater quantity. As the effect of weather on prices

Figure 1
Sensitivity of the Instrumental Variables Estimator to the Denominator

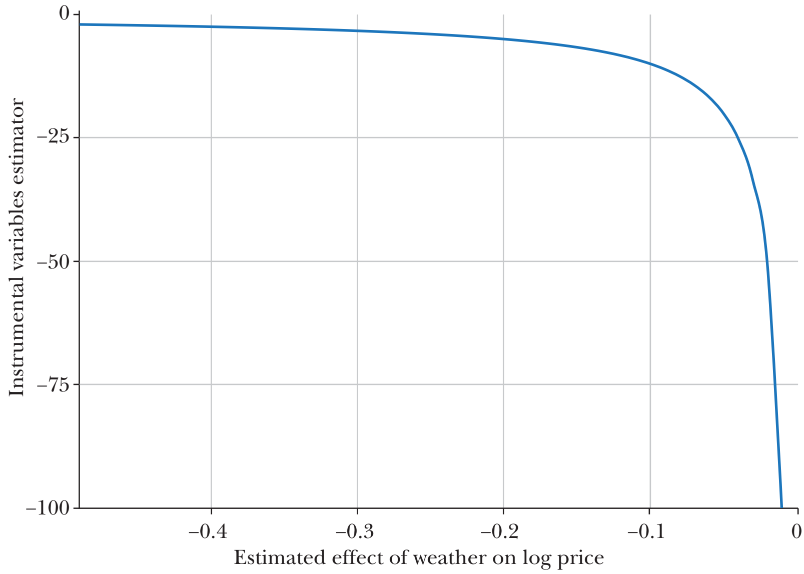

Source: Illustrative calculations by authors.
Note: The figure illustrates the instrumental variables estimator in a hypothetical case (with hypothetical numbers) where the numerator (the ordinary least squares coefficient from a regression of the log of quantity on the weather) is always 1 and we vary the value of the denominator (the ordinary least squares coefficient from a regression of the log of price on the weather). The figure shows that, as the denominator approaches zero, the estimator becomes very sensitive to small variations in the denominator.

approaches zero, the estimated elasticity more and more quickly approaches negative infinity.

The sensitivity of the instrumental variables estimator to its denominator is important because, like any statistical average, the denominator of the instrumental variables estimator is subject to random variation due to the particular sample of data (time periods in our example) included in the analysis. As Figure 1 suggests, when the denominator is close to zero, small sampling variation in the denominator can induce large, even explosive, sampling variation in the estimator.

The resulting statistical variability in the instrumental variables estimator is not captured well by the standard errors and confidence intervals that economists have often used to describe it. The reason is that such tools are often based on assumptions under which the estimator follows a bell curve or normal distribution, at least approximately. The huge sensitivity to the denominator that occurs when the denominator is small can make such approximations very poor.

The situation we are describing is a special case of a broader phenomenon called *weak instruments*, which arises when the instrument has only a small impact on the variable it is meant to affect. Interest in weak instruments among economists picked up in the 1990s when it was observed in some important research settings (Bound, Jaeger, and Baker 1995; Staiger and Stock 1997). By the time Isaiah began working on this topic in 2010, it was a mature area with a substantial body of theoretical research and practical tools.

The phenomenon of weak instruments is a special case of the more general phenomenon of *weak identification*. Intuitively, one way to think about such a situation is that the data are not sufficiently informative to clearly distinguish a single value of the parameter that is most consistent with the economic model. In our demand equation example, this manifests through the fact that when weather has very little relationship to prices, a large range of elasticities of demand will imply very little correlation between weather and unobserved determinants of demand and will therefore all be nearly consistent with the exclusion restriction.

Although Isaiah has made many important contributions to the study of weak instruments, his earliest work on the topic concerned weak identification in nonlinear models estimated via Generalized Method of Moments (GMM). GMM captures the implications of an economic theory through a set of moment conditions, which specify that some function of the data and parameters should, on average, be zero if the theory is correct. Estimation proceeds by finding the value of the parameter that makes the average value of the moment condition as close to zero as possible in the given dataset. In practice, this is often done by minimizing an objective function defined as a squared deviation of the moment condition from zero.

In our demand estimation example, a key moment condition is that the correlation between the weather and the unobserved determinants of the log of quantity is zero. One way in which instrumental variable regression is special is that this moment condition can be written as linear in the structural parameter (the elasticity of demand). If the weather is totally uncorrelated with the log of price, then the objective function is completely flat and so minimizing it does not yield a unique value of the elasticity of demand. If the weather is only weakly correlated with the log of price, then the objective function is nearly flat close to its minimum, making the solution to the minimization problem, and hence the instrumental variables estimator, very sensitive to sampling variation. Thus, in the case of the linear model, weak identification arises because there is an object (the correlation of the weather and the log of price) that must be nonzero in order for the data to pin down a unique value of the parameter of interest that minimizes the objective function. As this object approaches zero, identification becomes weak, and the data become uninformative about the parameter of interest.

But what about in nonlinear settings? The earliest work on weak identification in GMM settings attempted to generalize directly, and very literally, from the linear model. This work imagined that there is some object (analogous to the correlation of weather and the log of price) that determines whether it is possible to pin down a unique value of the structural parameters, and that the degree of "weakness" can be assessed by the proximity of this object to the value that leads to non-uniqueness (zero in the case of the correlation).

Figure 2
Illustration of Weakly Identified GMM

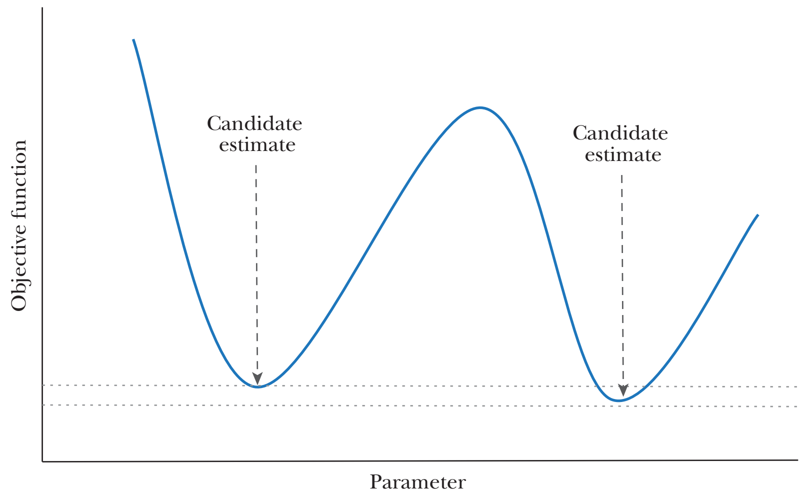

Source: Illustrative drawing by authors.
Note: The figure illustrates a hypothetical GMM objective function with two local minima. Different samples could cause either the right or the left minimum to be the lower one—and thus the parameter estimate, chosen as the minimizer of the objective, may differ substantially across sample realizations.

A central insight and contribution of Andrews and Mikusheva's work was the recognition that this analogy is incomplete. In nonlinear GMM settings, issues of weak identification can arise even if there is always a unique value of the parameter of interest that minimizes the objective function. Consider Figure 2, in which the horizontal axis shows a range of possible values for the parameter being estimated, and the vertical axis shows the value of the GMM objective function at each possible value of the parameter. The objective function exhibits a unique global minimum but two local minima. Because these different local minima achieve similar values of the objective function, small perturbations to the data can lead the global minimum to "jump" between two regions, and as a result, estimates may differ wildly across different datasets.

In addition to showing that the analogy between weak instruments and weak identification was incomplete, Andrews and Mikusheva also found that the analogy was unnecessary in the sense that it was possible to develop tools for statistical inference under weak identification without assuming a structure analogous to that under weak instruments.

In [1] and [2], Andrews and Mikusheva explore the role of weak identification in the estimation of dynamic stochastic general equilibrium models, which are one of the workhorse tools of modern macroeconomics. In these papers, Andrews and Mikusheva propose a new method of diagnosing identification weakness (based on Fisher's information matrix), and [2] proposes a method of hypothesis testing that is valid under weak identification. It is noteworthy that [2] contains an extensive discussion of practical details of implementation designed to help applied researchers use the ideas in the paper. A concern with applicability is a hallmark of much of Isaiah's research.

In [3] and [4], Andrews and Mikusheva turn their attention to broad classes of nonlinear Generalized Method of Moments (GMM) models that encompass many economic applications, such as dynamic stochastic general equilibrium models, Euler equation models, and quantile instrumental variables models. In [3], Andrews and Mikusheva develop an elegant geometric representation of the behavior of some key statistics in these models. Using this representation, the paper develops new methods for testing statistical hypotheses that are asymptotically valid even under weak identification.

In turn, [4] delves into the underpinnings of weak identification in GMM. In the case of the instrumental variables estimator, the key source of uncertainty is in the denominator, which is a single number. In the more general case, the key source of uncertainty is the shape of the entire objective function, as suggested by Figure 2. Andrews and Mikusheva show how to characterize this uncertainty as a stochastic process, which is a mathematical tool for describing a random function. Using this characterization, [4] develops a form of hypothesis test that is valid under weak identification and can be seen as generalizing earlier tools (Moreira 2003) to a broad class of nonlinear models. A challenge raised by [4] is how to make precise inferences on parameters of interest in the presence of an unknown function that determines the strength of identification. In [20], Andrews and Mikusheva suggest an approach that is based on Bayesian methods and allows a researcher to perform reliable estimation and inference in settings where information about parameters of interest is limited.

Isaiah has also made extensive contributions to the literature on weak instruments, focusing especially on important open problems relevant to empirical researchers. One such question is how to conduct hypothesis tests in situations where there are more instruments than there are unknown structural parameters. No one best hypothesis test applies in this situation, with each type of test tending to underperform in some settings. One way to handle this situation is to switch between different tests depending on features of the setting, to try to leverage the strengths of each test and avoid their weaknesses. Isaiah pursued such an approach in [6], his first sole-authored publication. Isaiah shows how to use information from one test statistic to determine how much weight to give each of two other test statistics. This leads to an approach to hypothesis testing that avoids being too conservative when identification is strong and avoids being too cavalier when identification is weak. It also clarifies some important connections among different ideas in prior work. In [9], Isaiah develops an approach to confidence interval construction suggested by the approach to hypothesis testing in [6] and includes a helpful user's guide designed to aid empirical researchers implementing the approach. Shapiro can report being a satisfied customer (Hastings, Kessler, and Shapiro 2021).

Although most of Isaiah's work on weak instruments and weak identification concerns issues of inference (such as hypothesis testing and confidence interval construction), along the way he also found time to develop an original approach to estimation in collaboration with Tim Armstrong. Recall Figure 1 and consider what would happen if we were not sure that favorable weather conditions lead to lower agricultural prices. Then the sign of the denominator of the instrumental variables estimator would be unknown. Holding fixed the numerator, a switch in the sign of the denominator would lead to even more explosive behavior of the estimator than what Figure 1 suggests, because it would lead the estimator to go from approaching negative infinity to approaching positive infinity! Fortunately, in many situations such as our wheat example, it is possible to make a reasonable economic assumption about the direction of the effect of the instrument (weather) on the variable it instruments (log price). In [7], Andrews and Armstrong show that in such situations, it is possible to construct an instrumental variables estimator that is asymptotically unbiased even when the instrument is weak.

In addition to advancing the frontiers of knowledge, Isaiah has worked to synthesize current knowledge and make it accessible to a wider audience of economists. In 2018, Isaiah and James Stock gave a series of Methods Lectures at the NBER Summer Institute devoted to weak instruments (2018). In [12], Andrews, Stock, and Sophie Sun review theoretical and practical considerations around weak instruments in linear regression settings. They find that weak instruments arise frequently in practice, showing the importance of continuing to develop tools for applied researchers. For scholars interested in understanding current best practices around weak instruments, the Methods Lectures and subsequent review article provide an important resource.

The Exclusion Restriction

When seeking to estimate the elasticity of demand, the weather is an appealing instrument for agricultural prices in part because it is unlikely to directly affect the quantity demanded—people probably do not get hungrier just because growing conditions are more or less favorable. Imagine, though, that a researcher proposes to use the prevailing agricultural wage as an instrument for the price of wheat. Like the weather, the agricultural wage influences supply conditions and would therefore seem likely to satisfy the relevance condition: all else equal, wheat will be more expensive when agricultural labor is scarce. But agricultural wages are also likely to be higher when the economy is doing well, and a strong economy may itself lead to greater demand for food. Therefore, a change in the agricultural wage may relate to the quantity sold both through an effect on the supply of wheat and through an effect on demand via consumers' disposable income. In this case, the exclusion restriction would be violated, and the instrumental variables estimator would be misleading.

Notice that while the relevance condition concerns the relationship between the instrument and an observed variable (in our example, prices), the exclusion restriction concerns the relationship between the instrument and an unobserved variable (in our example, unobserved determinants of demand). Correspondingly, while the relevance condition is usually testable, the exclusion restriction is often not testable and therefore must be justified on economic or other a priori grounds. The same applies to many other kinds of moment conditions employed in Generalized Method of Moments (GMM).

The fact that the exclusion restriction can be untestable leaves open the possibility that different researchers will disagree about the plausibility of the exclusion restriction, even in a given setting with a given instrument. For example, one researcher might find agricultural wages a very reasonable instrument for the price of wheat, believing that demand for wheat is mostly insensitive to variation in income. Another researcher might find agricultural wages a very unreasonable instrument for the price of wheat, believing that the changes in the economy signaled by a change in the agricultural wage may bring myriad effects on household demand for different goods. Although it may be possible to bring some data to bear on these questions, in many situations data alone cannot resolve such disagreements. In our experience, disagreements about exclusion restrictions are the subject of some of the liveliest exchanges in academic seminars and conferences.

Although, in general, we cannot tell whether the exclusion restriction is violated, we can often say something about how much a given violation of the exclusion restriction would be expected to distort the instrumental variables estimator. In our example, if the instrument is correlated with unobserved determinants of demand, then the instrumental variables estimator is biased, and the extent of the bias is given by the ratio of two coefficients: one the coefficient from a (hypothetical) regression of the unobserved determinants of demand on the instrument and the other the coefficient from a (feasible) regression of the log of price on the instrument.

Conley, Hansen, and Rossi (2012) characterize this bias in a more general set of linear instrumental variables models and advocate using economic intuitions about the correlation between the instrument and the unobservable to adjust inferences about the parameter of interest (in our example, the price elasticity). Conley, Hansen, and Rossi's approach applies to linear models, but, as we noted earlier, many important settings in economics call for estimation of nonlinear models. Within a couple of years, Gentzkow and Shapiro (2014) had muddled through the outlines of a generalization of Conley, Hansen, and Rossi's (2012) result to nonlinear models. Gentzkow presented this work-in-progress at an MIT seminar at which Andrews was present. These and subsequent interactions made clear that Andrews had a rather superior understanding of the issues involved—to adopt a degree of understatement suitable for these pages.[1] Andrews joined the project.

In [8], Andrews, Gentzkow, and Shapiro show how to quantify the effect of violations of an exclusion restriction, or other important economic assumptions, in a class of nonlinear economic models. The paper advocates that researchers report

[1] In what appears to have been Gentzkow's first email to Shapiro on the subject of Isaiah Andrews, Gentzkow wrote, "Did I mention that [Isaiah] is awesome?" (March 21, 2014).

a statistic called *sensitivity* which, when multiplied by a measure of the degree of violation of an exclusion restriction or other moment condition, yields a prediction of the resulting systematic error or bias in the estimator. This recommendation is shown to be practical, in the sense that in many estimation frameworks, such as GMM, sensitivity can be readily calculated, often based on objects a researcher will have already computed for other purposes.

Say, for example, that a researcher uses both the weather and agricultural wages as instruments for the price of wheat in order to estimate the elasticity of demand. By multiplying the sensitivity of the estimator by, say, a guess about the direct effect of agricultural wages on demand, a reader can arrive at an estimate of the likely bias in the estimator. In [8], Andrews, Gentzkow, and Shapiro show how to use sensitivity to quantify the effect of violations of exclusion restrictions in the context of estimating automobile demand.

While finishing [8], Andrews and Shapiro were also working on what became [17]. In that paper, Andrews and Shapiro consider a situation in which a researcher reports statistics to an audience of individuals who have different prior opinions, for example on the likely value of an unknown parameter. Andrews and Shapiro introduce a notion of the quality of a given statistical report, called *communication risk*, which captures the value to the audience of the information in the researcher's report. Andrews and Shapiro show that viewing scientific research as a communication problem can help explain many common research practices that might otherwise seem puzzling, such as reporting an estimate that violates a known sign constraint on the underlying parameter or refusing to report any estimate at all.

Ongoing work on [17] influenced the argument in [8] that reporting sensitivity is helpful to readers of a research article, because it allows them to interpret the estimator in terms of their own beliefs about the plausibility of the exclusion restriction or other important assumptions. The discussion in [8] presents this feature as an example of increasing transparency, in the sense of making the research useful even to readers who do not share the researcher's assumptions. In [15], Andrews, Gentzkow, and Shapiro formalize this connection by introducing a mathematical notion of the transparency of a statistical report based on the notion of communication risk in [17] and show a formal sense in which reporting sensitivity, as advocated in [8], can improve transparency.

A violation of the exclusion restriction distorts the instrumental variables estimator in an especially intuitive way, but the behavior of some estimators may be less intuitive. In [14], Andrews, Gentzkow, and Shapiro show that in such situations, it may be possible to offer guidance on the quantitative importance of different assumptions by connecting the estimator to other more intuitive statistics. Specifically, [14] envisions a situation in which a researcher estimates a parameter of an economic model based on a moment condition that is potentially violated. For a given degree of potential violation of the moment condition, it is possible to characterize the greatest amount of potential bias in the resulting estimator. Now suppose the researcher presents the estimates of some descriptive statistics that are connected to the parameters of interest through the economic model. If we assume that the model correctly describes the economic connection between the parameter of interest and the descriptive statistics, this reduces the potential for bias in the estimator. The extent to which the potential bias falls in this case is a measure of what Andrews and coauthors call the *informativeness* of the descriptive statistics for the parameter estimate. In [14], Andrews, Gentzkow, and Shapiro advocate that researchers report the informativeness of descriptive statistics alongside estimates of structural parameters. The framework in [14] provides a possible rationale for the common practice of reporting descriptive or summary statistics alongside formal estimates of economic parameters, a practice that is also studied (and formalized) in [15].

Like sensitivity, informativeness is often straightforward to calculate, making the suggestion to report informativeness a practical one. In [14], Andrews, Gentzkow, and Shapiro illustrate with applications. One is to Gentzkow (2007), which estimates a model of newspaper demand and approaches endogeneity using both exclusion restrictions and panel variation. Andrews, Gentzkow, and Shapiro report that in this application, a statistic related to exclusion restrictions is much less informative for a key parameter than is another statistic related to panel variation.

Sometimes researchers have multiple candidate instruments available, and when this happens, there are many different ways to combine the information from each. In [10], Andrews shows that different ways of achieving this combination can have very different properties when the exclusion restriction or other modeling assumptions are violated. This finding may have important implications for applied economists' choice of estimation methods.

A consistent theme in Isaiah's work on transparency is to try to make research more useful to its intended audience while respecting possible differences in opinions and objectives among the members of that audience. An avoidance of dogma and a positive attitude towards differences of opinion are also central to Isaiah's warm and engaging personality and to his professional identity as a researcher, teacher, and communicator. We therefore see this research agenda as an especially good match between the research subject and the researcher.

Publication Bias and Selection among Estimates

Publication bias refers to a situation in which not all estimates that are calculated are reported in the scientific literature (Kasy 2021). For example, researchers or journals may choose not to publish articles whose estimates are not statistically significantly different from zero, or whose estimates fail to support a particular conclusion. If selective publication is not accounted for, a review of the published literature can lead to misleading scientific conclusions. In [11], Andrews and Maximilian Kasy propose a way to correct for publication bias. Suppose we have a model of the publication process that tells us how the probability of publication depends on some aspect of the estimate, such as its degree of statistical significance. In this case, tools of probability theory can tell us how to account for selection and make inferences about the true range of estimates, including those that were not reported.

The challenge, then, is to uncover a valid model of the publication process. In [11], Andrews and Kasy develop an original approach to this problem that importantly avoids placing strong functional form restrictions on the model of the publication process. The approach is developed in two settings. The first setting is a replication study. Here, a researcher or collection of researchers attempt to replicate exactly a set of published studies on a new sample. If there was no publication bias, then the distribution of estimates (and other statistics such as *p*-values) would be identical between the published studies and the replications. Departures from this type of symmetry can allow us to pin down a model of the publication process from the replication data.

The second setting studied in [11] is a meta-study. Here, a researcher collects existing estimates from the literature, all estimated from different samples but aiming to uncover the same parameter. Absent publication bias, more variable estimates (say, those based on smaller samples) should have a distribution that looks like what we would get if we added noise to less variable estimates (say, those based on larger samples). Departures from this pattern can allow us to pin down a model of the publication process from the metastudy data.

Andrews and Kasy apply these methods to data from Camerer et al. (2016), who replicated a set of experimental economics papers published in the *American Economic Review* and *Quarterly Journal of Economics* between 2011 and 2014. Andrews and Kasy find that estimates significant at the 5 percent level are over 30 times more likely to be published than are estimates not meeting this level of statistical significance! Reassuringly, the correction for publication bias reaches similar conclusions regardless of whether it is based on the replication study or the metastudy.

Publication bias leads to a kind of predictable disappointment: if only the papers with the "best" estimates are published, then replications of previous studies will systematically underperform relative to the published estimates. A similar phenomenon arises in other contexts. Say, for example, that a researcher uses data on test scores to estimate the effect of teachers on students' performance. The teacher with the largest estimated effect on students' test scores may be the best teacher or they may be someone who was lucky enough to have high-performing students. In [19], Andrews, Kitagawa, and McCloskey tackle the problem of learning about the underlying effects in situations where we are interested in selecting a "best" teacher, policy, or treatment. In [16], Andrews, Kitagawa, and McCloskey apply related statistical ideas to the problem of inference after structural breaks.

Conclusion

Amazingly, the areas we discuss above do not exhaust Isaiah's contributions. He has also managed to make contributions to microeconomic theory [5], to the study of external validity bias [13], and to the problem of inference with moment inequalities [18].

Conversations with Isaiah's coauthors on these and other projects reveal what both of us know from working with him: that he is a generous collaborator who treats his research not only as a means of expanding the frontier of knowledge, but also as an opportunity to teach and learn from others. We look forward to continuing to take advantage of these opportunities whenever we can!

References

Andrews, Isaiah, and James Stock. 2018. "Weak Instruments and What to Do About Them." Lectures, NBER Summer Institute 2018 Methods Lectures, Cambridge, MA, July 22, 2018.

Andrews, Marcellus W. 1986. "Competition and Macroeconomics: Some Implications of Schumpeter's Theory of Development for Macroeconomics." PhD diss. Yale University.

Bound, John, David A. Jaeger, and Regina M. Baker. 1995. "Problems with Instrumental Variables Estimation When the Correlation between the Instruments and the Endogenous Explanatory Variable Is Weak." *Journal of the American Statistical Association* 90 (430): 443–50.

Camerer, Colin F., Anna Dreber, Eskil Forsell, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler et al. 2016. "Evaluating Replicability of Laboratory Experiments in Economics." *Science* (351) 6280: 1433–36.

Cheyney University of Pennsylvania. 2021. Cheyney University. https://cheyney.edu/ (accessed November 16, 2021).

Conley, Timothy G., Christian B. Hansen, and Peter E. Rossi. 2012. "Plausibly Exogenous." *Review of Economics and Statistics* (94) 1: 260–72.

Gentzkow, Matthew. 2007. "Valuing New Goods in a Model with Complementarity: Online Newspapers." *American Economic Review* 97 (3): 713–44.

Gentzkow, Matthew and Jesse M. Shapiro. 2014. "Measuring the Sensitivity of Parameter Estimates to Sample Statistics." NBER Working Paper 20673.

Hastings, Justine, Ryan Kessler, and Jesse M. Shapiro. 2021. "The Effect of SNAP on the Composition of Purchased Foods: Evidence and Implications." *American Economic Journal: Economic Policy* 13 (3): 277–315.

Kasy, Maximilian. 2021. "Of Forking Paths and Tied Hands: Selective Publication of Findings, and What Economists Should Do about It." *Journal of Economic Perspectives* 35 (3): 175–92.

Moreira, Marcelo J. 2003. "A Conditional Likelihood Ratio Test for Structural Models." *Econometrica* 71 (4): 1027–48.

Smith, Cheryl I. 1984. "Theories of Fluctuations in the Monetary Economy: A Study of the Concepts of Money, Capital, and Investment in the Theories of Wicksell, Hayek, and Keynes." PhD diss. Yale University.

Smith, William C. 1966. "A Factor Analytic Study of the 'Test of Sound Reasoning.'" PhD diss. University of Missouri-Kansas City.

Smith, Joyce Annunciata Davis. 1971. "The Effects of Modeling Behaviors on the Communications of Clients in Counseling Situations." PhD Diss. University of Missouri - Kansas City.

Staiger, Douglas, and James H. Stock. 1997. "Instrumental Variables Regression with Weak Instruments." *Econometrica* 65 (3): 557–86.

Working, E.J. 1927. "What Do Statistical 'Demand Curves' Show?" *Quarterly Journal of Economics* 41 (2): 212–35.

Wright, Philip G. 1928. *The Tariff on Animal and Vegetable Oils*. New York, NY: Macmillan.