Turing Award Winner, Longtime ASA Member Publishes *The Book of Why*

Judea Pearl, a longtime ASA member, was interviewed in November of 2012 (see https://bit.ly/2LdNidA*) after receiving the Turing Award from the Association of Computing Machinery. He has recently published a book,* The Book of Why: The New Science of Cause and Effect *(with Dana MacKenzie), that aims to familiarize the general, nontechnical public with recent advances in causal inference. ASA Executive Director Ron Wasserstein interviews him again here to find out what message he thinks his new book sends to* Amstat News *readers.*

The Book of Why is making a splash in statistics, as well as in machine learning and other data-intensive sciences. I would like to start with a question that you have probably heard many times: What brought you to write the book?

I have official and unofficial answers to this question.

The official answers: First, I have found it both timely and exciting to lay before the public the amazing story of a science that has changed the way we understand scientific claims and yet has remained below the radar to the general public. As we enter the era of big data and machine learning, it is important to share with the public our current understanding of how this new science is likely to affect our lives in the 21st century.

Second, as a part-time philosopher, I have found it intriguing to narrate the history of statistics as viewed from the special lens of its orphaned sister: causation. The story of this "forbidden love" was never told before and, believe me, it is full of mystery, intrigue, personalities, dogmatic orthodoxy, and heroic champions of truth and conviction.

Finally, my unofficial reason is to incite a rebellious spirit among rank-and-file statisticians, so the excitement that currently fuels causality research in academia percolates down to education and to practice. In other words, I am impatient with the slow pace at which the tools of causal inference are becoming an organic part of statistical thinking.

Judea Pearl's *The Book of Why* is a nontechnical book for the general public that discusses recent advances in causal inference.

You expressed a similar impatience in our interview six years ago. And you have initiated the ASA Causality in Statistical Education Award to close the growing gap between research and education. Hasn't this initiative met your expectations?

It has. But, with age, my impatience grew stronger and less forgiving. Of course, the availability of instructional material made it easier for instructors to introduce aspects of causal inference in graduate

courses, but it was not sufficient to change the curriculum of undergraduate classes. Nor was it sufficient to reshape the minds of practicing statisticians or high-profile academics who are too busy to sort out what all the causal inference "hype" is about.

What *The Book of Why* is doing can be described as "the democratization of causal inference." It awakens the untrained students to the realization that "it's easy and who needs the 'experts' and all their quibbles?" As a result, the book is accomplishing what I have failed to achieve in the past 30 years through hard labor and scholarly discussion with the leading statisticians of our time—a mass uprising of common sense.

I have read that some statisticians find your claims to be "hard to swallow," especially your characterization of causal inference as "The Causal Revolution" and your depiction of statisticians as antagonistic to causal thinking. Can you comment on these sentiments?

These are not only sentiments but natural complaints voiced by practicing statisticians who are genuinely surprised by how the history of statistics is viewed from the causal lens.

Take for instance the mantra "correlation does not imply causation," which every statistics student has learned to chant, demonstrate, and internalize. *The Book of Why* dissects this mantra to far-reaching conclusions that seem indeed "hard to swallow," even to seasoned statisticians.

First, it can be strengthened to assert that no causal conclusion can ever be obtained without some causal assumptions (or experiments) to support the conclusion. This is hard to swallow because it sounds circular, and because if you look at the statistical literature from 1832 to 1974, you will find many ideas about what is needed to substantiate causal conclusions (e.g., Yule, Fisher, Neyman, Hill, Cox, Cochran), but not one causal assumption—at least not formally.

This raises an interesting question: Why could not these giants of statistics come up with a simple principle, telling us what assumptions are needed for establishing a given conclusion, and let us judge—for any given situation—whether it is plausible to make those assumptions? And here comes the second surprise that is even harder for people to swallow: Even if they knew the needed assumptions, statisticians could not have articulated them mathematically—they simply did not have the language to do so.

Readers refuse to accept this linguistic deficiency until I ask them to write down a mathematical expression for the sentence, "The rooster crow does not cause the sun to rise." Failing this elementary exercise drives people to realize a totally new notational system is needed; the beautiful and powerful language of probability theory and its many extensions cannot make up for this deficiency.

The needed notation first came into being in 1920, when the geneticist Sewall Wright put down on paper a new mathematical object: a causal diagram. Thus, statistics was separated from causality, not by antagonism or disdain, but by a language barrier—the toughest barrier for humans to acknowledge and to cross. Now that the barrier is behind us, it is only natural we should call the crossing a "Causal Revolution."

These are interesting theoretical points, but I wonder if they are likely to have significant impacts on the practice of statistics or on statistical education.

The most significant practical impact of the Causal Revolution would probably be a continuous erosion of the supremacy of randomized clinical trials (RCT) in the development and evaluation of drugs, therapeutical procedures, and social and educational policies. Last year, for example, the editors of one of the two leading medical journals in America stated that authors should not talk about causation unless they have conducted a randomized clinical trial.

Miguel Hernan of Harvard and several other specialists in public health vigorously protested this restriction, and Hernan wrote, "The biggest disservice of statistics to science has been to make 'causal' into a dirty word, the C-word that researchers have learned to avoid."

Indeed, considering the practical difficulties of conducting an ideal RCT and its inherent sensitivity to sample selection bias, observational studies have a definite advantage: They interrogate the target populations at their natural habitats, not in artificial environments choreographed by experimental protocols.
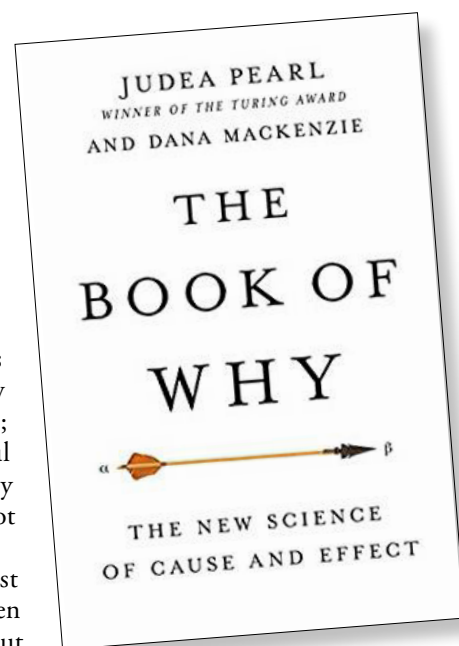
> The book is accomplishing what I have failed to achieve in the past 30 years through hard labor and scholarly discussion with the leading statisticians of our time—a mass uprising of common sense.

The development of a new toolkit that allows scientists to estimate causal effects from observational studies now opens a wide variety of applications—from medicine to social science to ecology—free from problems of ethics, costs, and external validity that plague randomized clinical trials.

True, observational studies are necessarily sensitive to modeling assumptions that must be defended on scientific grounds. However, the transparency with which those conceptual assumptions are displayed, coupled with the ability of testing them against data, now make observational studies serious contenders to RCTs.

I would like to go back to education and ask what you believe would induce a typical statistics instructor to introduce aspects of causal inference in a standard statistics class.

Curious students who read *The Book of Why* will make it impossible for statistics instructors to skip such aspects.

Take for instance Simpson's paradox, a phenomenon discussed in every statistics class, usually for the purpose of demonstrating that "correlation is not causation." The discussion usually ends with a song of praise to statistical tables for showing us that the reversal can indeed occur in the data, hence the paradox does not exist. Done. Some instructors go a bit further and praise the table for protecting us from naïve beliefs in miracle drugs that are good for men, good for women, and bad for the population.

Now imagine an inquisitive student raising his/her hand and asking the very obvious question: So, what do we do if we find Simpson's reversal in the data? Shall we believe the aggregated data or the disaggregated data? I do not believe any instructor

would in good faith be able to evade this question, suspecting the student knows the answer; it takes a few lines to describe. In other words, instructors would not be able to skip the causal implications of Simpson's paradox, as their professors did to them.

The same applies to Lord's paradox, spurious correlations, instrumental variables, confounders, and other causal concepts that were used to embarrass statistics instructors in the past.

The graphical approach you advocate in the book is but one of several approaches currently used in causal inference. Would a reader versed in potential outcome analysis feel comfortable with your methodology?

Not only comfortable, but enlightened and liberated. Researchers entrenched in potential outcome analysis will discover, to their amazement, that the following three notorious weaknesses of potential outcomes can easily be overcome:

- Assumptions of "conditional ignorability," which currently underlie every potential outcome study, can be made not because they facilitate available statistical routines, but when they are truly believed to hold in the world. They are, in fact, vividly displayed in our model of the world (i.e., the causal diagram), where they can be scrutinized for plausibility, completeness, and consistency.

- When assumptions of "conditional ignorability" do not hold, it is not the end of the world; the analysis can continue, and causal questions answered using ot her types of assumptions the model may license.

- Modeling assumptions need not remain opaque or data-blind; they can be tested for compatibility with the available data, and the model tells us how.

Making these three bullets available to researchers from the potential outcome camp will break through a wall of cultural isolation and enable them to communicate with the rest of the research community in a common, unified language.

To summarize, the democratization of causal inference is bringing about a globalization of common sense and a breakdown of cultural barriers. I am gratified to see *The Book of Why* contributing to this process. ∎