Is causal inference gendered? Examples from a male-dominated field

Candice Johnson, Michigan State University

In occupational epidemiology, we study how workplace exposures affect health. Historically, our field has focused on hazards found in male-dominated jobs, such as coal miners or construction workers, while neglecting hazards in female-dominated jobs. As a result, the methods and concepts used in our field in pursuit of causal inference were developed almost exclusively based on studies of men. In the United States, women and people of other gender identities now participate in the workforce at high rates and are of increasing interest to occupational epidemiologists. Can we use our field's traditional methods and concepts to study these populations, or are our methods and concepts gender-specific? In this presentation, I will present examples from our research in which the methods and concepts used routinely in our field did not work as expected when applied to studies of women. These results help us to better understand how gender can play an underappreciated role in causal inference.

Patients Don't Care About the Interaction Term (Should We?)

Lucy D'Agostino McGowan, Wake Forest University

Powering a trial to detect subgroup-by-treatment interactions is, in most realistic scenarios, simply not feasible. At the design stage, investigators presume clinical equipoise, that is they do not expect treatment to harm any group of patients. The overall effect is the weighted average of subgroup-specific effects and the expectation of clinical equipoise sharply constrains how different those effects can be. Under these assumptions, the sample sizes required to detect interactions are dramatically larger than those needed to detect the overall treatment effect. In a trial designed for 80% power on the overall effect, detecting an interaction of comparable magnitude would require at least four times the sample size, and even under ideal conditions, interaction test power hovers around 29%. However, power to detect a subgroup-specific effect is not the same as power to detect an interaction, and these reflect fundamentally different scientific goals. Typically, a patient does not care whether their treatment effect is statistically distinguishable from another subgroup's, they care whether there is a clinically and statistically meaningful benefit within their own subgroup. Meaningful, valid estimates of subgroup-specific effects are achievable without relying on formal interaction tests, and pre-specifying subgroups and reporting within-subgroup estimates directly is both statistically sound and more aligned with how patients and clinicians actually use trial results.

The effect of state-funded family planning policies on birth outcomes

Maria DeYoreo, RAND

State-funded family planning (SFP) programs extend reproductive health coverage — including contraception, STI screening, and preventive care — to low-income women excluded from Medicaid due to income, immigration status, or other factors. Despite their reach, evidence on the causal impact of these programs on reproductive and economic outcomes remains limited. We estimate the effect of SFP adoption on birth rates, contraceptive access, and state and federal health care spending across five states (IL, OR, SC, TX, UT) using a stacked difference-in-differences design with explicit matching on competing policy environments. We explore different inferential approaches to account for few clusters (states), and assess sensitivity to modeling assumptions. This study provides the first comprehensive multi-state assessment of SFP programs, enabling comparison of alternative program design features and informing future state investment decisions.

Causal Inference in Clinical Practice: Learning Personalized Treatment Decisions from Data

Lu Wang, University of Michigan

Personalized healthcare often requires balancing competing objectives, such as improving survival while minimizing treatment burden. Causal inference provides a principled framework for learning individualized treatment strategies from data, but in practice there is often more than one clinically acceptable option. In this talk, I introduce the concept of a tolerant dynamic treatment regime, which provides a set of feasible decision rules that achieve outcomes within a prespecified tolerance of the optimum. I will present an interpretable tree-based causal learning approach for estimating the optimal decision rules depending on the decision-maker's preferences and illustrate its application to sequential chemotherapy decisions for advanced prostate cancer. This work demonstrates how causal inference can support flexible, patient-centered decision-making in routine clinical practice.