

Commentary on Madole and Harden "Building causal knowledge in behavior genetics"
Behavioral and Brain Sciences, Volume 46, 2023

Polygene risk scores and randomized experiments

Lauren N. Ross, Kenneth S. Kendler and James F. Woodward

Abstract

We explore Madole & Harden's (2022) suggestion that single-nucleotide polymorphism (SNP)/trait correlations are analogous to randomized experiments and thus can be given a causal interpretation.

We commend Madole and Harden (hereafter MH) (this issue) for their lucid discussion of the sense in which genes or SNPs may legitimately be regarded as causes of behavioral traits. We agree with much of what they say but welcome clarification on some issues.

MH adopt a broadly “interventionist” treatment of causation—the minimal condition for some factor C to count as a cause for an outcome E is that if, hypothetically, unconfounded manipulations of C were to be performed these would lead to changes in E . In the familiar case of a randomized experiment, this leads to the conclusion that an average causal effect (ACE) is a legitimate causal notion. MH observe that an ACE can be present even though C does not have a uniform effect, even though a similar ACE may not be present in populations different from the population from which the experimental sample was drawn, and even though the experiment tells us nothing about the mechanism by which C s cause E s. We agree.

MH suggest that because of the random nature of meiosis, SNP/trait correlations from genome-wide association studies (GWAS) and/or the polygenic risk scores that incorporate these (or more precisely, such correlations among full siblings) can be likened to ACEs and hence given a causal interpretation. We explore this claim.

Consider a set of fertilized eggs immediately after conception drawn in a representative fashion from some population. Suppose this set is divided randomly into two groups, such that at a particular SNP position, one nucleotide is experimentally imposed, say A , while for the other group a different nucleotide, e.g. C , is imposed. Also suppose that the environments E are uniform across the two groups. Then, any difference in the incidence of some trait T across the two groups can be regarded as the average causal effect of having A rather than C in that population and environment

This is not an experiment that is currently technologically possible or morally acceptable. We introduce it only to provide some intuition for what a randomized experiment involving SNP manipulation that provides information about an ACE would look like. If we consider SNP/trait correlations from a GWAS, there are critical differences with the experiment just

described. Even putting aside population stratification, the random nature of meiosis does not ensure that individuals with A at some locus in comparison with those with C at that locus are causally similar in other respects (as a genuine randomized experiment does). This is because of linkage disequilibrium-- the A/C difference is very likely correlated with other causally relevant differences (often unobserved) nearby in the subjects' genomes that affect trait T. Indeed, the evidence is that most SNPs reported in a GWAS are not causal for traits of interest but are rather merely correlated with factors that are causal-- a point recognized by MH when they suggest that most SNPs have the status of "indicator" variables, tracking through correlations other factors that are causal.

Moreover, there is another, more subtle disanalogy with the randomized experiment described above. In that experiment, a single treatment--e.g., A vs C-- is randomly imposed on the population. Assuming the random nature of meiosis, a GWAS corresponds to a huge number of *different* randomized treatments in the population: e.g., A versus C at SNP1, G versus T at SNP2 and so on. An analogy would be an experiment in which a large number of different drugs D1..Dn are simultaneously randomly assigned to subjects with unknown correlations among the assignments. Indeed, matters are even more complex since haplotypes are randomized not SNPs. We might perhaps conceptualize this as the assignment of randomized bottles to subjects, each containing a mixture of different drugs. Neither of these scenarios has the straightforward causal interpretation of a standard randomized experiment.

Are these problems ameliorated if, as MH suggest, one only compares full siblings? This will help with confounds having to do with population stratification and also help, at least somewhat, with potential environmental confounds (to the extent the sibs are exposed to similar environments). However, the challenges posed by genetic linkage remain-- given a correlation between e.g. the presence of A at some SNP and trait T, we still don't know whether A is causal for T or merely correlated with some genetic factor that is causal. MH acknowledge this, suggesting that we should regard the causal factors as whole haplotype blocks.

One problem with this is that haplotype blocks are overly broad candidates for causes, in the sense that although these will contain causally relevant factors, they will also contain many more factors that are causally irrelevant, with no information about which is which. In this respect, citing a haplotype block as a cause seems analogous to saying that something unknown in my refrigerator causes an odor -- not false but not particularly informative. Moreover, we wonder whether such a causal interpretation of SNP/trait correlations is necessary. As MH suggest, one important role for such information is as a control; allowing us to see the causal role of other non-genetic (environmental) variables. Correlational information not having a straightforward causal interpretation can function as such a control as long as it is correlated with the genuinely causal confounds that need to be controlled for. A binary variable indicating whether a voter was in the U. S. south was often used as a control variable in investigations of the causal influences on voting in the mid-twentieth century. Residence in the South is not in any ordinary sense a causal variable but because it tracks or indicates genuinely causal factors (e.g., racial attitudes) that influence voting complicated ways influence voting behavior it can be used

as a control to isolate the causal role of other variable such as income. Perhaps we should think of PRSs as functioning similarly¹.

Kendler, K. and Woodward, J. (under review) "Polygene Risk Scores: A Philosophical Exploration" *Philosophy of Medicine* 4

¹ For additional discussion, see Kendler and Woodward, (2023) .