Polygene Risk Scores: A Philosophical Exploration

Kenneth Kendler
Virginia Institute for Psychiatric and Behavioral Genetics
Virginia Commonwealth University

James Woodward
History and Philosophy of Science
University of Pittsburgh

PRS: The Basic Idea

- Polygenic risk scores (PRSs) use associations between individual single nucleotide polymorphisms (SNPs) obtained from genome wide association studies (GWAS), and traits of interest, such as height, liability to a disease or educational attainment (EA).
- PRSs are constructed by regressing a trait measurement onto different variants at relevant SNPs to obtain a regression coefficient reflecting the magnitude of the association of these variables. In the simplest case, these coefficients are then summed to obtain the PRS for the trait. The score thus has the form

$S = \Sigma Bj$

where the Bj are the regression weights associated with various SNPs.

PRSs Replicate qua Predictors on Similar Samples

- The scores are first estimated on a sample drawn from a population and then “tested” by applying them to another sample, ideally from a similar ethnic group. Scores that successfully predict on the new sample are generally found to replicate in further applications on similar samples. Successful prediction is understood in terms of percent of variance explained on new samples which is a population-level measure.

- PRSs are increasingly widely used in medicine, psychiatry, genetics and social science –for disease prediction, EA etc.
- They raise a number of interesting philosophical and methodological questions.
- As noted most obviously used for predictive purposes but some authors (e. g., Harden and Madole) suggest that PRSs (or the SNP/trait correlations that go into them) can be interpreted casually or at least that they can be so interpreted in special circumstances--for example via an analogy between randomized experiments and meiosis in the case of siblings.

Questions this Raises

- What criteria need to be met for SNP/ trait correlations and PRSs to count as "causal"?
- Do SNP/trait correlations or PRSs typically satisfy such criteria?
- If instead PRSs are merely predictive (non-causal) does this mean that they have no useful role to play in causal analysis?
- Does the widespread use of PRSs signal the replacement of a kind of science that aims at causal analysis with one that aims at mere prediction?

Our Answers

- SNP/outcome correlations are not directly causal and PRSs do not directly incorporate causal information but instead summarize predictively useful but non -causal information.
- Nonetheless, although not directly causal themselves, PRSs can play an important role in causal analysis: one can sometimes leverage the associational information in PRSs to reach causal conclusions involving other variables different from those incorporated in the PRS– for example, conclusions about the causal role of environmental variables.

- This general theme– that information that is not causal can nonetheless be exploited to reach causal conclusions is interesting in its own right and not as widely appreciated as it should be in philosophical discussion

Some Background

Long before the development of modern high through-put genotyping, twin and adoption designs apparently established that many traits and disorders were substantially heritable. With the advent of modern molecular techniques, it seemed reasonable to look for “the genes” underlying this heritability. The initial expectation was that most of these genes would be in coding regions of the genome, that some of these would have relatively large effects and that much of the genetic variance associated with various traits would be captured by a small number of genes.

This expectation turned out to be incorrect. It became clear that nearly all complex traits and diseases resulted from a very large number of variants that typically had very small effects.

Claims about particular single candidate genes that causally influence complex behavioral traits like those involved in EA or in mental illness have very commonly failed to replicate. In many cases effects of any single gene even if real are so small that sample sizes in candidate gene studies are too small to identify them

- By contrast, PRSs do successfully replicate (again in the sense of percent of variance explained) when studies with adequately large samples are conducted on the same population—a central part of their appeal.

The fact that so many genetic variants of very small effects contribute to many complex traits makes the task of achieving a full “mechanistic understanding” of the genetic causation of such traits difficult. The use of PRS can be seen as a reaction to (or adaptation to) these facts. It is a strategy for dealing with what would otherwise be overwhelming complexity. Rather than attempting to directly identify individual genes or other variants that are causally responsible for various traits and to elucidate their mechanisms of action, polygenic scores bypass this in favor of a single aggregate measure that is predictively useful.

Are SNP/trait Correlations and PRSs Causal?

- A common view among users of PRSs is that the individual SNPs that go into the scores are rarely in the coding region of genes and in most cases are not direct causes of the traits they predict. Instead the SNPs are correlated in virtue of LD with variants that *are* causal. For example, two recent large-scale GWAS studies for depression and schizophrenia found that under 10% could be identified as likely causal (Levey et al, 2020, Ripke et al. 2020). A recent survey article claims flatly that SNPs “typically do not cause the [associated] trait” (Schaid et al. 2018)

Linkage disequilibrium is one way in which this can happen —if G is a genuinely causal factor for P and it is in linkage disequilibrium with a variant S1 at some SNP, then the latter will be correlated with P even though it may have no causal effect on P. In this case, the correct causal story may take the following form, with the arrows (→) representing causes :

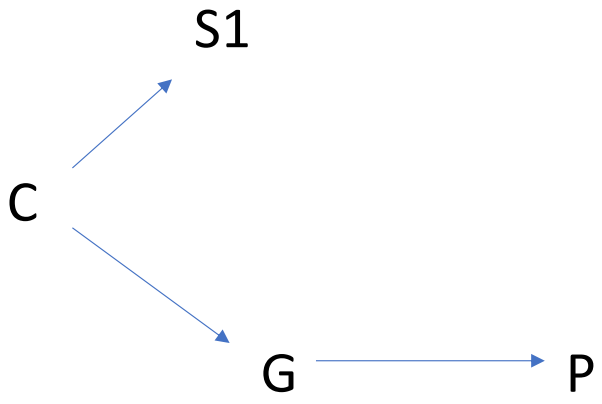

Overview of Causal Interpretation of PRSs

Some SNP trait relationship in PRSs may be genuinely causal (in a minimal interventionist sense) but many such relationships involve non-causal associations and typically we don't know which is which. Moreover, the scores do not tell us anything about the mechanisms or intervening variables by which SNPs or genes causally affect traits. Finally, at least in a number of cases, the scores are not highly stable even in a predictive sense when applied to populations that are different from those on which they were originally estimated. They are also relatively non-specific

Nonetheless PRSs Can Sometimes be Used in Causal Analysis

- These considerations may seem to suggest that PRS are of rather limited usefulness when it comes to causal analysis. We think this assessment is incomplete.

- The associational information in a PRS, although not straightforwardly causal, can sometimes be used to reach interesting causal conclusions.

Kong et al study on the “nature of nurture”

- As an illustration, Kong et al., 2018 studied the effect of the environment provided by parents and perhaps other close relatives on children’s EA. In a broad sense of genetic influence/causation which incorporates both active and reactive gene/environment correlation, children’s genotype clearly influences their EA. However, there are also important influences due to the environment provided by parents which is influenced by the parents' own genes, some of which are transmitted to their offspring and some of which are not transmitted.

- Thus the cumulative effect on children's EA reflects (1) genes transmitted from their parents, including (1a) transmitted genes the presence of which in the parents affects the environment parents provide and (2) genes of the parents that affect the environment they provide but that are not transmitted to their child (non-transmitted heredity). A standard study that looks only at the relation between parental genotypes transmitted to the child and child's educational attainment will not disentangle these factors.

- Kong et al. used a novel method to address this question. They obtained PRS related to EA for children as well as parental PRS related to EA involving alleles *not* transmitted to their children. They assumed that the direct effect *d* of the transmitted alleles on child's EA (that is, the effect not mediated by parental nurture) can be estimated via $d = \Theta(T) - \Theta(NT)$ where $\Theta(T)$ and $\Theta(NT)$ are respectively estimates of the effects of the transmitted and non-transmitted alleles. As the authors remark, calculating this difference allows one to cancel out or control for genetic nurturing effects, as well as other potential confounds. Using this methodology, the authors find that "the average estimated effect of the non-transmitted alleles [on EA] is 34.2% of that of the transmitted alleles" (425).

- Thus, there is a substantial genetically based nurture effect on children's EA that is based on non-transmitted alleles as well as on those alleles that are transmitted. By way of contrast, parental alleles associated with height and BMI predict these characteristics in children only insofar as the children actually possess those genes (cf., Koellinger and Hardin, 2018)—a result in accord commonsense causal expectations.

- This example provides an illustration of the use of PRS to disentangle, at least in part, different paths by which aggregate effects come about and thus the way in which “predictive” or not straightforwardly causal information can be used in causal analysis. Analysis does not provide detailed information about how the route that operates outside the children's bodies works (for example, it does not tell us just which factors in the environments that parents provide causally influence children’s EA) but it does tell us that there are such factors and that they are influenced by parent’s genes. This is important because it suggests the value of follow-up research to determine what those factors are and because the factors themselves may be targets for interventions that would improve EA.

- Note this disentanglement works even if the PRS for EA incorporates correlations that do not have a direct causal interpretation.
- This because PRS tracks or is correlated with factors that *are* causal for EA. Thus looking at scores associated with non transmitted EA can allow us to “see” the influence of parental environment that is not associated with genes transmitted to their offspring.

Causal Interpretation Redux

- What about the claim that conditional on parental genotype (e.g., considering pairs of full sibs) SNP/outcome relations will not be confounded because of the "random" nature of meiosis? (Harden and Madole's claim)? Here the idea is that in this case, we have a close analogy with a randomized experiment, with the causal conclusions (about average causal effects) that such an experiment licenses,
- It is true that conditioning on parental genotype will eliminate certain forms of confounding such as those due to population structure and will make others less likely (e.g., environmental confounds, assuming the sibs have similar environments)

- However, the possibility of confounding due to genetic linkage (which operates over considerably larger genomic distances than linkage disequilibrium) remains. That is, because of linkage, even if it is random whether a sib gets A or a, and random whether the sib gets B or b, it does not follow that there will be no correlation between whether the sib gets A and whether she gets B if they are even moderately nearby on a chromosomal arm. Thus despite the random nature of meiosis if, say, A/a is causal for trait P, and B/b is not, B/b can still be correlated with P because linkage is present and hence this relation can "look" causal. Conditioning on parental genotype does not solve this problem.

- Here is another disanalogy: in an ordinary RCT values of a single treatment variable are randomly assigned. By contrast when SNPS for sibs are compared, one in effect has a comparison of the combined upshot of a very large number of *different* treatments (one for each SNP) for pairs of sibs. Because of linkage one is comparing large collections of different treatments (chunks of chromosomes) rather than treatments corresponding to individual SNPs.

Thanks for listening!