Causation and SNP Heritability

Pierrick Bourrat

Macquarie University,

Department of Philosophy

North Ryde, NSW 2109, Australia

The University of Sydney,

Department of Philosophy & Charles Perkins Centre

Sydney, NSW 2006, Australia

Email: p.bourrat@gmail.com

Causation and SNP Heritability

Abstract

Genome-wide association studies (GWAS) of human complex traits have provided us with new estimates of heritability. These estimates foreground the question of genetic causation. After having presented in simple terms the rationale underlying this way of estimating heritability, I assess the extent to which relationships between genes and phenotypes established with GWAS satisfy several dimensions of causal relationships – namely range of influence, specificity, and stability – distinguished within the interventionist account of causation. The upshot is that if these relationships are causal in some sense, my analysis shows the extent to which they do not represent paradigmatic causal relationships.

Keywords: Heritability, Causation, Interventions, Specificity, Stability, GWAS, SNP

1. Introduction

Heritability is standardly defined as the ratio of genetic variance to phenotypic variance (Falconer and Mackay 1996). This statistical measure is used in a variety of contexts from breeding programs to behavioural genetics. It has long been recognized that there is not one but several concepts underlying the word 'heritability.' Perhaps the most well-known distinction is between 'broad-sense' and 'narrow-sense' heritability (Downes 2009). Broad-sense heritability is the ratio of total genetic variance to phenotypic variance, while narrow-sense heritability considers only the additive component of genetic variance to phenotypic variance. Another distinction rests on the concept of the gene underlying a particular heritability definition. In classical quantitative genetics, the genetic component typically refers to an evolutionary concept of the gene, where a 'gene' is a postulated physical basis for the phenomenological relationships between phenotypes in successive generations. In this context a 'gene' could have a number of possible physical substrates, including but not limited to DNA (Griffiths and Neumann-Held 1999; Griffiths and Stotz 2013; Lu and

Bourrat 2018). Other material substrates include transgenerational epigenetic materials which do not have a DNA-sequence basis. In the context of molecular biology – including GWAS, as we will see – it refers to a more restricted notion of the gene, such as a particular DNA sequence. In addition to the different issues surrounding the theoretical definition(s) of heritability, a gap often exists between a definition and the available methods for estimation.

Because of these different moving parts, heritability is a rich target for philosophical investigation. Previous philosophical analyses on this notion have aimed at assessing the extent to which it tracks genetic causation (e.g., Sesardic 2005; Lynch and Bourrat 2017; Bourrat in press), in particular when a gene-environment covariance or interaction exists for a trait. Prompted by new methods of estimation, other philosophical analyses have tackled a narrower issue, namely the missing heritability problem (Bourrat and Lu 2017; Bourrat, Lu, and Jablonka 2017; Matthews and Turkheimer forthcoming; Turkheimer 2011), of which one aspect is that heritability estimates obtained from GWAS are much lower than those obtained from classical family studies (e.g., twin studies, parent-offspring regression).

The aim of this paper is to evaluate in what sense and to what extent heritability estimates obtained from GWAS track genetic causation from correlations. As rightly pointed out by Turkheimer (2011), the notion of genetic causation used in the context of GWAS and heritability is often left to the intuition of the reader. I aim to clarify it. To do so, I rely on the interventionist account of causation, more particularly the analysis proposed by Woodward (2010). I start by introducing the interventionist account. I then show how this account can be applied to single nucleotide polymorphism (SNP) heritability in an ideal setting. Finally, I show to what extent the relationships identified by SNP heritability estimates depart from what represent paradigmatic causal relationships within the interventionist account. In particular, the lack of known gene to phenotype mechanisms is an obstacle for these relationships to be paradigmatic.

2. The interventionist account of causation

Following the interventionist account of causation, a variable C is a cause of another variable E, if changing the value of C by means of an ideal intervention, say from c_1 to c_2, is associated with a change in the value of E, say from e_1 to e_2 at a later point in time (see Woodward 2003; 2010). This provides a minimal criterion for causation within this account. An ideal intervention on C is defined as a change in the value of C at a given point in time which is associated with no other change in any other variable at that time.

Although this minimum criterion permits us to distinguish causes from non-causes, it provides no guidance to select the most salient cause(s) *among* different causes involved in a causal explanation. This is unsatisfactory because any given effect has infinitely many causal precedents. Yet, when one provides a causal explanation, one only cites a subset of this infinity of causes – those that are considered as the most relevant or important in some sense. The problem of choosing which of the causes should be cited in a causal explanation has been called the 'problem of causal selection' (Hesslow 1988). Within the interventionist account, the minimal criterion for causation can be supplemented by a number of 'causal dimensions' permitting us to systematically rank causal relationships according to these dimensions with the idea that, everything else being equal, the higher a causal relationship scores on any given dimension, the 'better', 'stronger', 'more salient' or 'more relevant' the causal relationship.

Following the analysis provided by Woodward (2010), the causal dimensions of 'range of influence', 'specificity', and 'stability' are all relevant to ranking causal relationships.[1] Range of

---

[1] Note that Woodward (2010) gives different names to some of the dimensions I distinguish here. Notably what I call 'range of influence', he calls the 'INF' (for influence) notion of specificity, and what I simply refer to as 'specificity', he calls 'one-to-one specificity'. I use different names for these two notions here for clarity purposes and because they have sometimes been confused (see Bourrat 2019a; 2019b). What I call 'stability' has often been lumped together with other, distinct causal dimensions. Following Pocheville, Griffiths, & Stotz (2017), I refer to one as 'stability' and the other as 'invariance' to avoid using invariance for both. I do not treat invariance in this restricted sense here. Note also that Woodward distinguishes several other dimensions which I do not treat. This is not to claim that they are not important. I simply take them to be less relevant for the question at hand.

influence concerns the extent to which one can intervene in a fine-grained way on a causal variable to produce a fine-grained effect. One example commonly used is the tuning dial of a radio, which has a larger range of influence on what we hear when compared to the on/off switch. The on/off switch has only two states leading to two effects: we hear something, or we don't, and when we hear something there is no control over what we hear. In contrast, the tuning dial determines which of many channels we hear when the radio is on.

Specificity is the extent to which there is a one-to-one mapping between cause and effect variables. Sometimes the same effect is obtained from different causal states, or the same cause produces different effects. In such cases, the specificity of the causal relationship is low. Contrasting examples here are an enzyme which catalyses a wide range of chemical reactions – low specificity – and one facilitating a single product – high specificity.

Finally, stability concerns the extent to which a causal relationship holds in different backgrounds. The higher the range of backgrounds against which a causal relationship holds the more stable it is. For instance, Huntington disease occurs in 100% of individuals when a gene on the short arm of the fourth chromosome has a particular sort of genetic abnormality (Walker 2007). Whether the disease occurs in such a situation is insensitive to the interventions made in the background of individuals with the genetic abnormality. Malignant hyperthermia is another condition under genetic control (Rosenberg et al. 2007). It leads some individuals with a mutation – generally on the long arm of chromosome 19 – to manifest, among other things, a strong fever and muscle rigidity, *only* during general anaesthesia. If left untreated, the condition can be lethal. This condition results from anaesthetic agents perturbing a physiological mechanism leading to muscle contraction and generation of heat. Individuals with the mutation are otherwise normal. The causal relationship between the mutation and the disorder is much less stable than in the case of Huntington disease, since it only occurs in a particular background, namely being exposed to some anaesthetic agents.

With some exceptions, the higher a causal relationship scores on each of the three dimensions discussed, the more explanatory the cause(s) involved in a relationship. To see that, consider that, everything else being equal, an explanation invoking a causal relationship scoring higher on range of influence, specificity, or stability, would generally be preferred over one invoking a causal relationship scoring lower on any one of these dimensions. For this reason, I refer to causal relationships scoring maximally on the three dimensions as 'paradigmatic causal relationships'. For the same reason, I refer to causal relationships scoring low on the three dimensions as 'marginal causal relationships' (see Figure 1). The former ones have a high explanatory power. The latter ones satisfy the minimal criterion, but their explanatory power is greatly diminished.
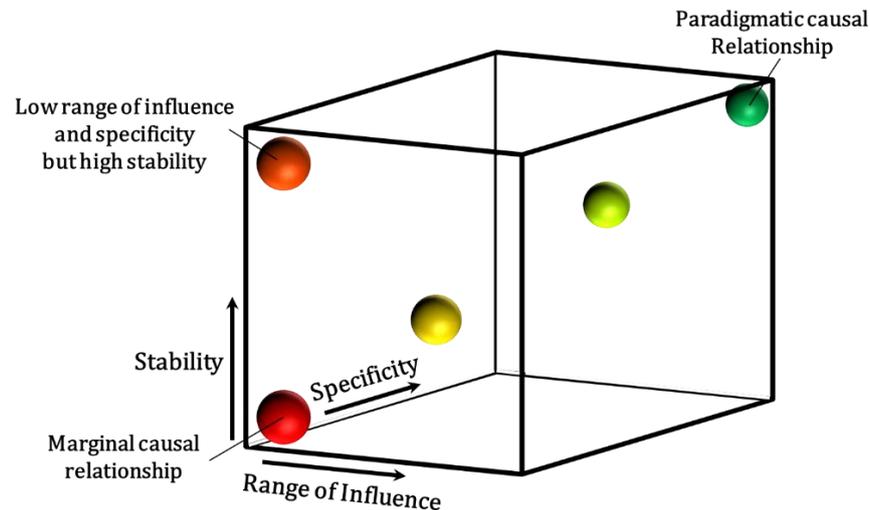

Figure 1. *Representation of three important dimensions of causal relationships for causal explanation: 'range of influence', 'specificity', and 'stability' (see main text for definitions of these dimensions). The higher a causal relationship scores on each of these dimensions, the more paradigmatic it is. In contrast the lower it scores the more marginal it is. It is a general feature of sciences to seek paradigmatic causal relationships.*

Having presented some of the most important features of the interventionist account, in the next section I apply it to an ideal case where heritability estimates obtained from GWAS could be

causally interpreted straightforwardly and would capture paradigmatic causal relationships. In Section 4, I present a number of limitations to that interpretation.

3. GWAS heritability estimates: An ideal case

At the core of GWAS are single nucleotide polymorphisms (SNPs); substitutions of a single nucleotide. When such a mutation occurs, a given nucleic base A, C, T or G, is replaced by another one in the genome. In some cases, because the genetic code has some redundancy, this type of mutation leads to no change in phenotype – the amino acid coded for in protein synthesis by the non-mutant and mutant DNA sequences are the same. Such mutations are called 'synonymous substitutions.' In other cases, a substitution will lead to phenotypic change. The aim of GWAS is to unravel from correlations the substitutions making differences to a complex phenotype – that is polygenic – such as being affected by asthma.

To do so, when a GWAS is conducted, individuals of a population are classified into two groups, namely those having the diseases and those not having it.[2] If a given SNP is more strongly associated with the disease state, one can conclude – with some confidence – that the difference in phenotype is due in part to the nucleotidic variant.[3] When this happens, the variant is said to be 'causal' (Yang et al. 2017, 1305)

Despite being purely observational studies as opposed to experiments in which interventions are performed, GWAS emulate the minimal criterion of causation, a claim made explicitly in the literature (see Lee and Chow 2013). Although no ideal interventions are performed on nucleotides in GWAS, one can assume, following the central limit theorem, that the phenotypic differences

---

[2] The method can be tweaked to accommodate quantitative phenotypes such as height. I simplify greatly here to keep only the essence of the approach (for more details see Bourrat and Lu 2017; Bush and Moore 2012).

[3] To be clear, in and of themselves, these associations are not necessarily difference makers: they might simply be located near a site which is a difference maker (these non-causal variants are said to be in "linkage disequilibrium"). There are however ways to mitigate this problem (e.g., increase sample size, perform GWAS in different populations) and, by using sophisticated statistical methods, distinguish a variant which is a difference maker from one which is merely correlated with a difference maker.

made by the genetic background and the environment of the individuals of the two groups are, on average, the same when the population is large and the factors are independent. It should be clear however that because one is dealing with correlations rather than causation *qua* ideal interventions here, despite all the precautions one might take to assess whether an SNP is causal, there is no guarantee that it truly is so. Some confounding factors might be the reason an SNP is associated with a phenotypic difference, such as when there is some population stratification (see, for instance, Mathieson and McVean 2012). Putting the problems surrounding confounding variables to one side, one can estimate the (SNP) heritability of a trait from associations between SNPs – assumed to be causal – and phenotype.

Early attempts to estimate heritability from GWAS have aimed to discover all causal variants for a given phenotype by conducting a series of GWAS and summing the variance explained by each of the different causal variants identified for a trait (Yang et al. 2017). Once this sum is divided by the total phenotypic variance, we obtain a heritability estimate. It came as a surprise to geneticists that these estimates account for less than 5% percent of the phenotypic variation for human height – the most studied human phenotype –much less than the estimates obtained by traditional family studies (around 80% for height). This observation motivated the quest for the now-famous missing heritability problem (Bourrat and Lu 2017; Bourrat, Lu, and Jablonka 2017; Maher 2008; Matthews and Turkheimer forthcoming; Turkheimer 2011). Using a method based on maximum likelihood estimation which takes into account all the information from SNPs at once rather than independently as with the early GWAS estimates, Yang et al. (2010) were able to explain about 45% of the variance in height, with the hope that new, more powerful, studies will close the remaining gap with family studies' estimates.[4]

---

[4] For reasons why this might be difficult to attain see Bourrat and Lu (2017) and Matthews & Turkheimer (forthcoming).

4. SNP heritability and causal dimensions

In the previous section, I presented an ideal situation in which heritability estimates obtained from GWAS capture the extent to which SNPs make differences to a given phenotype. I have already mentioned that the existence of confounding factors might render an estimate spurious. Besides this problem, there might be different reasons why an estimate refers to causal relationships that depart from what I called 'paradigmatic causal relationships' (see Figure 1). To refer to paradigmatic causal relationships, SNPs would have to be difference-makers of phenotypes, scoring high on range of influence, specificity and stability (and certainly other dimensions not treated here). Is that the case?

Let us start with range of influence. Because SNPs represent the smallest possible intervention-like changes one can make on the genome, this may be interpreted as a high range of influence from SNP to phenotype. In spite of this, the range of influence here is nevertheless restricted to only a subset of what is meant by 'genetic' change, in particular when 'genetic' is understood from the perspective of the evolutionary gene concept.

To see this, recall that SNPs refer to only one type of mutations, namely substitutions. However, other types of mutations occur on the genome. The most common are insertions and deletions – mutations which, as the names indicate, insert or delete a number of nucleotides on a DNA strand. We know that these types of mutations are involved in a number of diseases such as some cases of cystic fibrosis (Iannuzzi et al. 1991). In so far as SNPs concern only substitutions, the genetic variance they refer to is a *restricted* conception of the molecular gene, as opposed to a notion of the molecular gene that would include insertions and deletions (see Figure 2).
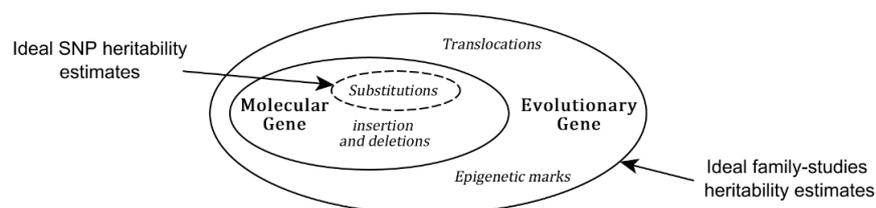
Figure 2. *SNP heritability refers to a restricted notion of the molecular concept of the gene, which is itself less general than the concept of the evolutionary gene used in classical quantitative genetics.*

Furthermore, other types of changes in the cells can have phenotypic effects. Some concern DNA itself, others do not (see Figure 2). For instance, it is well known that a particular DNA sequence being at one position rather than another on a chromosome can make a phenotypic difference. The position-effect variegation observed in *Drosophila melanogaster* in which mutants express a variegated eye pattern, as opposed to a uniformly red colour in non-mutants (see Elgin and Reuter 2013 for details on this fascinating case) is one such example. The cause for the phenotypic difference is due to a chromosomal translocation (more particularly an inversion) on the X chromosome. Because of the change in position, the gene responsible for eye colour is located in the heterochromatin – a condensed part of nuclear DNA difficult to access for RNA polymerases – in mutant individuals, as opposed to the euchromatin – less condensed and thus more easily accessible for RNA polymerases – in wild individuals. This change in position explains why the gene is expressed in some cells but not others in mutant individuals: when polymerases reach the gene, it is expressed, otherwise it is not. This difference is 'genetic' in the sense that it is heritable and thus corresponds to the evolutionary notion of the gene, but not genetic from the perspective of the DNA sequence underlying the character since this sequence is the same (but with an inversion) in mutant and wild individuals. Causal variants identified in GWAS do not capture these factors.[5]

These different cases indicate that if the causal relationships between SNPs and phenotype, and consequently their aggregation and normalisation into SNP heritability estimates, satisfy – or more accurately emulate – the minimal criterion of causation, the type of 'genetic' causal

---

[5] However, note that some variants might be *correlated* with non-substitution mutations or non-DNA changes and, as such, be indicators of genetic causation in an evolutionary sense and prompt new studies to identify causal factor(s).

relationships they capture have restricted ranges of influence. They refer to a narrower notion of the gene than that referred to in classical quantitative genetics. The extent to which these relationships have restricted ranges of influence would, however, be difficult to estimate and is ultimately an empirical question that will probably lead to a different answer for different phenotypes.

Having shown in what sense SNP heritability departs from a paradigmatic genetic causal relationship with respect to range of influence, I now move on to specificity. We saw that SNPs can be conceptualized as emulations of idealized interventions. Substituting one nucleotide for another by means of an ideal intervention in one individual (at the zygote stage), would, following the interventionist account, make a phenotypic difference, if this nucleotide is causal. But one might legitimately ask whether in general the focal phenotypic difference made would be the only one, i.e. whether SNP-to-phenotype causal relationships are specific. Were those relationships unspecific, claims about SNPs being important causes of phenotype would be somewhat undermined.

The example of the gene *FOXP2* is a case in point. This gene has been presented as a gene for language in popular media. This is because a number of mutations on this gene – located on the 7th chromosome in humans – are associated with language impairment (for a review including the history of the discovery of this gene see Nudel and Newbury 2013). Yet, one might reasonably ask whether speech impairment is only one of the outcomes of a number of other phenotypes. For instance, it is well known that very precise motor coordination is necessary to be able to speak. Research on the *FOXP2* gene shows that mutations on this gene can lead to difficulties with fine motor coordination such as individuals having difficulties tying their shoelaces or buttoning their clothes (Morgan et al. 1993). It is also clear that *FOXP2* is involved in the development of many parts of the brain, as well as the development of the lungs and gut. Knockout of the gene in mice leads to reduced life expectancy (Vargha-Khadem et al. 2005).

It is also important to recognize that without a mechanism explaining the mode of action of differences made by a nucleotide change, assessing the degree of causal specificity of causal relationships from SNPs to phenotypes will be undermined.[6] This is however mitigated by measures of genetic correlation between different traits which permit estimates of the pleiotropy for a given genetic variant – that is the extent to which a gene influences more than one trait, a form of causal specificity. Based on such estimates, Visscher et al. (2017) report that pleiotropy is pervasive for many human complex traits. Finally, and related to that point, the extent to which a relationship from SNP to phenotype is specific will, in many cases, depend on the way a phenotype is defined. A complex trait such as hypercholesterolemia would perhaps appear as a number of different traits – different subtypes of hypercholesterolemia – if one knew all the different mechanisms involved in the development of the condition.[7]

In sum, even though an SNP might be causally linked to some of the variation in a particular phenotype, it might only be so very non-specifically. Without details of the molecular mechanisms underlying the expression of a phenotype however, it would be hard to assess the extent to which the variation results from a specific causal relationship.

Moving on to stability, it should be stressed that all associations between SNPs and phenotype happen in particular populations. Why does that matter? Simply because establishing a causal relationship between an SNP and a phenotype in one population does not mean that this relationship would hold in a different background – especially in the absence of knowledge of underlying mechanisms. It could very well be the case that the same variant in one population has the opposite effect or no effect in a different population.[8] Recall the case of malignant hyperthermia. The disorder occurs as a result of a single mutation and in very particular

---

[6] The interventionist account has some resources to assess whether a causal relationship refers to a mechanism. See for instance Woodward (2002).

[7] This point is also related to the notion of proportionality of causal explanation, see Woodward (2010) and Pocheville et al. (2017).

[8] This is a version of the famous 'locality problem' (Bourrat in press; Lewontin 1974; Sesardic 2005).

circumstances. From this example, it is possible to imagine that many of the causal variants influence complex phenotypes only under specific conditions (e.g., diet, climate, cultural practices) varying between different populations. Were this to be the case, SNP heritability estimates would refer to unstable causal relationships (see Figure 1). The possibility of such phenomena is recognized in the literature, and it is hoped that studies of populations of non-European descent will soon be available (Visscher et al. 2017) that permit us to assess the stability of the SNP heritability estimates.

5. Conclusion: Where are we at?

I have shown that applying the interventionist account of causation to the notion of SNP heritability can help us understand some of the challenges surrounding the claim that SNP heritability captures causal relationships between genes and phenotypes. More particularly, if (strictly speaking) it is true that causal variants represent loci of genetic causation for phenotypes (insofar as they emulate ideal interventions), they generally do not correspond to the full range of influence classically associated with genetic causation. It is also doubtful that the causal relationships identified in GWAS have high specificities in general – although there might be some exceptions – especially in the absence of known mechanisms underlying the relationships (see Matthews and Turkheimer forthcoming). Finally, it is unclear whether all the associations found in populations of European descent will remain once samples from a wider range of populations become available. Having framed some of the problems surrounding the notion of heritability from the perspective of the interventionist account, I expect the ins and outs of other problems surrounding this concept – especially those pertaining to causation in the context of the advent of new biotechnologies – to be more easily assessed thanks to this framing.

Acknowledgments

References

Bourrat, Pierrick. in press. 'Heritability, Causal Influence and Locality'. *Synthese*.
———. 2019a. 'On Calcott's Permissive and Instructive Cause Distinction'. *Biology & Philosophy* 34 (1): 1. https://doi.org/10.1007/s10539-018-9654-y.
———. 2019b. 'Variation of Information as a Measure of One-to-One Causal Specificity'. *European Journal for Philosophy of Science* 9 (1): 11. https://doi.org/10.1007/s13194-018-0224-6.
Bourrat, Pierrick, and Qiaoying Lu. 2017. 'Dissolving the Missing Heritability Problem'. *Philosophy of Science* 84 (5): 1055–67.
Bourrat, Pierrick, Qiaoying Lu, and Eva Jablonka. 2017. 'Why the Missing Heritability Might Not Be in the DNA'. *BioEssays* 39 (7).
Bush, William S., and Jason H. Moore. 2012. 'Chapter 11: Genome-Wide Association Studies'. *PLOS Computational Biology* 8 (12): e1002822. https://doi.org/10.1371/journal.pcbi.1002822.
Downes, S. M. 2009. 'Heritability'. *Stanford Encyclopeadia of Philosophy*. https://plato.stanford.edu/entries/heredity/.
Elgin, Sarah C.R., and Gunter Reuter. 2013. 'Position-Effect Variegation, Heterochromatin Formation, and Gene Silencing in Drosophila'. *Cold Spring Harbor Perspectives in Biology* 5 (8).
Falconer, Douglas S., and Trudy FC Mackay. 1996. *Introduction to Quantitative Genetics (4th Edn)*. Essex: Longman.
Griffiths, P. E., and E. M. Neumann-Held. 1999. 'The Many Faces of the Gene'. *Bioscience* 49: 656–62.
Griffiths, P. E., and K. Stotz. 2013. *Genetics and Philosophy: An Introduction*. New York: Cambridge University Press.
Hesslow, Germund. 1988. 'The Problem of Causal Selection'. *Contemporary Science and Natural Explanation: Commonsense Conceptions of Causality*, 11–32.
Iannuzzi, Michael C., Robert C. Stern, Francis S. Collins, Catherine Tom Hon, Noriko Hidaka, Theresa Strong, Lisa Becker, et al. 1991. 'Two Frameshift Mutations in the Cystic Fibrosis Gene'. *American Journal of Human Genetics* 48 (2): 227–31.
Lee, James J., and Carson C. Chow. 2013. 'The Causal Meaning of Fisher's Average Effect'. *Genetics Research* 95 (0): 89–109. https://doi.org/10.1017/S0016672313000074.

Lewontin, Richard C. 1974. 'The Analysis of Variance and the Analysis of Causes'. *American Journal of Human Genetics* 26: 400.

Lu, Qiaoying, and Pierrick Bourrat. 2018. 'The Evolutionary Gene and the Extended Evolutionary Synthesis'. *The British Journal for the Philosophy of Science* 69 (3): 775–800. https://doi.org/10.1093/bjps/axw035.

Lynch, Kate E., and Pierrick Bourrat. 2017. 'Interpreting Heritability Causally'. *Philosophy of Science* 84 (1): 14–34. https://doi.org/10.1086/688933.

Maher, Brendan. 2008. 'Personal Genomes: The Case of the Missing Heritability'. *Nature News* 456 (7218): 18–21.

Mathieson, Iain, and Gil McVean. 2012. 'Differential Confounding of Rare and Common Variants in Spatially Structured Populations'. *Nature Genetics* 44 (3): 243–46. https://doi.org/10.1038/ng.1074.

Matthews, Lucas, J., and Eric Turkheimer. forthcoming. 'Across the Great Divide: Pluralism and the Hunt for Missing Heritability'. *Synthese.*

Morgan, Angela, Simon E. Fisher, Ingrid Scheffer, and Michael Hildebrand. 1993. 'FOXP2-Related Speech and Language Disorders'. In *GeneReviews®*, edited by Margaret P. Adam, Holly H. Ardinger, Roberta A. Pagon, Stephanie E. Wallace, Lora JH Bean, Karen Stephens, and Anne Amemiya. Seattle (WA): University of Washington, Seattle. http://www.ncbi.nlm.nih.gov/books/NBK368474/.

Nudel, Ron, and Dianne F Newbury. 2013. 'FOXP2'. *Wiley Interdisciplinary Reviews. Cognitive Science* 4 (5): 547–60. https://doi.org/10.1002/wcs.1247.

Pocheville, Arnaud, Paul E. Griffiths, and Karola Stotz. 2017. 'Comparing Causes – an Information-Theoretic Approach to Specificity, Proportionality and Stability'. In *Proceedings of the 15th Congress of Logic, Methodology and Philosophy of Science*, edited by Hannes Leitgeb, Ilkka Niiniluoto, Elliott Sober, and Päivi Seppälä, 260–75. London: College Publications.

Rosenberg, Henry, Mark Davis, Danielle James, Neil Pollock, and Kathryn Stowell. 2007. 'Malignant Hyperthermia'. *Orphanet Journal of Rare Diseases* 2 (April): 21. https://doi.org/10.1186/1750-1172-2-21.

Sesardic, Neven. 2005. *Making Sense of Heritability.* Cambridge: Cambridge University Press.

Turkheimer, Eric. 2011. 'Still Missing'. *Research in Human Development* 8 (3–4): 227–241.

Vargha-Khadem, Faraneh, David G. Gadian, Andrew Copp, and Mortimer Mishkin. 2005. '*FOXP2* and the Neuroanatomy of Speech and Language'. *Nature Reviews Neuroscience* 6 (2): 131–38. https://doi.org/10.1038/nrn1605.

Visscher, Peter M., Naomi R. Wray, Qian Zhang, Pamela Sklar, Mark I. McCarthy, Matthew A. Brown, and Jian Yang. 2017. '10 Years of GWAS Discovery: Biology, Function, and Translation'. *American Journal of Human Genetics* 101 (1): 5–22. https://doi.org/10.1016/j.ajhg.2017.06.005.

Walker, Francis O. 2007. 'Huntington's Disease'. *The Lancet* 369 (9557): 218–28. https://doi.org/10.1016/S0140-6736(07)60111-1.

Woodward, James. 2002. 'What Is a Mechanism? A Counterfactual Account'. *Philosophy of Science* 69 (S3): S366–77. https://doi.org/10.1086/341859.

———. 2003. *Making Things Happen: A Theory of Causal Explanation.* New York: Oxford University Press.

———. 2010. 'Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation'. *Biology & Philosophy* 25 (3): 287–318.

Yang, Jian, Beben Benyamin, Brian P. McEvoy, Scott Gordon, Anjali K. Henders, Dale R. Nyholt, Pamela A. Madden, et al. 2010. 'Common SNPs Explain a Large Proportion of the Heritability for Human Height'. *Nature Genetics* 42 (7): 565–569.

Yang, Jian, Jian Zeng, Michael E. Goddard, Naomi R. Wray, and Peter M. Visscher. 2017. 'Concepts, Estimation and Interpretation of SNP-Based Heritability'. *Nature Genetics* 49 (9): 1304–10. https://doi.org/10.1038/ng.3941.