Indirect Genetic Causes

Kate E. Lynch
Department of Philosophy and Charles Perkins Centre
The University of Sydney
NSW, 2006
Australia
kate.lynch@sydney.edu.au

Indirect genetic effects occur when genes exert their influence on traits via the environment. Sometimes this occurs when individuals modify their own circumstances, other times people when impose circumstances on others. In both cases genes indirectly lead to phenotypic differences, yet causal interpretation of the two types of cases differs. This paper suggests a reason for this difference: environmental confounding. This point of difference depends upon the way causal variables are described. Invoking proportionality to determine appropriate variable description maintains the difference. The dependence of variable description helps to illuminate the relationship between gene-environment correlation and gene-environment interaction.

Abbreviations
GEI: Gene-environment interaction
GEC: Gene-environment correlation
GWAS: Genome-wide association study
SNP: Single-nucleotide polymorphism

1. Introduction

Causality is an important topic in human behavior genetics. Studies reveal relationships between genetic variation and differences in human behaviors, abilities and achievements. Philosophers and scientists continue to debate whether the current available evidence is sufficient for genetic causal inference (Lewontin 1974; Sesardic 2005; Plomin 2018; Kaplan and Turkheimer 2021; Madole and Harden 2022; Burt 2022). Part of this discussion concerns whether genes influence outcomes *in the right way*. Examples of gene-environment interplay challenge intuitions about genetic causal attribution. In these cases, *how* genes exert their effects seem to be important to whether they should be counted as meaningful causes of human difference.

This paper addresses these cases by providing an analysis of gene-environment interplay within causal-modelling framework. The paper has two aims. The first is to give a principled account of different types of gene-environment interplay. This extends the work of Lynch (2017) and Lynch and Bourrat (2017), taking into consideration a broader set of causal information, producing different causal conclusions. These conclusions better cohere with intuitions about problem cases, and with scientific practice within the behavior genetics community. The second aim is to illustrate the relationship between different kinds of gene-environment interplay. Gene-environment interaction (GEI) and gene-environment correlation (GEC) have traditionally been considered distinct phenomena. In this paper I show that they can be made conceptually equivalent given a particular understanding of 'environment'. This demonstrates the dependency of this distinction on

variable description, and illustrates an implicit, but generally accepted, definition of environment used within behavior genetics.

The paper proceeds as follows. In §2, I describe different kinds of gene-environment interplay within a behavior-genetics framework, and survey interpretations from the literature (§3). In §4 I propose a new way to causally model these cases which illuminates a distinction in line with causal intuitions raised in §3. In §5 I consider an alternative to this approach which highlights the relationship between GEI and GEC. In light of this, §6 argues for the approach suggested in §4 by appealing to additional criteria for variable description. I conclude with an account of how to causally interpret cases of gene-environment interplay that better coheres with causal intuitions and scientific practice.

2. Gene-Environment Interplay

Genes and environment act in combination throughout development to produce phenotypic effects. These developmental interactions are distinct from gene-environment interplays, whereby genetic and environmental influences interact at the population level to produce phenotypic variance that cannot be separately attributed to genetic or environmental differences. To understand the implications for behavior genetics, one must first understand the assumptions underlying behavioral genetics methods.

2.1 Heritability

Heritability studies have traditionally used family resemblances to infer the relative influences of genes and environment on phenotypic differences[1]. For example, height has a relatively high heritability. Differences in height are observed between those with different genetic backgrounds, and similar phenotypes are found between those with similar genetics (such as family groups). Conversely, language(s) spoken, which is largely attributable to environmental differences, has a low heritability, because environmental variation is primarily responsible for population level differences. To estimate heritability genes are not studied directly but are inferred from familial relatedness.

This contrasts with genomic methods like genome-wide association studies (GWAS) which infer genetic sources of phenotypic difference by correlating large numbers of genetic variants arrived at via sequencing. GWAS study a large set of single nucleotide polymorphisms (SNPs): nucleotides which have more than one variant (termed an allele) present in the population, at some common frequency (usually >1%). It is important to note that SNPs are not genetic variants with known phenotypic outcomes or molecular pathways. They are genetic markers which may be part of a gene which directly influences the trait of interest, may exert their influence via regulating other regions of the genome, or may simply be associated with a difference making gene. Crucially, family-based studies or GWAS study genetic variants with mechanisms known to be related to differences in

[1] These methods are also used in non- human animal populations, in this paper I focus on human populations.

phenotype. GWAS employ a linear regression model to associate SNP differences with phenotypic differences in large populations. Millions of SNPs are treated as independent variables to see which significantly associate with trait differences[2]. Like family-based methods, GWAS can be used to estimate heritability: the degree to which genetic variation accounts for variation in phenotype.

To estimate heritability (H^2) the relative contributions of environmental and genetic differences are encompassed by an environmental variance term (V_E) and a genetic variance term (V_G) respectively, which sum to produce the total phenotypic variance in a population (V_P). Heritability is the proportion of total variance that can be accounted for by genetic variance.

(1) $V_P = V_G + V_E$

(2) $H^2 = V_G/V_P$

Both family-based and GWAS methods assume the 'additive' heritability model (equation 1), where genetic and environmental variance influence the phenotypic variance independently. There are two ways in which this assumption can be violated, which lead to heritability estimates which contradict intuitive judgements about genetic causation. I turn to these below.

[2] Individually, common SNPs account for a miniscule amount of variance. When significant effects are combined, some phenotypic variance can be accounted for. Including non-significant SNPs explains slightly more. Even then, the estimate of genetic effects is much smaller than predicted by family studies. This has been termed the 'missing heritability' problem.

2.2 Gene-Environment Interactions

When genes and environments act additively the effect of the environment influences a trait to the same extent (on average), no matter the genetic background. For instance, children receiving poor post-natal nutrition will have slower growth rates than those receiving better quality nutrition, not matter the genetic background of the child. The additivity assumption is violated in cases of gene-environment interaction (GEI), when environments have differential modifying effects depending on the genetic background of the individual. The classic, though controversial, case in the behavior genetics literature is GEI for antisocial behavior. According to Caspi et al. (2002), individuals with a low monoamine-oxidase (*MAOA*) genotype are more likely to present with antisocial behavior compared to high *MAOA* individuals, when raised within an environment of childhood maltreatment. However, the opposite trend is true for individuals raised in environments without childhood maltreatment. Thus, the direction of the environmental effect changes depending on an individuals' genetics. A schematic of this result is represented in Figure 1.

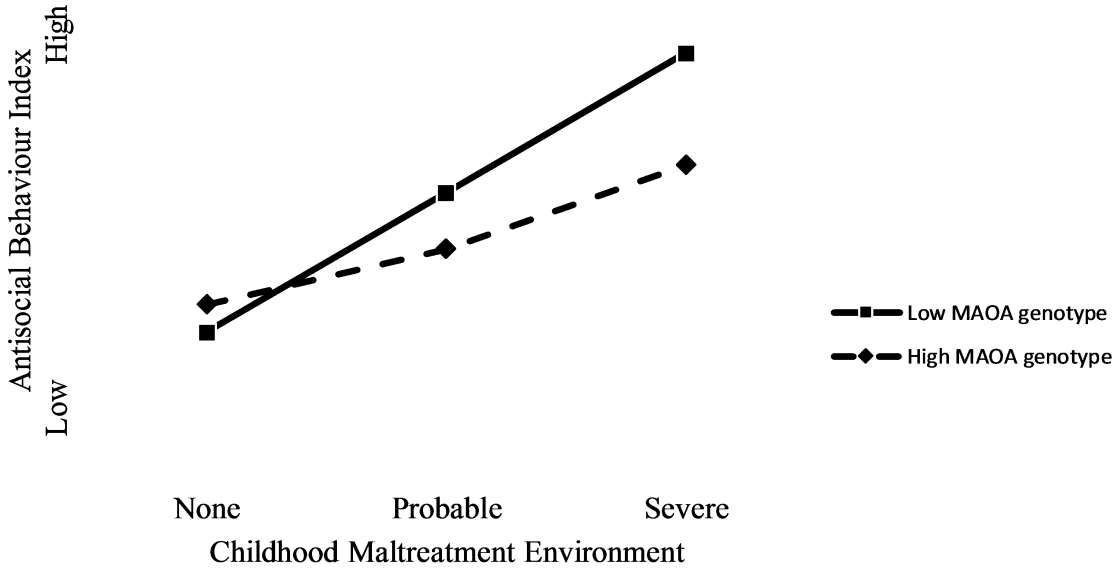

Figure 1. The interaction of the MAOA genotype and childhood maltreatment on antisocial behavior.

Phenotypic variance produced from interactions can be accounted for statistically, adding an extra term (V_{GxE}) to equation (1). However, even when such adjustments are made, GEI highlights the possibility of non-additivity for untested genotype-environment combinations. There could an untested environment in which genotypes develop in discordance to a usual trend, such as a genotype in which fetal undernutrition produces an increase in growth rate. This problem was first raised by Lewontin (1974) and has been termed the locality objection (Sesardic 2005). The objection points to a limitation of the explanatory power of heritability estimates beyond the population under study. Lewontin

and his supporters believe that this limitation undermines the utility of heritability estimates for tracking genetic causation.

GEI is recognized as theoretically troublesome for estimating heritability (Tabery 2014). However, the degree to which GEI makes a difference in human populations has been seriously questioned. The results of Caspi et al. have failed to replicate in more recent studies and meta-analyses (Figlio et al. 2017). Estimates of GEI using GWAS have also been attempted and found no significant contribution to phenotypic variance (Allegrini et al. 2020). I will return to a discussion of GEI in §5.

2.3 Gene-Environment Correlations

A second violation of the additive heritability model is gene-environment correlation, also termed gene-environment covariance (GEC). In these cases, genes correlate with differences in the external environment. The heritability estimates attained in these situations can lead to conflicting intuitions about genetic causation, again leading some to call for a rejection of heritability for tracking genetic causation (Block 1995; Jencks 1972; Sober 2001).

GEC are generally split into three different types: passive, active, and reactive (also called 'evocative') (Plomin et al. 1977; Scarr and McCartney 1983). Passive GEC occurs when children 'inherit' environments from their parents that are correlated with their genetics. For example, children who inherit genes related to sporting ability are also more likely to be raised in environments where sporting activity is encouraged and coached. At a population level, these genes and environments are correlated, making it difficult to

147 separate the effects of genes and environment. The causal interpretation of passive GEC is
148 uncontroversial, and the separation of the two sources of variance can be achieved using a
149 variety of experimental designs. I direct the reader to Lynch (2017) and Lynch and Bourrat
150 (2017) for an exposition of the causal structure of these cases. This paper focuses on the
151 two more controversial cases of GEC and their implications for genetic causal claims:
152 reactive and active GEC.

153 When reactive GEC occurs, the environment of an individual is imposed upon them by
154 others as a result some genetically caused phenotypic difference. The best-known
155 (hypothetical) example is thanks to Jencks (1972, 66-67) [3]:

156 > If, for example, a nation refuses to send children with red hair to school, the genes that
157 > cause red hair can be said to lower reading scores... If an individual's genotype affects
158 > his environment, for whatever rational or irrational reason, and if this in turn affects his
159 > cognitive development, conventional methods of estimating heritability attribute the
160 > entire effect to genes and none to environment.

161 Assuming that red-hair has a genetic basis, then one genotypic group (red-haired children)
162 is subjected to a different schooling environment than another (say, blonde-haired children)
163 – correlating genotypes and environments at the population level. Underlying the

[3] This is a toy example which is still being discussed today (e.g., Matthewson & Griffiths 2017, Kaplan & Tukheimer, 2021; Harden 2021; Burt 2022). Actual family-based and GWA studies examine the effects of many different genes on a heterogenous population rather comparing two groups based on a single feature. Keeping with this example allows me to cleanly illustrate the relevant causal dynamics.

correlation is a causal process going from genes to reading ability, indirectly via hair color and educational environment.

Though Jenck's example is hypothetical, it represents the varied ways in which prejudices occur against individuals with genetically based phenotypes like skin color, facial features, or physical manifestations of disability. In these cases, it seems obvious that prejudice (an environmental factor) is a cause of discriminatory outcomes, however a strictly causal interpretation will also include genes as an indirect cause.

Figure 2 represents the causal structure of these cases using a directed-acyclic-graph (DAG). DAGs represent causal models, which can be used to illustrate different types of causal structures. This technique has been particularly influential when used in concert with difference making accounts of causation such as the interventionist account (Woodward 2003). Under an interventionist account, a cause (C) causes an effect (E) if an intervention on C, changing the nature of C, produces a change in the nature of E (Woodward 2003). This is most easily understood when the causal relata are variables, and changes result in a variable taking a different value. For instance, for two variables, C and E, C has the possible values c_1 and c_2 and E the values e_1 and e_2. For C to be a cause of E, an intervention on C, changing its value from c_1 to c_2, results in a change in the value of E, from e_1 to e_2. Thus, the expression 'C causes E' is true if an intervention on C changes the value of E. An intervention is a particular kind of manipulation in which the change of value for variable C occurs without changing the value of any other variable, Z, which could itself be a cause of E. In this way the change in the value of E is fully accounted for by the change of the value of C. For a relationship to be causal all that is required is that

some intervention on C results in a change in E. There may be some interventions which result in no change. Finally, for a relationship to be causal no actual intervention need be performed, or even be possible. The account captures causal relations as those where, if an intervention on the cause were to occur, there would be a resulting change in the effect. (Woodward 2003).

Representing relata as variables is useful in scientific contexts as scientists routinely examine how changes to variables, often induced via manipulations, produce changes in effect(s). Scientific investigation often also includes controlling for other background variables that might confound the causal relationship of interest (Z above). The importance of both manipulation and controlling for confounds id further discussed in §4 and §5.

In DAGs causal relata are represented as variables and the causal relationship represented by directed arrows going from causes to effects. A directed path is represented by a sequence of arrows pointing in the same direction, indicating the direction of causation. The remainder of this paper will examine gene-environment interplays using a causal framework in the spirit of Woodward (2003). DAGs will be employed to illustrate the causal structure of cases and their changes under interventions on variables in the system.

Figure 2 represents active GEC at the level of the population, where variables represent variance in causes and effects.

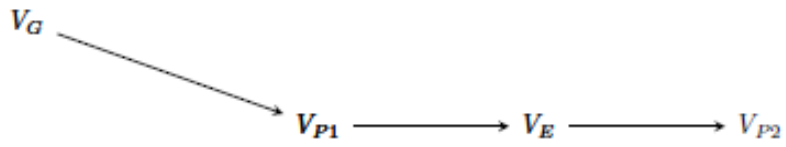

Figure 2. Reactive GEC V_G (genetic variance) causes V_{P2} (variance in reading ability) via variance in hair color (V_{P1}) causing variance in educational resources (V_E).

Active GEC occurs when an individuals' genotype causes them to modify the environment of their own accord, with resulting phenotypic effects. For example, imagine that blonde-haired children were more genetically predisposed to seek out enriching educational environments for themselves by way of books and extra study. Red-haired children, by comparison, lacked such motivation, exposing themselves to less educational resources. The two genotype groups (identified by different correlated physiology) are also subject to different educational environments. But now the environmental differences are due to the children themselves, rather than an imposition from others. A similar causal process also underlies the resulting correlation between genes and environment. Genes cause differences in the propensity to seek out educational resources, and those extra resources in turn cause differences in reading ability (Figure 3). The indirect causal structures represented in Figures 2 and 3 are in line with the causal interpretation given by Lynch (2017), which she has explicated in accordance with Woodward's (2003) interventionist account of causation.

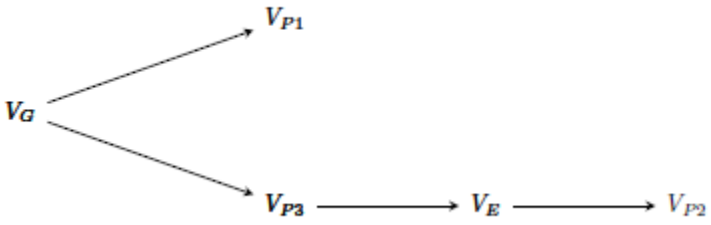

Figure 3. Active GEC. V_G (genetic variance) causes V_{P1} (variance in hair color) independently of causing V_{P3} (variance in resource seeking). This causes variance in educational resources (V_E), causing variance in reading ability (V_{P2}).

3. Interpretations of Indirect Causation

Both reactive and active GEC occur when genes indirectly cause phenotypes. Figures 2 and 3 demonstrate an identical causal structure linking gene to phenotype. Yet reactive cases are generally more unpalatable than active ones, and causal interpretations among philosophers and scientists differ between the two types of GEC.

It is generally accepted that reactive GEC cases are not good representations of genetic causation, even when the heritability estimate is high. Block (1995) claims that such cases are in 'violent conflict' with common sense ideas about causation, and most others agree that the resulting phenotypic variance in these examples are not logically attributable to genetic causes (Eaves et al. 1977; Jencks et al. 1972; Sober 2001; Kaplan and Turkheimer 2021; Coop and Przeworski 2022; Burt 2022, though see Harden 2021, 2022).

Active cases, however, are often interpreted as caused by genes, with some claiming that environmental effects are merely extensions of the genotype (Eaves et al. 1977; Harden 2021, 2022). Though some others maintain that phenotypic differences in these cases should not be thought of as genetically caused (Block 1995; Plomin 1987; Sober 2001).

This results in a problem. Active and reactive GEC appear to be causally homologous, yet intuitions about causal attribution for the two types differ. In response to this problem, most authors have accepted the causal homology between active and reactive GEC, and argued for one of three options: a blanket rejection of the use of heritability estimates to track genetic causation (Block 1995; Sober 2001), a common-sense approach to the acceptance or rejection of estimates under GEC conditions (Sesardic 2003; 2005, Kaplan and Turkheimer 2021), or a consistent application of the heritability metric as causal, which in some cases results in a violation of common-sense causal intuitions (Lynch and Bourrat 2017; Harden 2021, 2022).

The first and last approaches are motivated by the importance of homologous causal structures between 'active' and 'reactive' cases. As such, these cases are treated the same, despite differences in the causal intuitions they elicit. Lynch (2017) agrees that the cases are structurally identical but identifies other features that account for differences in interpretation. She argues that blame-worthiness, agency, and phenotypic features impact the way causal intuitions about genetic causation are elicited. This is supported by evidence of psychological influences that impact causal reasoning in genetic contexts (Lynch et al. 2019).

An alternative way to explain the interpretative differences between active and reactive GEC is to contest their causal homology. This is the approach taken by Burt (2022) who suggests new ways of distinguishing these kinds of cases. Burt terms Jencksian cases involving prejudice and discrimination 'downward causation' and argues that they are distinct from both active and reactive GEC. For Burt, active and reactive GEC involve individuals selecting or evoking an environmental response (upwards causation), while downward causation involves a social force acting upon individuals based on their phenotypes. For the reading ability example, active GEC for could involve children intentionally seeking out educational environments, reactive GEC could involve children evoking a response in teachers which leads them to pay more attention to them in class, and downward causation could involve teachers withholding of educational resources to red-haired children based on a reaction to their hair color.

Burt also argues that downwards causation is causally because it does not capture what she calls 'authentic' genetic traits: those which produce their effects via biological pathways within the body[4,5]. The notion of 'authentic genetic causation' is similar to Block and Dworkin's (1976) notion of 'direct' and 'indirect' genetic effects:

[4] It is unclear whether Burt believes GEC cases should be treated as genetically caused at all. They are not cases of downward causation, yet they do violate the 'authenticity' condition by exerting their influence outside of internal biochemical pathways.

[5] Harden (2021) also appeals to the importance of gene expression in the brain as evidence for a genetic influence on educational attainment. Though elsewhere she states that indirect genetic influences do not undermine their causality (Harden 2022).

> We shall call the effect of a gene on a phenotypic characteristic a *direct* genetic effect just in case the gene affects the characteristic by means of an internal biochemical process initiated by its product. A gene affects a characteristic *indirectly* when it produces a direct effect which in turn produces or affects a feature of the environment (including the immediate environment) which itself affects the characteristic' (Block and Dworkin 1976, 480-481)

This criterion fails to single out just downward cases, or even just reactive GEC, as there are plenty of situations where no GEC is present which exhibit indirectness. Block and Dworkin discuss the example of height as a trait which is supposed to be genetically influenced in Burt's 'authentic' sense: via internal biochemical pathways. However, to take their example, we know that the influence of genes on height also transverses outside of the body. The *GHRL* gene produces ghrelin, a molecule that affects the hypothalamus, resulting in an appetite response, which in turn affects how much nutrition the body receives via human action in their own environment (Schwartz et al. 2000). The expression of ghrelin is also mediated by environmental cues, such as over or under nutrition. These cues are themselves affected by ghrelin expression via appetite, feeding back into the causal pathway involved (Burger and Burner 2014). Further, the *GHRL* gene is variable at a population level, with multiple different SNPs at this region thought to correspond with variation in height (Gueorguiev et al. 2007).

This example shows that both Block and Dworkins 'direct genetic effects', and Burts' 'authentic genetic variants' are too restrictive criterion for genetic causation. One response could be to reject the authenticity criteria and simply exclude cases of downward causation

from those involving upward causation. For Burt, downwards causes (which capture the Jencksian examples) must occur through socio-environmental factors. To properly distinguish upward and downward causation further development is needed clarifying what kind of environments count as 'socio-environmental forces', and to distinguish what it takes to 'evoke' or 'select' an environment, compared to being subject to such forces. One could argue that red hair also evokes a response, though the evoking phenotype is not intentional and the response is not justified.

So, what makes the children evoking increased attention from teachers different to the red-haired children evoking responses from society? Is the significant difference that one involves a response to physiology (hair color), and the other to a psychological trait deemed more relevant to education (e.g., studiousness). Is a socio-environmental force dependent upon a large number of people reacting to a group in a certain way? Some of these features are captured by Lynch (2017): An individual intentionally evoking a response appears more agential and blameworthy and thus more causally responsible, compared to the red-haired children who are unintentionally reacted to. Lynch also points to the perceived relevance of the intermediate phenotype to the phenotype of interest. Studiousness seems conceptually closer to reading ability compared to red hair.

While Burt's account requires some development, it does raise an interesting question as to whether reactive, active and passive are the best ways to categorize GEC. If one takes seriously the causal homology between active and reactive cases, something more is required to explain why genetic causal attribution seems problematic in some cases and not

others. This 'something more' may not necessarily align with the current active/reactive distinction.

Both Lynch (2017) and Burt (2022) point to qualitative factors to explain this 'something more'. For Lynch it is features of the variables within the indirect causal pathway (blameworthy agents, relevance to phenotypes), for Burt it is the type of cause exerting its effect (evoking or selecting versus the action of socio-environmental forces). In the following sections I explore another point of difference. I argue against Lynch and Bourrat's (2017) assertion that all active and reactive (including Jencksian cases) GEC are causally identical. I will show that there is a crucial piece of the causal modelling of these cases that has been missing, which I believe accounts for why they are interpreted as causally distinct by many philosophers and social scientists. This difference is likely to influence causal reasoning alongside the qualitative factors suggested by Lynch (2017) and Burt (2022).

4. Confounding Background Conditions

The previous section recognized that the current division between active and reactive cases may not be the best reflection of the causal differences among GEC examples. In the following two sections I will compare a paradigmatically reactive GEC case with a paradigmatically active GEC case to illustrate a salient difference in causal structure. While the differences between two examples coheres with the traditional active/reactive divide, the salient difference I demonstrate might also be used to draw different

distinctions between types of GEC. I remain neutral on the 'right' number and way of categorizing these cases, though think this is an interesting area for future research. In this paper I simply aim to demonstrate one important and overlooked causal characteristic that corresponds with differences in intuitive causal attribution and interpretative practice within the behavior genetics community.

4.1 Jencksian Reactive Gene-Environment Correlation

Implicit in Jenck's red-hair example is the contrast that children without the red-haired phenotype are *not* denied educational access. If all children were not sent to school, then there would be no difference between in reading scores, and genes (at least those relating to hair color) would not be implicated as a cause of reading ability difference. To understand what is going on in this example, and related GEC cases which evoke uneasy causal intuitions, the contrast case between those with genetic and phenotypic differences must be made explicit.

In this scenario, a fairer and more 'common sense' investigation of the causes reading ability would be to equal the playing field between children and study the genetic differences in a population where none are subject to educational neglect. This situation involves a change to the background conditions so that environmental differences are controlled between groups. In Jenck's reactive GEC example, educational limitation and prejudice are causal background variables which confound an assessment of the causal relationship between genetic variance and variance in reading ability.

Controlling for or conditionalizing on environmental confounds is common practice in experimental and epidemiological situations looking to uncover causal relationships. For instance, Tielsch et al. (1990) observed significant differences in the prevalence of blindness between racial groups in Baltimore (USA), with blindness in Black-Americans doubling that of Caucasians. One might at first conclude that this difference stems from racially based genetic differences between the two groups, however, Sommer et al. (1991) found a second significant difference between the groups - the history of cataract surgery. Cataract surgery was 43% more common among Caucasian individuals than Black individuals of the same age, leaving a larger proportion of Black Americans with conditions leading to blindness. To accurately determine whether racial differences cause blindness differences, irrespective of confounds, researchers would need to keep fixed or account for differences in the availability or propensity to undertake cataract surgery between the two groups.

The reason that it seems problematic to attribute the reading ability of red-haired children to genes is that in Jencks' example the background variables between genotypic individuals systematically differ – leading to a difference in the causal structure of the cases. Figure 4 illustrates this difference by modelling the causal dynamics inclusive of the causal background.

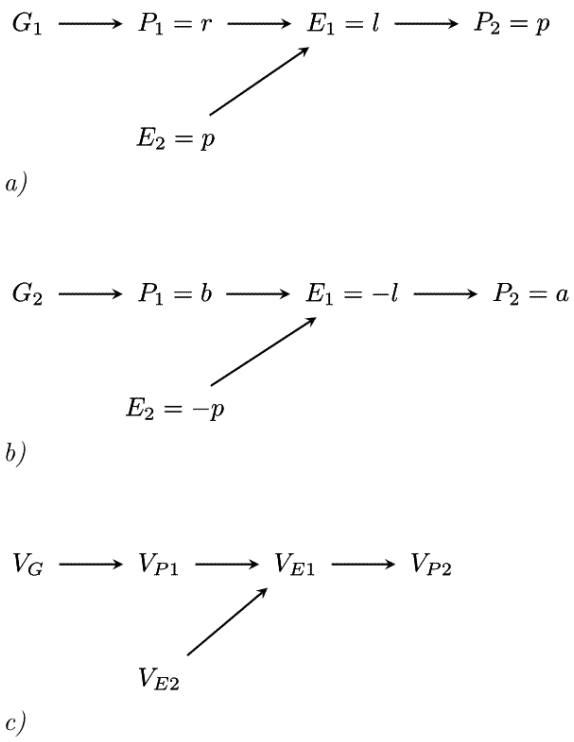

379

380 **Figure 4: Causal Background Conditions Differ Between Groups in Reactive GEC.**
381 A) illustrates the causal process in the red-haired children group, where genotype 1 (G_1)
382 causes phenotype 1 (P_1) to be red hair (r) which, along with a societal prejudice (p) in the
383 social environment (E_2) causes limited educational access (l) in the educational
384 environment (E_1), which causes poor (p) reading ability (P_2). B) shows the causal process
385 for blonde haired children, which have a different genotype (G_2), causing blonde hair (P_1 =
386 b), and as they are not subjected to prejudices (E_2 = -p) have no educational limitations
387 (E_1= -l) and attain average reading ability (P_2= a). C) shows how this translates into
388 variance for a heritability model: variance in genotype (V_G) produces variance in hair color
389 (V_{P1}), which, when combined with variance in societal prejudices (V_{E2}), produces variance
390 in reading ability (V_{P2}).

Causal background conditions systematically differ between two compared groups, with E_2 having the value p (prejudice), in the first group, and -p (no prejudice) in the second. Figure 4c represents this at the population level. Variation in background conditions (V_{E2}) confounds the causal relationship between the variables of interest (V_G, V_{P1}, V_{E1}, and V_{P2}). This confounding parallels the way cataract surgery confounded the causal relationship between race and blindness. Confounded causal relationships of this kind are likely to elicit the intuition that the causal relationship under question is suspect, as is found almost universally for Jencks-style examples.

In the cataract case the presence of a systematically varying variable was recognized as problematic, and the true causal relationship between race and blindness was arrived at by controlling for differences in the causal background. This practice is paralleled in the criteria for causation spelled out by the interventionist account: Where a cause C, causes an effect E, when a change in C produces a change in E, *when all other variables (Z) are held fixed* (Woodward 2003). In the reactive GEC case fixing other causal background variables such as societal prejudices should allow the causal relationship between the variables of interest (in this case V_G and V_P) to become clear[6].

[6] For some this may raise the question: What to fix these variables to? It may appear that this account requires a 'normal' background to explain causal reasoning (Hitchcock and Knobe 2009). An appeal to norms can be problematic as statistical, moral, social, and norms of proper functioning do not always overlap. These difficulties should not concern us here. So long as the environment selected is not systematically biased *between* treatment

To properly conduct a study on the genetic basis of reading ability differences, any systematic variation between genotype groups in the societal environment (E_2) should be eliminated. Either all children go to school, or no children go to school. Both of these possibilities are presented in Figure 5, again at the population and at the group level.

groups it does not matter whether the background conditions are considered ‘normal’ or not.

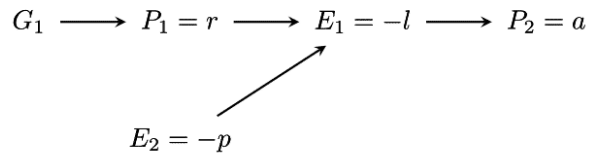

a)

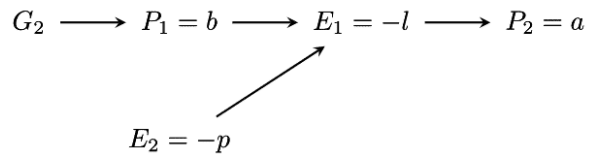

b)

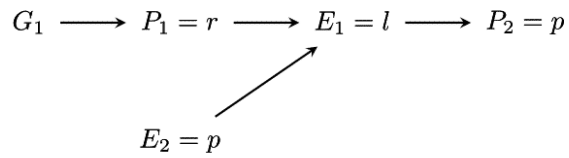

c)

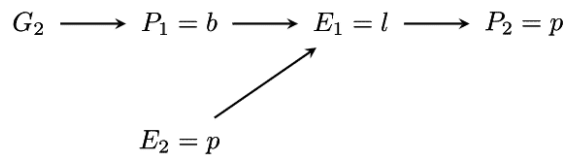

d)

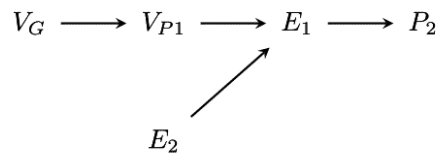

e)

412

413 **Figure 5 Conditionalizing on the Causal Background in Reactive GEC.** Educational

414 environment (E_2) is kept the same between genotype groups G_1 and G_2: either fixed at no

prejudice (-p) (A and B) or prejudice (p) (c and d). As a result, either both groups have no educational limitation (E_1 = -l), leading to average reading ability (P_2 = a) (A and B), or both groups suffer from educational limitation (E_1 = l), leading to poor reading ability (P_2 = p) (C and D). This means that there is no difference in reading ability between genotype groups, as though genotypes vary (V_G), leading to variance in hair color (V_{P1}), the societal environment does not vary (E_2), meaning that the educational environment does not vary (E_1), and there is no variance in reading ability (P_2) (E).

This fixing of background conditions illustrates why reactive GEC cases have been interpreted as problematic for a causal interpretation of heritability -the examples present an implicit confound to the causal relationship under investigation. When reactive GEC occurs there is systematic variation between genotype groups in some form of environmental variation (E_2 in this example) which impacts the causal relationship between V_G and V_P[7]. So, it is not surprising that heritabilities for cases of this kind are deemed inappropriate or invalid.

4.2 Active Gene-Environment Correlation

Let us now turn to what have traditionally been called active GEC cases. Recall that the interpretation of active GEC is less clear, with disagreement as to whether the resulting variance should be subsumed under V_G or not. To continue with the same example,

[7] In the figures V_P as it would be included in a heritability model is represented as V_{P2}.

consider again two genotypic groups, though this time those with blonde hair are genetically predisposed to alter their educational environment by seeking out extra resources and those with red hair are not. This leads to an increase in reading ability in the blonde-haired group, and thus variance in readings scores in the population (Figure 3). Figure 6 illustrates this situation with the background variables included. A disposition to seek resources (P_3 = r) in G_2 children leads (through their own volition) to development in a more enriched educational environment (E_1 = e), compared to G_1 children. G_1 children do not seek out additional resources (P_1= -r) and so instead develop in a less enriched environment (E_1 = -e). G2 children have average reading ability (P_2 = a), which is better than those with no preference to resource seeking (G1), who read poorly (P_2 = p).

Figure 6c shows this situation in terms of variance, where variance in genotype (V_G) causes variance in resource seeking (V_{P3}) causing variance in the enrichment of the educational environment (V_{E1}), in turn causing variance in reading ability (V_{P2}). V_G is an indirect cause of V_{P3} in this example, paralleling the causal structure in reactive G-E covariance cases. However, unlike the reactive scenario presented above, there is no difference in the causal background between the two genotype groups. In this example the background conditions (E_2) are assumed to be the same, or at least not to systematically differ between the two genotype groups. The background conditions in this example have the value n, for normal, which can be assumed in this case to mean no prejudice against either group or a symmetrical access to resources. Thus, the causal structure between groups does not display the same differences as the structure between the two groups as the Jencksian reactive GEC case does.

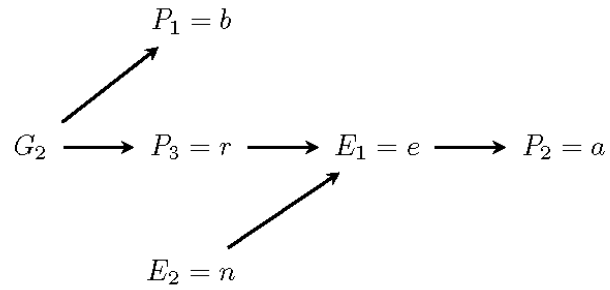

A)

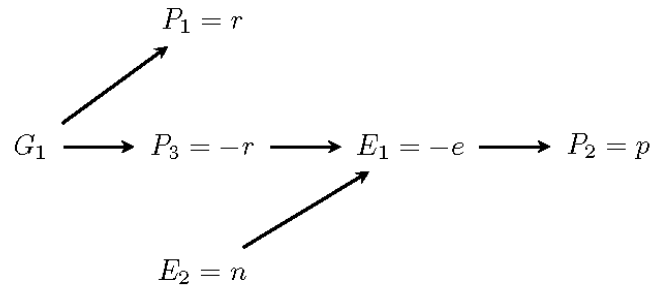

B)

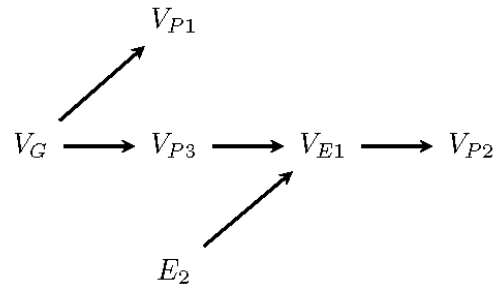

C)

456

457 **Figure 6 Causal Background Conditions Fixed Between Groups in Active GEC.**

458 Different genotypes (V_G: G_1, G_2) cause differences in hair color (V_{P1}: P_1) with the values

459 blonde (b) and red (r), and differences in resource seeking (V_{P3}: P_3); with the values of

460 resource seeking (r) or no resource seeking (–r). This in turn causes variance in educational
461 environment (V_{E1}: E_1); with the values enriched (e) or not enriched (–e). As a result, there
462 is a variance in reading ability (V_{P2}: P_2), with the values average (a) in the G_1 group, and
463 poor (p) in the G_2 group. The causal background conditions (E_2) are the same between
464 groups, with the value normal / no prejudice (n).

465

466 This shows how traditional examples of active GEC differ from Jencksian examples of
467 reactive GEC. In the Jencksian cases there are systematic environmental differences
468 between genotype groups which need to be held fixed (Figures 4 and 5) in order to
469 ascertain the causal relationship between V_G and V_P. In active GEC cases these differences
470 in the causal background are not present (Figure 6). This suggests that the differences in
471 intuition and interpretation between these kinds of active and reactive GEC cases is due to
472 differences in how the causal background conditions systematically vary between
473 genotypic groups. In reactive GEC there are often other, non-genetic variables which
474 explain variation in phenotype, which differ as V_G differs. In this example the presence or
475 absence of societal prejudice systematically varied between genotype groups. Variables of
476 this kind are causal background conditions, forming part of the environment (measured as
477 V_E in heritability estimates). In active GEC these kinds of systematic differences are not
478 usually present. In the example shown in this section, both G_1 and G_2 children were
479 exposed to the same background, where neither were subject to societal prejudices, and
480 both groups had equal educational access. Therefore, it seems that a significant point of
481 difference between some active and reactive GEC is the presence (in reactive cases) or

absence (in active cases) of systematic variation in a causal background variable between genotype groups. I believe that this explains why there are intuitive differences about how active and reactive GEC cases should be interpreted in a heritability framework.

5. Variable Description

Although the section above demonstrated a difference in causal background structures between two types of GEC, this point of difference can be dissipated by a re-description of the causal variables in the system. The illustration above is only one way of representing the systems involved. I will show in this section how a re-description of these variables can eliminate the above distinction just established. In §6 I then argue that despite this, the interpretation consistent with the results of §4 is the best way to handle these cases in behavior genetics.

An alternative way of understanding the background variables in Jencks' red-hair example is as a system of norms. That is, instead of understanding the background environment as either containing prejudice or not, the relevant background variable is represented at a more coarse-grained level, as a societal norm. This is represented in Figure 7 as 'Norm 1'. Norm 1 is the societal norm to be prejudiced against red-haired children and to abstain from granting them educational access. Societies in which Norm 1 is part of the causal background limit the educational resources of red-haired children, yet do not treat blonde-haired children in the same way. In this instance the description of the environment is

502 determined at a population level, where the values of the variable are determined by how
503 they apply to the whole population, not to individuals or groups.

504 As shown in figure 7a, when Norm 1 is combined with the presence of a red-haired
505 phenotype it causes educational limitation (E_1 = l), resulting in a decrease in reading ability
506 (P_2 = p) for G_1 children. If a blonde-haired child was to be raised with identical
507 background conditions, then in this sense they would live in same society as the red-haired
508 child, the society which has Norm 1 (Figure 7b). Controlling for causal background
509 conditions when variables are described at the societal-norm level of explanation simply
510 means that both groups live in a Norm 1 society. However, Norm 1 specifies only the
511 educational limitation of red-haired children, and when combined with the phenotype of
512 blonde hair no educational limitation occurs (E1 = – l), meaning that average reading
513 abilities are attained (P = a) for G_2 children.

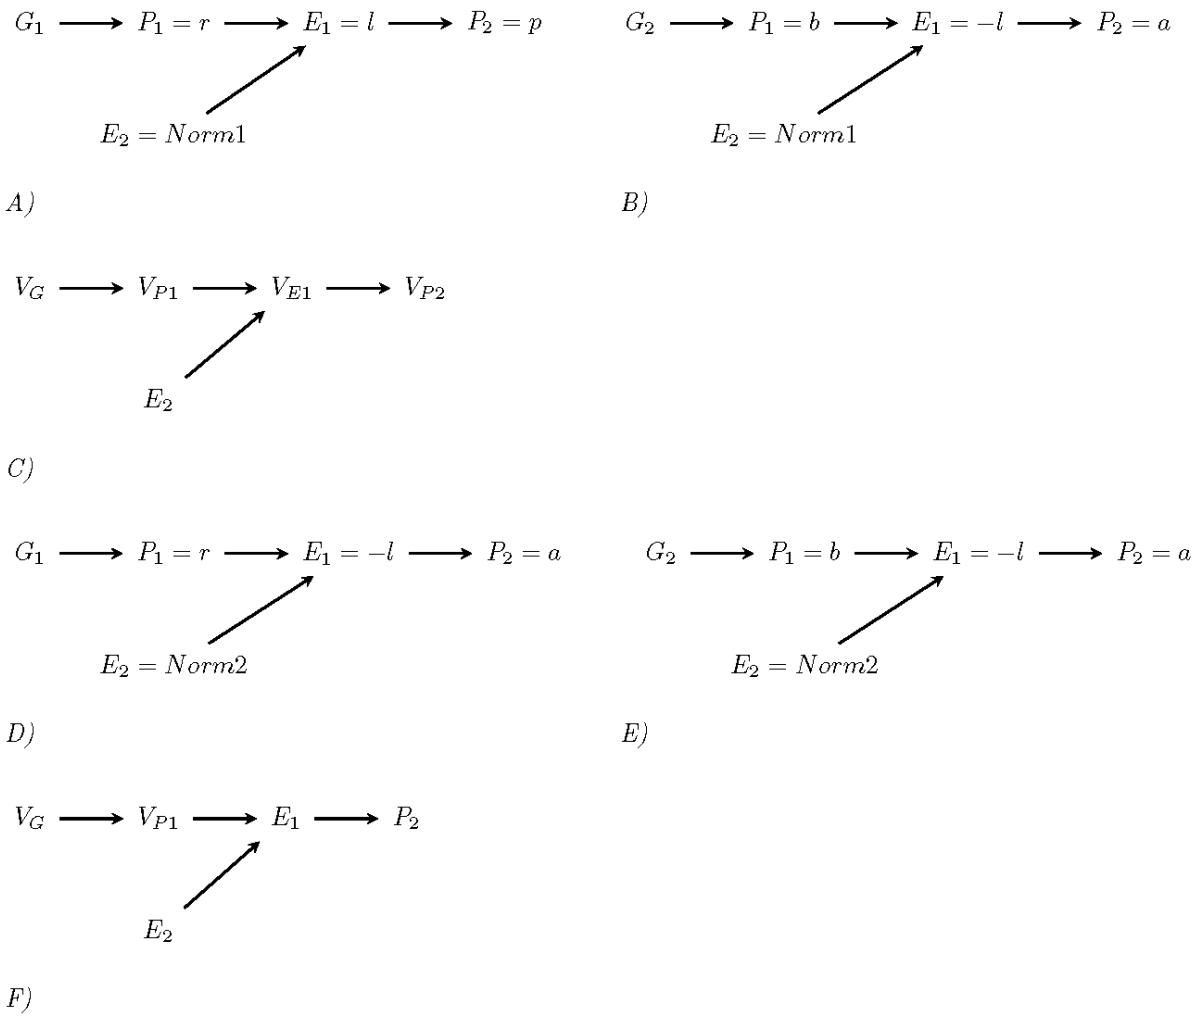

514

515 **Figure 7 A Re-description of the Variables in Reactive G-E Covariance Cases.**

516 Different genotypes (V_G: G_1, G_2) cause differences in hair color (V_{P1}: P_1); with the values
517 of red (r) or blonde (b). Both genotype groups are subjected to the same societal norms
518 (E_2), either Norm1 (figures a-c): where there is prejudice against red-haired children, or
519 Norm 2 (figures d-f): where there is no prejudice against any children. In A) G1 children
520 develop in a Norm 1 environment, so their educational is limited (E_2 = l) (as they are red-

haired), and they attain poor reading abilities (P_2 = p). In B) blonde haired children develop in a Norm 1 environment and their education is not limited (E_2 = -l), and they attain average reading abilities (P_2 = a). In D) and E) G1 and G2 children develop in a Norm 2 environment, where no prejudice is subjected to any children. As a result, both groups have no educational limitation (E_1= -l) and both attain average reading abilities (P_2 = a). Figures C) and F) show the differences between these two groups in terms of variance, where V_G represents variance in genotypes, E_2 represents a non-varying environment of Norm 1 or Norm2, V_{P1} represents variance in hair color, V_E variance in environment, E_1 no variance in environment, V_{P2} variance in reading abilities, and P_2 no variance in reading abilities.

When represented in this way, variation in background conditions (E_2) disappears between the two groups (Figure 7c). The causal structure returns to one identical to the active GEC case described in §4.2. Thus, a re-description of the background variables to occur at a more coarse-grained level of analysis undoes the differences observed between active and reactive GEC illustrated in §4.

Figure 7a, b, and c illustrate what happens when Norm 1 is maintained across the two groups – variation in reading ability occurs. But what about when this environmental background takes a different value? Imagine a society now with Norm 2: the norm for no prejudice against any children. This leads to an ample supply of educational resources for both groups. This is shown in Figure 7d, e and f. Here, differences in hair color, when

combined with Norm 2 (which is fixed between groups) produces no differential effects on E_1 (educational access) or P_2 (reading ability).

Figure 7 shows two different scenarios: In a, b and c the causal background conditions are fixed to Norm 1 –the norm for prejudice towards red-haired children. This results in differences in reading ability between the two genotype groups. In Figures 7 d, e and f the causal background conditions are also fixed, this time to Norm 2 –where no children are prejudiced against. This results in no differences in reading ability between the two groups, as neither have their education limited. What appears to be happening now that the background variables have been redescribed is that variance in phenotype (reading ability) depends upon a combination of genotype (G1 or G2) and the kind of environment present for that genotype to develop in (Norm 1 or Norm 2). In other words, reactive GEC, when considered over a range of possible background conditions (characterized at this level of description), now describes GEI.

5.1 Correlation as Interaction

Recall that GEI occurs when change in phenotype depends upon the combination of genetic and environmental background conditions. §2 described the example of MAOA interacting with childhood maltreatment to produce differences in antisocial behavior. A characteristic of GEI is that, when represented as a reaction norm, the genotypic curves diverge or cross over one another, indicating a differential effect of the environment that is dependent upon the genotype.

Jenck's reactive GEC case can be represented in this same way when the environment is characterized at a coarse-grained level of description. When the representation from Figure 7 is translated into a reaction norm, you can see the effects of two different genotypes (G1, G2) across two different environmental variables – in this case the environmental norms (Figure 8).

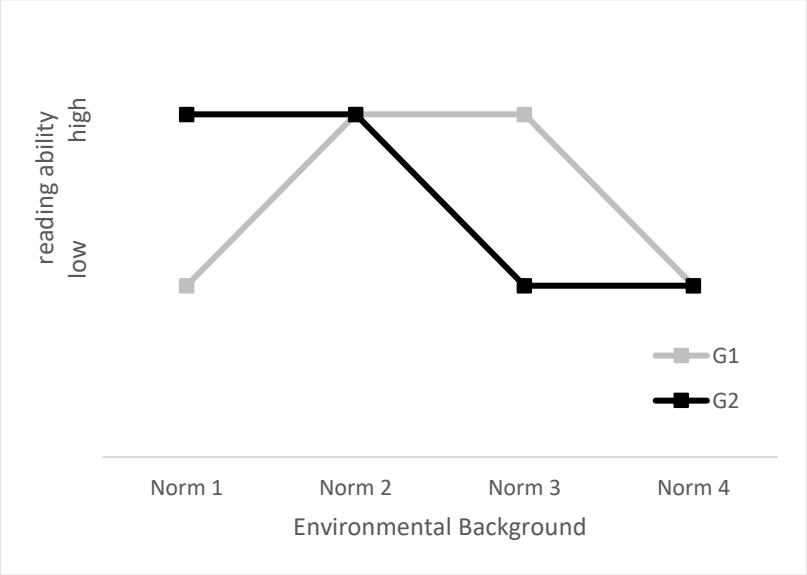

Figure 8 Norm of reaction for societal norms on reading. Norm 1 is a society which discriminates against red-haired children, Norm 2 is a society that does not discriminate against any children, Norm 3 is a society which discriminates against blonde-haired children, and Norm 4 is a society which discriminates against both blonde-haired and red-haired children. G1 represents red-haired children, and G2 represents blonde-haired children. The y axis represents reading ability, with points further up the axis corresponding to greater ability.

This is significant because traditionally GEC and GEI have been considered distinct forms of gene-environment interplay, which can be addressed with distinct experimental and statistical methods. To date, no one has considered that they can both represent the same phenomenon at different levels of description. That is, when environments are characterized at a certain level of description, phenotypic differences are best explained by appeal to either GEC, or GEI.

Related to this is the dependence of GEC on GEI. For both active and reactive GEC to occur, certain environment circumstances must be present. Children who actively seek out educational resources are only able to do so if they are available. If no resources are present, then the genetic propensity to seek out certain environments will cease to make a difference and contribute to heritability in those populations. Similarly, children can evoke extra attention from teachers only when they are in susceptible teaching environments with individuals inclined to respond.

What I have shown in this section aligns with, but goes beyond this dependence. At a narrower level of environmental description involving prejudice GEC is either present or absent in the Jencksian example, depending upon the environment. An additional contribution is that when described at a broad level of environmental description involving norms, GEC disappears, and resulting phenotypic differences between groups can *only* be understood as GEI.

In this section I have shown that the confounding background differences which distinguish between some active and reactive GEC cases disappear under certain ways of

describing background variables. When this occurs, GEC becomes GEI. I now consider which approach is most appropriate for characterizing causal background variables, and the implications this has for interpreting cases of gene-environment interplay.

6 Appropriate Variable Description

If the distinction between active and reactive GEC hinges on a particular level of description for environmental variables, it must be shown why one form of description is more appropriate or useful than another.

To do this, let us abstract away from human genetics and imagine a heritability study involving plants, in which two different plant genotypes (G1 and G2) are grown under different nutrient conditions. The nutritional environment in the study could be described in two ways. The first way divides the environment into 'high nutrient load' and 'low nutrient load'. The environment in this case is carved up according to the factors which are hypothesized to make a difference to plant growth. An alternative description of environments is as norms. For instance, 'Norm 1' is that G1 plants receive a high amount of nutrients, while G2 plants receive low nutrition. An experiment described this way would return the result that G1 and G2 plants displayed different on average growth patterns. However, these results would be interpreted as completely genetically caused, as the 'environment' was held fixed between groups.

615 This conflicts with the way in which heritability is used and analyzed in its most basic
616 form and is clearly the wrong way to conceptualize and utilize environmental variables.
617 Common-sense causal intuitions indicate that the environmental condition 'nutrient level'
618 makes a difference to plant height. Similarly, in the reactive GEC scenario, societal
619 prejudice is an intuitively important difference maker. G1 children are subject to prejudice,
620 whilst G2 children are not, leading to differences in reading ability. When I attempted in
621 §5 to re-describe these differences as norms, eliminating differences between groups, this
622 did not seem to be the right level of description for environmental background variables.

623 The sense of some levels of explanation 'seeming more right' can be captured by
624 proportionality. Proportionality is invoked to understand the best level of explanation for a
625 cause, given an effect of interest. Informally, a proportional cause is one that is described
626 as neither inappropriately broad or general, nor overly narrow (Yablo 1992). Put in
627 interventionist terms, Woodward (2010) described proportionality **(P)** as follows:

628 (**P**) There is a pattern of systematic counterfactual dependence (with the dependence
629 understood along interventionist lines) between different possible states of the cause and
630 the different possible states of the effect, where this pattern of dependence at least
631 approximates to the following ideal: the dependence (and the associated
632 characterization of the cause) should be such that (a) it explicitly or implicitly conveys
633 accurate information about the conditions under which alternative states of the effect
634 will be realized *and* (b) it conveys *only* such information –that is, the cause is not
635 characterized in such a way that alternative states of it *fail* to be associated with changes
636 in the effect. (298)

637 Woodward refers to an example taken from Yablo (1992) where a pigeon is trained to peck
638 at a target when presented with any shade of red. Given this, there are three possible causal
639 explanations[8]:

640 (1) The presentation of a scarlet target caused the pigeon to peck.

641 (2) The presentation of a red target caused the pigeon to peck.

642 (3) the presentation of a colored target caused the pigeon to peck

643 Claim (2) in this case is the most appropriate causal explanation. Under (P), (1) fails
644 condition (a), as it fails to specify alternative states of affairs where the effect will be
645 realized (like the presence of other shades of red). It also fails (b), as (under what
646 Woodward calls the "natural interpretation") it suggests that changes from scarlet to non-
647 scarlet shades of red may affect the occurrence of pecking. (3) also fails condition (a) as it
648 implies that the pigeon will peck at the target no matter the color presented, that is, it does
649 not convey accurate information about when alternative states of the effect (i.e., not
650 pecking) will be realized.

651 To apply the proportionality requirement to Jenck's reactive GEC example the two levels
652 of explanation considered are:

653 (1) An environmental value of prejudice causes poor reading ability G1 children.

[8] Woodward considers just two explanations (1) and (2) in his paper, as he invokes proportionality to defend 'higher level' explanations more generally. I have adapted his analysis with the addition of (3) which shows that some higher-level explanations can also fail P.

(2) An environmental value of Norm 1 causes poor reading ability in G1 children.

Claim (1) best fulfils the criteria given in (P), as it conveys (a) the conditions under which alternative states of the effect will be realized – if prejudice is not present then G1 children will not have poor reading ability. In this example it seems natural that the contrast class for (1) is not-prejudice, and a change to this value will produce a change in the effect. This claim also fulfils condition (b), as it does not contain any other irrelevant or misleading information.

Claim (2) on the other hand fails (P) under condition (a), as it does not convey how G1 children could have higher reading abilities – the contrast class for Norm 1 is not clear. If contrasted with Norm 4, in which all children are subjected to prejudice, then the effect for these children does not change. Norms 2 and 3 (see Figure 8) will change the value of the effect (increasing reading ability for G1 children), but it is not obvious that these are the most appropriate contrasts for this causal claim. Additionally, claim (2) fails condition (b) of the P criteria, as irrelevant information is included in the claim. Claim (2) relates to reading ability of G1 (red-haired) children, yet the explanation of these abilities (Norm 1) encompasses information about the treatment of blonde-haired children as well. Thus, the most appropriate description of background variables, is the one given in §5, which describes whether there is prejudice is experienced or not. This supports the claim that the differences between the reactive and active GEC cases considered are due to the confounding effects of environmental difference.

Appeals to proportionality have most often been used to support the selection of higher-level explanations over lower ones (e.g., Oftedal 2022). In these situations, lower-level variables can be absorbed into upper-level variables without a loss of difference making information for the effect of interest (Woodward 2021). It is clear from the Jencksian examples that absorbing prejudice towards one group and not another within a broader 'Norm' description does result in a loss of difference making information, as this difference accounts for differences in reading abilities between the two groups.

By maintaining the level of description given in §5, a causal difference is illuminated between the intuitively problematic Jenckisan cases and the more causally intuitive active GEC cases. This level of description also maintains the distinction between GEI and GEC for these kinds of examples. This explains current differences in interpretation between active and reactive cases and may also be in-line with recent suggestions for distinguishing between cases such as proposed by Burt (2022).

7 Conclusion

The interpretation of 'indirect' genetic causes, described as active and reactive GEC, has been a source of disagreement for over 40 years. Lynch and colleagues used causal modelling to suggest that these two types of cases have identical causal structures. This has left the options for distinguishing between them to appeals to common sense, psychological factors, or using a notion of within-body directness.

I have shown, using the same causal framework as Lynch and Bourrat (2017) that the two forms of GEC traditionally divided into active and reactive are not in fact causally identical. As such, I have provided an additional condition on which to distinguish and causally interpret these cases based on causal information. This has the advantage of appealing to pragmatic and psychological factors as per Lynch (2017), as those factors can vary between individuals and contexts. As assessment of implicit background confounding provides systematic criteria to address 'violently conflicting' causal intuitions and separates the examples by a feature uncontroversially regarded as important in scientific practice.

This analysis also uncovered an unexpected relationship between two types of gene-environment interplay: gene-environment correlation and gene-environment interaction. While traditionally thought of as conceptually and empirically distinct, I have shown how they are related in virtue of variable description. This is an important finding and is likely to be related to the way that levels of variable description impact upon scientific findings of statistical significance.

8 References

Block, Ned. 1995. "How heritability misleads about race." *Cognition* 56(2): 99-128.

712 Block, Ned and Gerald Dworkin. 1976, "IQ, heritability and inequality." In *The IQ*
713 *Controversy: Critical Readings, ed.* Ned Block and Gerald Dworkin, 410-540. New
714 York: Pantheon.

715 Burger, Kyle S. and Laura A. Berner. 2014. "A functional neuroimaging review of obesity,
716 appetitive hormones and ingestive behavior." *Physiology and Behavior* 136:121-127.

717 Burt, Callie H. 2022. "Challenging the utility of polygenic scores for social science:
718 Environmental confounding, downward causation, and unknown biology". *Behavioral and*
719 *Brain Sciences.*1-36.

720 Caspi, Avshalom, Joseph McClay, Terrie E. Moffitt, Jonathan Mill, Judy Martin, Ian W.
721 Craig, Alan Taylor, and Richie Poulton. 2002. "Role of genotype in the cycle of violence
722 in maltreated children" *Science*, 297: 851-854.

723 Coop, Graham and Molly Przeworski M. 2022. Lottery, luck, or legacy. A review of "The
724 Genetic Lottery: Why DNA matters for social equality" *Evolution* 76(4): 846-853.

725 Eaves, Lindon J., Krystyna Last, Nicholas G. Martin, and John L. Jinks. 1977. "A
726 progressive approach to non-additivity and genotype-environmental covariance in the
727 analysis of human differences" *British Journal of Mathematical and Statistical Psychology*
728 30(1): 1-42.

729 Figlio, David N., Jeremy Freese, Krzysztof Karbownik, and Jeffrey Roth. 2017.
730 "Socioeconomic status and genetic influences on cognitive development." *Proceedings of*
731 *the National Academy of Sciences* 114(51): 13441-13446.

Hitchcock, Christopher, and Joshua Knobe. 2009. “Cause and Norm.” *Journal of Philosophy* 106(11): 587-612.

Griffiths, Paul E. and Karola Stotz. 2013. *Genetics and Philosophy: An Introduction*, Cambridge: Cambridge University Press.

Gueorguiev, Maria, Steven Wiltshire, Edwin A. Garcia, Charles Mein, Cecile Lecoeur, Brigitte Kristen, Rebecca Allotey et al. 2007. "Examining the candidacy of ghrelin as a gene responsible for variation in adult stature in a United Kingdom population with type 2 diabetes." *The Journal of Clinical Endocrinology & Metabolism* 92(6): 2201-2204.

Harden, Kathyrn P. 2021. *The genetic lottery: why DNA matters for social equality*. Princeton University Press.

Harden, Kathryn P. 2022. “On genetics and justice: a reply to coop and przeworski (2022).” *Evolution.*

Kaplan, Jonathan M., and Eric Turkheimer. 2021. “Galton's Quincunx: Probabilistic causation in developmental behavior genetics.” *Studies in History and Philosophy of Science*, *88*:60-69.

Jencks, Christopher, Marshall Smith, Henry Acland, Mary Jo Bane, David Cohen, Herbert Gintis, Barbara Heyns, and Stephen Michelson. 1972. *Inequality: A Reassessment of the Effect of Family and Schooling in America.* New York: Basic Books.

Lewontin, Richard C. 1974. “Annotation: The Analysis of Variance and the Analysis of Causes.” *American Journal of Human Genetics* 26 (3): 400-411.

Lynch, Kate E. 2017. "Heritability and causal reasoning." *Biology and Philosophy* 32(1):25-49.

Lynch, Kate E. and Pierrick Bourrat. 2017. "Interpreting heritability causally" *Philosophy of Science* 84 (1): 14-34.

Lynch, Kate E., James S. Morandini, Ilan Dar-Nimrod, and Paul E. Griffiths. 2019. "Causal reasoning about human behavior genetics: synthesis and future directions." *Behavior Genetics* 49(2): 221-234.

James W. Madole and Kathryn P. Harden. 2022. "Building causal knowledge in behavior genetics." *Behavioral and Brain sciences.* 1-76.

James Matthewson and Paul E. Griffiths. 2017. "Biological Criteria of Disease: Four Ways of Going Wrong." *The Journal of Medicine and Philosophy: A Forum for Bioethics and Philosophy of Medicine* 42(4): 447–466.

Lucas J. Matthews and Eric Turkheimer. 2021. "Three legs of the missing heritability problem." *Studies in History and Philosophy of Science* 93:183-191.

Oftedal, Gry. 2022. Proportionality of single nucleotide causation. *Studies in History and Philosophy of Science* 93: 215-222.

Plomin, Robert. 2018. *Blueprint: How DNA makes us who we are.* Penguin Random House / Allen Lane. London UK.

Plomin, Robert, John C. DeFries, and John C. Loehlin. 1977. "Genotype-environment interaction and correlation in the analysis of human behavior." *Psychological bulletin* 84(2): 309-322.

Scarr, Sandra and Kathleen McCartney. 1983. "How people make their own environments: a theory of genotype → environment effects." *Child Development* 54: 424-435.

Schwartz, Michael W., Stephen C. Woods, Daniel Porte, Randy J. Seeley, and Denis G. Baskin. 2000. "Central nervous system control of food intake." *Nature* 404(6778): 661-671.

Sesardic, Neven. 2003. "Heritability and indirect causation" *Philosophy of Science* 70(5): 1002-1014.

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

Sober, Elliot 2001. "Separating nature and nurture." In *Genetics and Criminal Behavior*, eds. David Wasserman, and Robert Wachbroit, 47-78. Cambridge: Cambridge University Press.

Sommer, Alfred, James M. Tielsch, Joanne Katz, Harry A. Quigley, John D. Gottsch, Jonathan C. Javitt, James F. Martone, Richard M. Royall, Kathe A. Witt, and Sandi Ezrine. 1991. "Racial differences in the cause-specific prevalence of blindness in east Baltimore." *New England Journal of Medicine* 325(20): 1412-1417.

Tielsch, James M. Alfred Sommer, Kathe Witt, Joanne Katz, and Richard M. Royall. 1990. "Blindness and visual impairment in an American urban population: the Baltimore Eye Survey." *Archives of Ophthalmology* 108(2): 286-290.

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.

Woodward, James. 2010. "Causation in biology: Stability, specificity, and the choice of levels of explanation.'' *Biology and Philosophy* 25(3): 287-318.

Woodward, James 2021. "Explanatory autonomy: the role of proportionality, stability and conditional irrelevance." *Synthese* 198: 237-265.

Yablo, Stephen. 1992. "Mental Causation." *The Philosophical Review* 101(2): 245-280.