Estimating heritability of psychological traits using the classical twin design; a gentle introduction to concepts and assumptions

Nikolai Haahjem Eftedal^{a*}

^aDepartment of Psychology, University of Oslo, Oslo, Norway

The research was funded by grants 0602-01839B and 231157/F10 from the Danish- and Norwegian Research Councils, respectively (to Lotte Thomsen).

*Corresponding author

E-mail: n.h.eftedal@psykologi.uio.no

Abstract

Heritability is a concept that is often misunderstood. The term is used somewhat differently by researchers from how it is used in common parlance. This paper is a gentle introduction to the scientific concept of heritability, and to how it can be estimated for psychological traits in humans through analyses of data from monozygotic and dizygotic pairs of twins. The paper then explores some of the assumptions of the classical twin design, and presents calculations of the consequences of breaking these assumptions. Lastly, the paper introduces multivariate twin modeling, with a focus on how heritability of traits can be impacted by causal effects from other traits, and how twin designs can be informative when making causal inferences.

1. Estimating heritability of psychological traits using the classical twin design; a gentle introduction to concepts and assumptions

In behavior genetics, the term heritability is defined as *the proportion of variability in a trait that is due to the influence of genes*. Thus, all traits that individuals vary on have a heritability (even if, hypothetically, it was zero). Examples include height, neuroticism, and years of schooling. Variability in traits can be seen as the sum of variability that is rooted in genes and variability that is not, and heritability is then the genetic proportion of this sum. In common terminology, the genetic variance component is called “A”, and the non-genetic component is called “E”. (Variance components called “C” and “D” can also exist; more on this later.)

While seemingly quite straight-forward, this heritability concept has some subtleties, and some divergences from how the term is used outside of behavior genetics. In fact, the term has been described (by critics) as “...one of the most misleading in the history of science” (Moore & Shenk, 2017). To explore what heritability does and does not mean, I will discuss a series of scenarios.

1.1. Illustrative examples

1.1.1 Roses with equal conditions. Imagine a large bed of roses. Interestingly, extreme care has been taken to ensure that all the roses have exactly the same environmental conditions. They all get exactly the same nutrition, the same amount of light, and so on. Yet, the roses still vary in height. How could this be? Well, if environmental conditions for all roses are *exactly* equal, then all differences between them could be due to genes. If so, the heritability of height is 100%. This holds regardless of whether the differences between all roses are miniscule, at the scale of only a few millimeters, or if they are substantial.

1.1.2. Cloned roses with equal conditions. Suppose we take one of the roses, and we use it to make a new flower bed that consists entirely of genetically identical roses. As before, extreme care is taken to hold the environment equal for all the roses. Undoubtedly, these roses would end up quite similar in height. But would they have *exactly* the same height? Perhaps not. It turns out that there can be an element of randomness¹ in how genotypes are translated into phenotypes. For example, the world’s first cloned cat, called CC (for “Copy Cat”), had a somewhat different coloration to the cat she was cloned from, called Rainbow (see figure 1). Any such differences between roses in how genotypes are translated into phenotypes then come out as purely non-genetic; the genomes are still all the same, so these kinds of differences are not due to genes.

If the roses did turn out to be *exactly equal* in height, down to the nanometer, then heritability cannot be estimated. (And it also cannot be estimated if our instruments of measurement are not precise enough to capture any differences.) Heritability is the proportion of variance that is due to genes; so if there is no variance, there can be no heritability. Here already, we get to an important difference between common parlance and “scientific” use of heritability. In a group of people where everyone has two arms, the heritability of “number of arms” is not defined (due to the lack of variance). However, most people would perhaps want to say that the number of arms a human has is *highly* heritable; after all, it is programmed into our genome. When a population of humans does have variance in the number of arms, so that heritability can be estimated, the heritability will likely come out as low. Presumably, when people are missing an arm, this has little to do with their genes; rather, missing an arm is typically due to bad luck, in the form of accidents or illness. Cases of people with three or more arms typically come about when an embryo starts out as conjoined twins, but one of the twins degenerate completely except for one or more limbs, which end up attached to the other twin.

¹ Perhaps not *actual* randomness, but something close to it.

So, whether one's number of arms is higher or lower than two, genes tend to play little part in explaining this.

Figure 2. CC and Rainbow.

Notes. CC (left) was cloned from Rainbow (right). Yet, they look somewhat different. They are dissimilar also when viewed from the same side.

Returning to our bed of genetically identical roses with equal environments: There are other potential sources of differences in the measured heights of these roses besides randomness in gene expression. There can also be mutations in the genome: this is different from randomness in gene expression, which refers to how the *same* genome can lead to different phenotypes. With mutations, the genome changes. So, any variations in height that are due to mutations would indeed be due to differences in genes, which would increase the heritability of height in the sample².

Lastly, differences in observed phenotypes can occur through measurement errors. In the case of measuring the heights of roses, these errors would probably not be large, since this is easy to measure. But when measuring psychological traits, for example, measurement errors could account for substantial amounts of variance. This variance from error is also not genetic, and it thus reduces heritability estimates.

All these three sources of variability would also have been present in the first bed of roses, where roses were genetically different. Thus, heritability might not actually have been at 100% in that bed either, due to randomness in gene expression and measurement error.

1.1.3 Cloned roses with unequal conditions. Disregarding variance that is due to genetic mutations, the cloned roses with equal growth conditions that we just looked at would have 0% heritability for height. What would happen to this heritability if the gardeners no longer took great care to hold environmental conditions equal? What if they went out of their way to make environmental conditions very *unequal*, so that some roses got lots of water, nutrition, and light, while others got almost nothing of these things?

Surely, the amount of variance to be explained would increase, perhaps by orders of magnitude. Nevertheless, *nothing would happen to the heritability estimate*. The roses are still clones, so genes still explain 0% of the variance in their heights. The amount of variance contributed by the environment has increased heavily, but it is impossible to explain more than 100% of variance, and so the non-genetic variance component remains unchanged.

² In twin studies, which will be described later, it is simply assumed that mutations contribute no variance to the observed phenotypes. This assumption appears to be mostly justified, as mutations typically account for very little of the variation we observe in most traits. But whenever mutations do have an effect, this would come out as environmental variance in a twin study, rather than a genetic one.

1.1.4. Roses with unequal conditions. What about the very first bed of roses we looked at, with roses that were not clones? When growth conditions were equal, the heritability of height was at almost 100% (non-genetic variance can still come from randomness in gene expression and measurement errors). What would happen to this heritability if the gardeners created severely unequal growth conditions here?

In a common-sense understanding of heritability, the term means something like “the extent to which a trait is genetically fixed”. In such an understanding, the number of arms on a person should be highly heritable. And the amount of environmental input on something should not affect heritability. If heritability indexes how resistant something is to environmental influences, it should stay the same regardless of how much environmental input is actually present. However, heritability in behavior genetics does *not* index the extent to which something is genetically fixed. At least not directly. Rather, as mentioned, it shows the percentage of variance that is due to genetics. When the variance from the environment goes up, as it does when the gardeners start treating roses unequally, heritability goes down. Genes then account for a smaller proportion of variance.

Still, the common-sense understanding of heritability and the “scientific” one are not *completely* unrelated. We can imagine two species of roses, A and B, where height is more “common-sense heritable” in species A than it is in species B. In other words, species A tends to grow to its genetically specified height whether the environment is favorable or not. The environment would have to be either really bad or really good for the height of these A-roses to deviate from what they would have otherwise been (the A-roses can still differ amongst themselves in height, but these differences are mostly unaffected by the environment). The B-roses, on the other hand are highly sensitive to gardening. If you treat them somewhat well, they will all grow substantially taller than if you treat them somewhat poorly. Even slight differences in their growth conditions will impact their relative heights. If the gardeners go out of their way to cause differences in the heights of the A-roses, they will not have much success unless they resort to quite extreme measures. And so, heritability in the scientific sense will remain high, whether conditions are equal or not. Genes will still account for most of the differences between roses, even if environmental conditions vary substantially. For B-roses, on the other hand, any heritability of height that is there when gardeners treat plants equally will very quickly go away when gardeners start treating them unequally. So, if scientific heritability remains high even if you can observe large differences in how individuals are treated, this suggests that the “common-sense” heritability is high as well. And if scientific heritability is low even if most individuals seem to be treated quite similarly, this suggests that the trait is sensitive to even small differences in environmental conditions.

1.1.5. Two additional scenarios, in quiz format

Here are two more scenarios to illustrate some additional counterintuitive aspects of the way behavior geneticists use heritability. They are presented in quiz format.

1. Roses are growing in a bed, and environmental conditions are not tightly controlled. The roses are either purple or yellow; there are about 50-50 of each color. Roses of different colors do not differ in their average height, nor in how their heights are distributed. Most roses are about equally tall, with variations being mostly on the scale of a few centimeters or less. The heritability of height is at about 40%.

Then, the gardeners come up with a new practice. They decide to cut the heads off all the purple roses, causing them to be substantially shorter.

Will this cause the heritability of height in this bed to go up, to go down, or to stay the same?

Answer: The heritability of height goes up. (Given that color is determined by genes.) The new practice causes the variance in height to increase. And when a head is cut off, this is always the end-point of a causal chain that starts with genes. And so, all this new variability in height is causally rooted in genes, even if it comes about through the way the environment (the gardeners)

reacts to this genetically grounded variation (cf. Lynch, 2017).. Genes then explain a larger share of variability than before, and heritability goes up.

2. There are two separate beds of roses. Both beds have the same species of roses in them (though they are not all clones). In both beds, environmental conditions are quite tightly controlled. So heritability of height in both beds is quite high, at ~70% in both beds; environments do not explain much of the variance in each bed. However, environmental conditions are quite different between the two beds. In the first bed, the roses get excellent conditions for growth, while in the second bed conditions are quite bad. So, even though the roses are the same species, those in the first bed are on average 10 cm taller than those in the second (the standard deviation of height within each bed is smaller than this, at about 1 cm, let's say).

If the two beds of roses were combined to create a single sample, while the differences in environmental conditions remained in place, would the heritability of height in the sample as a whole be higher, lower, or about the same as the heritability of height when each of the beds are looked at separately?

Answer: The heritability of height would go way down as compared to what it was when the beds were looked at separately. Due to the differences in environmental conditions, the combined sample has much more variance in height than the separated samples. And all this new variance is due to the environmental differences. So the environment explains a much larger proportion of the variance in the combined sample than in the separated samples, meaning that the heritability of height in the combined sample is much lower.

This scenario can also be flipped, in a sense: As discussed earlier, the heritability of “number of arms” in a group of humans is very low. And so too would this heritability be in a group of squids; any deviations from the genetically specified number of arms would be likely to be due to bad luck³ etc. However, if the heritability of number of arms was to be estimated for a sample with 100 humans and 100 squids, for example, then it would be close to 100%. There is now lots of new variability in our variable, and it is all due to differences in genes.

An approach to estimating heritability in humans is the *classical twin design*, which is discussed in the remaining sections. Interestingly, if the heritability of number-of-arms was to be estimated with this design, which uses samples of twin-pairs, then all the genetically caused variance would, absurdly, be misinterpreted as originating from the so-called *shared environment* (comprising all things *except genes* which can cause twins to be similar to each other; more on this to come). A sample with both humans and squids would not fulfill the assumption that there is no assortative mating on the trait in question (also discussed later). In fact, such a sample would break this assumption to the fullest extent, meaning that *all* genetic variance is attributed to the shared environment.

1.2. Heritability of traits in humans; twin studies

Let us suppose that a specific human trait can only be influenced by the following two things: Additive genetic effects and randomness. If we identified such a trait, and we were able to accurately measure it, then the heritability of the trait could be estimated with the right kind of sample. One such right kind of sample is one containing only pairs of twins, with good amounts of both monozygotic and dizygotic pairs. With such a sample, we would be able to tell how much of the variability in the trait was due to additive genetic effects, and how much was due to randomness.

Let us first suppose that the trait in question is purely genetically determined. To have a name for the trait in question, let us say it is “musicality”. This is a trait that people vary on, and we will suppose that it can be measured with perfect accuracy. And, again, variance in musicality, in this contrived example, is *solely* due to variance in genes: in the hypothetical world for this example,

³ For both humans and squids there could still be a genetic component here: genes could make someone more likely to seek out the kinds of situations where they might lose an arm, for example.

people have a genetically pre-determined level of musicality that is never affected at all by their experiences and environmental conditions.

Furthermore, let us say that musicality is a highly *polygenic* trait. This is a realistic supposition: It turns out that most human traits are quite highly polygenic (Plomin et al., 2016). This simply means that they are affected by a large amount of genes, and, typically, each of these genes has only a very small effect. As mentioned, we also assume that these genetic effects are purely additive. This means that the effect of each gene on the trait in question is fully independent of the effects of the other genes. It is conceivable that phenotypes are *not* simply additive summations of genes, but rather emergent attributes depending on the total constellation of genes. Whenever phenotypes are “more than the sum of their genetic parts”, the assumption of additivity is broken. So it then seems unlikely that the assumption of additivity is *strictly* true, in most cases. Researchers often think that it is close enough to being true to be workable (but there have been suggestions that this is not always warranted, such as from Daw, Guo, and Harris (2015)). In any case, musicality, in this example, is stipulated to be fully determined by additive genetic effects.

If musicality is 100% heritable, how are we supposed to be able to tell that this is the case with a sample of twins? Well, if variation in a trait is purely due to genes, then the correlation between monozygotic twins should be perfect, since they have identical genomes. However, twins also have a so-called *shared environment*: most or all of them have grown up in the same house, on the same street, in the same city, and so on. So, just from a correlation between monozygotes, we cannot rule out the hypothesis that some or all of the variation we see is due to effects of this shared environment.

This shared environmental variance component is called “C”, and it comes in addition to the additive genetic variance component “A”, and other environmental variance “E”. For simplicity, we will exclusively be dealing with variables that are standardized, so that their total variance is 1. A, C, and E will then always be numbers between 0 and 1, and the three of them will never add up to a number greater than 1.

The key to distinguishing A from C is dizygotic twins. While monozygotes have the exact same genomes, dizygotes only share on average 50% of their segregating genes. What does this mean? It does not mean that only 50% of their DNA is the same. After all, humans share ~99% of our DNA with chimps; yet dizygotic twins tend to be more alike than a random pairing of a human and a chimp. Rather, it means that among the genetic effects that push a person’s score on a trait away from the population average, their dizygotic twin will share on average half of these effects. Thus, the correlation between dizygotes on the given trait will be half that of monozygotes. For a trait that is solely shaped by the shared environment, on the other hand, the correlation between dizygotes will be as it was for monozygotes, at 1. Both mono- and dizygotic twins have equal shared environments (as per the definition of shared environment).

If both genes and the shared environment impact a trait, the correlation between monozygotes (“ r_{MZ} ”) will remain perfect. The correlation between dizygotes (“ r_{DZ} ”) will end up somewhere between perfect and .5. Per Falconer’s (1960) formulas (which depend on a series of assumptions that we will soon look at) the expected correlations between twins for a trait that has specified A and C components is as follows:

$$r_{MZ} = A + C$$
$$r_{DZ} = \frac{A}{2} + C$$

What happens when there are other things than genes and the shared environment impacting a trait? In other words, what happens when the *unique environment* contributes to trait variance? In the real world, monozygotes will tend to not be exactly equal on most traits. These differences between monozygotes must be due to factors that affect only one twin, or which affect the twins differently. The A-component is defined as all things that make monozygotes more similar than dizygotes; it is assumed that only genes can do this. The C-component is defined as all things except genes that make

twins similar to each other. All remaining influences on the trait in question must then be things that make twins *less* similar to each other: all of these influences are categorized as part of the E-component, which can also be called the unique environment or the non-shared environment. This E-component can contain a surprisingly large array of different things, as we have already touched on in our discussion of roses. Most prototypically, E contains unique experiences, that are had by one twin but not the other; if one twin is given music lessons while the other is not, for example, and this causes the twins to differ more on musicality, then this is counted as E-variance. E also contains unique effects from the same experience: if both twins are given music lessons, but only one of them develops higher musicality from this, then the music lessons still become part of E even if both twins get them. E contains all things that make twins different except genes⁴, and so if music lessons somehow cause twins to become more different, then they are part of E even if both twins get them. As we have touched on, E also contains all kinds of measurement error. And, as we have also touched on, with the example of CC the cat, E contains so-called stochasticity in gene-expression. The processes translating genotypes into phenotypes are complex and unpredictable; thus, the same genome can give varying phenotypes even when environmental conditions are equal.

Now, to answer the question that started off the previous paragraph: When E contributes variance in a trait, the correlations between monozygotes (r_{MZ}) become lower than 1. In fact, the E-component will be equal to the distance between r_{MZ} and 1. Thus, we can make formulas for estimates of all three of A, C, and E (estimates are distinguished from true parameter values by the addition of a “hat” over the letters):

$$\begin{aligned}\hat{A} &= 2 * (r_{MZ} - r_{DZ}) \\ \hat{C} &= 2 * r_{DZ} - r_{MZ} \\ \hat{E} &= 1 - r_{MZ}\end{aligned}$$

The ratio between r_{MZ} and r_{DZ} then reflects how the variance that is not from E is divided between A and C. If r_{DZ} is exactly half of r_{MZ} , then the non-E variance is all A. If r_{DZ} is equal to r_{MZ} , then the non-E variance is all C. And if C is somewhere between r_{MZ} and $r_{MZ}/2$, then the non-E variance is divided between A and C. In fact, the way the space between r_{MZ} and $r_{MZ}/2$ is divided by r_{DZ} will exactly reflect the division of the non-E variance between A and C. For illustration, let us say r_{MZ} is .6 and r_{DZ} is .4. r_{DZ} is then one third of the way up from $r_{MZ}/2$ up to r_{MZ} . And thus C explains one third of the variance that is not explained by E. A then explains the remaining two thirds of non-E variance. E’s proportion of the variance comes to $1 - 0.6 = 0.4$. Of the remaining 60% of variance, one third is explained by C and two thirds are explained by A. So C is .2 and A is .4.

In practice, variance components for traits are not simply calculated from correlations as described above, but rather estimated using maximum likelihood estimation on structural equation models. This has several advantages, such as allowing the estimation of confidence intervals, increasing flexibility in how models are set up, and better handling of missing data.

Samples with mono- and dizygotic twins then allow partitioning the variance of a trait into the three components called A, C, and E. This model of trait variance can be illustrated like this:

Figure 3. Variance decomposition

⁴ E cannot simply be defined as “all things that cause twins to differ”: dizygotes differ genetically, which can contribute to their phenotypic differences.

The circles are latent factors representing all the influences from its specified source. They can be thought of as generators of random draws from a standard normal distribution (i.e. a normal distribution with a mean of zero and standard deviation of one): each person gets one number drawn for each of the three latent factors. The arrows represent causal effects: the latent factors all causally affect the observed phenotype, which is represented by the square labeled “Trait”. The small letters (a, c, and e) represent the size of the causal effects; they can be seen as regression coefficients. If the trait is to get a variance of 1 in the population, then $a^2 + c^2 + e^2$ must equal 1. The small letters are then square roots of the share of variance in the trait explained by the source of variance associated with that letter (e.g. if A explains 36% of variance in the trait, then the coefficient “a” becomes .6).

1.2.1. *Non-additive genetic effects*

By the logic I have now sketched out, the true value for r_{DZ} will be somewhere between r_{MZ} and half of r_{MZ} . But what if we find that r_{DZ} is outside this range? Can this be interpreted as anything other than sampling error? The answer here depends on whether r_{DZ} is below or above this range. If r_{DZ} is *higher* than r_{MZ} , then that would indeed not make much sense, and should in fact be interpreted as random error, in most circumstances. If r_{DZ} is *lower than half* of r_{MZ} , on the other hand, this can be explained by the introduction of a new kind of variance component, called D.

D stands for dominance genetic effects, and represents the phenomenon whereby one allele overrides the effects of the corresponding allele on the other copy of the chromosome. This would then be a genetic effect that is *not additive*. When such effects are present, monozygotes will still fully share them, but the correlation between dizygotes will be at only .25 (all correlations discussed here are Pearson’s r).

As detailed by e.g. Chen et al. (2015), there are many kinds of processes by which the effects of genes can be non-additive. In addition to D-effects, where there are interactions between alleles within the same locus, there can also be interactions between different loci (“epistasis”), as well as other types of higher-order interactions. Furthermore, particularly with regards to complex psychological traits, there are many plausible ways whereby genetic predispositions towards certain traits can create feedback loops, so that these traits are then amplified. If someone has a particular talent, they might actively cultivate it further, or they might seek out environments that are conducive to their abilities. And, conversely, if someone is particularly untalented at something, they might actively avoid doing it. As per how genetic effects are defined in this framework, the effects of this kind of active cultivation of traits are still defined as genetic, since the causal chain starts with the genome. But these kinds of effects might very well not be additive: they could cause correlations between monozygotes that are more than twice those of dizygotes.

All these kinds of non-additive genetic effects will lead to increased estimates of D, as they push r_{MZ} and r_{DZ} further apart than additive effects do. Yet, while these kinds of processes might seem plausible, it is still rare to see models with D-components as the final model in published research. According to some, such as Chen et al. (2015), this might be related to how D (as well as other non-additive effects) and C can counteract each other: D pushes r_{MZ} and r_{DZ} away from each other (by pushing r_{MZ} upwards) while C pushes them closer together (by pushing r_{DZ} upwards). So, if they are both present at the same time, r_{MZ} could very well end up being about twice as large as r_{DZ} , which is exactly the same pattern of correlations that is expected if *neither* C nor D are present; only A and E. Thus, C and D can mask each others presence: Models with all of A, C, and D present at the same time cannot be estimated with the classical twin design. The model is under-informed to allow estimation of more than one source of deviance from pure additivity (see e.g. Neale & Cardon, 2013). You can have ACE-models, CDE-models, and ADE-models. But you cannot estimate all of A, C, and D at the same time (i.e. an ACDE-model). With designs that include data for other family members, such as parents, siblings, half-siblings and so on, it is possible to make ACDE-models (e.g. Wang et al., 2011). These kinds of samples are less common, and it largely remains an open question how frequently C and D mask each other across different traits.

1.3. Assumptions of the classical twin design

With this introduction of D, we can now look at the full model in the classical twin design:

Figure 4. The classical twin design

Note. A, C, and D are all perfectly correlated for monozygotes. For dizygotes, C remains perfectly correlated, but A correlates at .5, and D correlates at .25. E's are always fully independent. The squares are observed phenotypes for the first and second twin in the pair. Variances for factors and traits are not drawn in, but they are all 1. (Note that this full ACDE-model is not identifiable with a standard twin sample; some of the parameters must be constrained.)

Here is a run-through of some of the assumptions that are inherent in this design, and a discussion of their implications, starting with further discussion of the assumption that A, C, and D are not all present at the same time (for a more complete overview of all assumptions, see e.g. Neale & Cardon, 2013) :

1.3.1. Assumption: A, C, and D not present all at the same time

To people familiar with path-tracing rules, the Falconer formulas for r_{MZ} and r_{DZ} can be discerned from this figure. Or, rather, updated versions of the formulas, which also include D, can be discerned. The small letters in the figure are regression coefficients, and their values are the squareroot of the amount of variance explained by their factor. Since all paths between Twin 1 and Twin 2 here will count both of a particular small letter, the letters are multiplied with themselves. For monozygotes, all correlation paths are at 1, so the expected correlation between trait scores for monozygotes (i.e. r_{MZ}) the sum of a^2 , c^2 , and d^2 , which again is equal to the amount of variance in the trait explained by A, C, and D. For dizygotes, variance components must be multiplied by correlation coefficients that are not equal to 1. The formulas then become:

$$r_{MZ} = A + C + D$$

$$r_{DZ} = \frac{A}{2} + C + \frac{D}{4}$$

Earlier, we derived expressions for A, C, and E as functions of r_{MZ} and r_{DZ} . Then, we simply assumed that D did not exist. It turns out that when we no longer make this assumption, we cannot make expressions for all four of A, C, D, and E just based on r_{MZ} and r_{DZ} . There are three unknowns but only two informative equations, so the model is unidentified. We are then forced to choose at least one out of A, C, and D to set to some specific value (usually zero), in order to get estimates for the other two. But what happens if all three actually exist? For example, what happens if our trait is actually shaped by D-effects (or other kinds of non-additive genetic effects), but we run analyses as if only A, C, and E are relevant? When D is assumed to be zero, the expressions for A, C, and E are as follows: $A = 2 * (r_{MZ} - r_{DZ})$; $C = 2 * r_{DZ} - r_{MZ}$; $E = 1 - r_{MZ}$. We can then simply replace r_{MZ} and r_{DZ} with their values when D is non-zero, to see what happens to estimates of A, C, and E.

$$\begin{aligned}\hat{A} &= 2*((A + C + D) - (\frac{A}{2} + C + \frac{D}{4})) = A + \frac{3*D}{2} \\ \hat{C} &= 2*(\frac{A}{2} + C + \frac{D}{4}) - (A + C + D) = C - \frac{D}{2} \\ \hat{E} &= 1 - (A + C + D) = E\end{aligned}$$

Thus, if D is assumed to not exist even though it does, all of D is shifted over to A in the analysis. Additionally, a portion of C equal to half of D will *also* be shifted over to A. E remains unaffected. In other words, A is overestimated and C is underestimated.

The same procedure can be repeated for the situation where an ADE-model is fitted even though C is non-zero, resulting in:

$$\hat{A} = A + 3*C \quad ; \quad \hat{D} = D - 2*C$$

All of C goes to A. E is left untouched. And, crucially, a share of variance equal to two times C is moved from D to A. Thus, it does not really take much C at all in order to completely hide D. For example, let us say the true variance components are as follows: [A = .1; C = .2; D = .4; E = .3]. Then, r_{MZ} becomes .7 and r_{DZ} becomes .35. This means that, whether in an ACE or ADE model, all the variance that is not E goes to A: $\hat{A}$ becomes .7, $\hat{E}$ becomes .3, and both $\hat{C}$ and $\hat{D}$ are 0. When r_{DZ} goes past half of r_{MZ} , as it can easily do when there is C-variance, D will come out as negative, meaning that ADE models will not even be considered. Similarly, when r_{DZ} is below half of r_{MZ} , models with C components will be dismissed.

1.3.2. **Assumption: No correlations between different variance components**

There are many correlational paths that could have been in the model but which are not. Specifically, all correlations between differing kinds of variance components are set to zero (i.e. there are no correlational paths between circles with different letters on them). We will look at two of these correlations that are assumed to be zero, and which are particularly interesting: the correlation between A and C, and the correlation between A and E.

1.3.2.1. *Correlation between A and C*

The lack of a correlation between A and C implies that one's genetic predispositions are unrelated to one's shared environment. For example, in the case of height, people who have a genetic predisposition towards being tall are assumed to not be any more or less likely than other people to be born into an environment that is also conducive to being tall. This could be quite accurate for height, but it is perhaps more questionable in the case of certain other traits. It could well be the case that genes for being physically active are more common among children born in environments that further stimulate that trait: parents of children with such genes will tend to have these kinds of genes themselves, which could influence the kind of environment they create. A similar story can be told for e.g. the musicality trait discussed earlier (more instruments around the house, etc) or for intellectual ability (more books around the house, etc).

When this correlation between A and C is assumed to not exist when it actually does exist, the influence of the shared environment on the trait in question will be *overestimated*, while the influence of both additive genetic effects and the unique environment will be *underestimated*.

When all assumptions in an ACE model are fulfilled, the variance “V” of the phenotype being analysed becomes equal to the sum of the squared regression coefficients, for both monozygotes and dizygotes: $V = V_{MZ} = V_{DZ} = a^2 + c^2 + e^2 = A + C + E = 1$. If we introduce a correlation of size r between A and C, however, the variance becomes: $a^2 + c^2 + e^2 + 2acr$. A new term is added to the total variance, which is $2*a*c*r$. The formulas for twin-correlations also need to be updated:

$$r_{MZ} = \frac{A + C + 2acr}{V_{MZ}} \quad ; \quad r_{DZ} = \frac{\frac{A}{2} + C + 2acr}{V_{DZ}}$$

Note that V_{MZ} and V_{DZ} are assumed to be equal to V , which is now equal to $1 + 2acr$. Both the nominator and denominator in both expressions are then increased by $2acr$ as compared to the formulas when A and C are not correlated (the denominator, V , is not typically included, because it is usually equal to 1). The effect of this is that both r_{MZ} and r_{DZ} increase. The amount they increase by depends on how large they were to begin with: the smaller they were, the more they will increase. To illustrate: an increase from $\frac{1}{10}$ to $\frac{2}{11}$ is larger than the increase from $\frac{9}{10}$ to $\frac{10}{11}$. Thus, since r_{DZ} is usually smaller than r_{MZ} , it will increase by more than r_{MZ} does, meaning that C will take a larger chunk of the variance left to A and C . Since r_{MZ} also increases, there will be more variance left from E to share between A and C . But the increase in r_{DZ} will always make it so that A decreases.

Here are formulas to calculate the new expected values for A , C , and E from the values they had before the introduction of a correlation between A and C :

$$A_{\text{new}} = \frac{A}{V} ; E_{\text{new}} = \frac{E}{V} ; C_{\text{new}} = \frac{C + 2acr}{V}$$

The new variance of the trait becomes $V = A + C + E + 2acr = 1 + 2acr$. A and E retain their original size, which now takes up a smaller portion of the total variance, so the new values for A and E become the old values divided by V . All of the new variance gets transferred to C .

1.3.2.2. Correlation between A and E

There are three ways in which correlations between genes and environments (r_{GE}) can come about: Passive r_{GE} ; active r_{GE} ; and evocative r_{GE} (Scarr & McCartney, 1983). Passive r_{GE} is basically the notion that genes are not randomly distributed across geographical and/or sociological locations. We have already dealt with passive r_{GE} in our discussion of correlation between A and C . Active and evocative r_{GE} rather refer to how genes can influence ones environments. In active r_{GE} it is the person who actively shapes their environments or actively seeks out new ones. Evocative r_{GE} is rather that something about a person's genotype influences how the world acts upon them, even though the person is in no way active in bringing this about, as in the previous example where all the purple, but not yellow, roses had their heads cut off..

As opposed to passive r_{GE} , which only shapes the shared environment, active and evocative r_{GE} can shape the unique environment as well. This would then correspond to a correlation between A and E . The effects of introducing a correlation between A and E on estimates of A , C , and E are in certain ways quite similar to those just discussed for correlations between A and C . The important difference is that for correlations between A and E , all the extra variance goes to A rather than to C . As before, the total variance increases by a constant. This time, this constant is $2*a*e*r$, or $2aer$, where r is the correlation between A and E . The twin correlations now become:

$$r_{MZ} = \frac{A + C + 2aer}{V_{MZ}} ; r_{DZ} = \frac{\frac{A}{2} + C + \frac{2aer}{2}}{V_{DZ}}$$

These formulas are just like those from before, with two important exceptions. Firstly, the constant, $2aer$, is now slightly different, in that it contains e and not c . Secondly, and perhaps more importantly, this constant is now divided by two in the formula for r_{DZ} . This change is due to how, as opposed to the case for shared environments, dizygotes only share 50% of segregating genes (i.e. A correlates at .5 rather than at 1 as for C). The effect of this is that all of the increase in the total variance in the trait goes to A rather than C . The formulas for new values of A , C , and E become:

$$A_{\text{new}} = \frac{A + 2aer}{V} ; E_{\text{new}} = \frac{E}{V} ; C_{\text{new}} = \frac{C}{V}$$

It has been argued by Lynch (2017) and others that environmental influences that follow causally from genes should simply be defined as genetic influences. If so, whenever A influences E, this E simply becomes part of A. And this is also what is already happening in the analyses: all the new variance in the trait from introducing a correlation between A and E goes to A.

A common assumption to hear about in discussions of twin studies is the *equal environments assumption*, which says that monozygotes and dizygotes experience equally similar environmental pressures with regards to the trait in question. In other words, it is assumed that the only reason monozygotes are more correlated than dizygotes is genes. If monozygotes are somehow more strongly pressured by their environments towards being similar than dizygotes are, then this assumption is broken. This assumption can be seen as a sub-category to the assumption being discussed in this section, which is that there is no correlation between A and E. If monozygotes are more correlated than dizygotes purely due to the environment, then the trait in question would be deemed to be heritable by twin studies, since this is the conclusion whenever $r_{MZ} > r_{DZ}$. There are, however, some reasons to think that the equal environments assumption is reasonable. For example, dizygotic twins that were mistaken for being monozygotes were found to be no more similar than regular dizygotes on a series of traits (Conley et al., 2013). Also, heritability of traits is increasingly being shown to exist also through methods that do not rely on the equal environments assumption (Plomin et al., 2016).

1.3.3. Assumption: No assortative mating

The assumption of no assortative mating entails that parents pair up completely at random with regards to the trait in question. This assumption is encoded through the correlation of .50 between A-factors for dizygotes: when parents pair up non-randomly, this correlation could very well diverge from .50. If genomes of parents are positively correlated for the trait in question, the correlation between A-factors for dizygotes can go above .50. The genes that are not copies of each other in pairs of dizygotes are then more likely than chance to pull in the same direction with regards to the trait studied. If parent-scores on the trait are negatively correlated (“opposites attract”) then the correlation can go below .50 (to the extent that the trait is heritable).

If assortative mating is taking place, and analyses are run as if it is not, what will happen to estimates of A, C, and E? Well, r_{DZ} will increase. The formula for r_{DZ} can be rewritten as $r_{DZ} = A*r + C$, where r is the correlation between A-factors. This is assumed to be equal to .50, so in the standard formula A is multiplied by $\frac{1}{2}$. But if the correlation changes, then A will need to be multiplied by a different number. Here are expressions for expected estimates of A, C, and E calculated from their true values and a correlation r between A-factors for dizygotes (the stronger assortativity there is on the trait, the further this r will diverge from its assumed values of .50):

$$A_{\text{new}} = 2A - 2Ar \quad ; \quad C_{\text{new}} = C + 2Ar - A \quad ; \quad E_{\text{new}} = E$$

As can be seen from these expressions, when r is .50, they simply turn into A, C, and E, meaning that estimates are then unbiased. The further r diverges from .50, the more biased the estimates will be. E will always remain unbiased, however, since r_{MZ} remains unaffected. All bias will be towards overestimating C at the cost of A (except in an unrealistic “opposites attract” scenario, where A is overestimated).

1.4. Multivariate twin models

Until this point, we have only considered so-called univariate heritability analyses. These are analyses of the variance composition of just a single trait. But it is also possible to do *multivariate* twin models, where two or more phenotypes are included in the same model. While univariate models are set up to estimate sources of variance in a given phenotype, multivariate models also provide estimates of the sources of *covariation* between phenotypes (Plomin et al, 2013). Just like variances can be decomposed into A, C (or D), and E, so too can covariances. Variables can covary due to causal effects in one or both directions, and/or due to confounding influences in the genes or the

environment. We will discuss causal effects in more detail later. For now, it suffices to say that when a variable X has a causal effect on another variable Y , the variance components of X (except measurement error) will be carried over into Y , to an extent reflecting the size of the effect. Variables can also covary in absence of any causal effects between them, for example due to having partly shared genetic substrates, or due to how certain environments can influence both traits. For instance, the high comorbidity between anxiety and depression appears to mostly reflect overlap in genes that contribute to liability to these disorders, more so than environmental factors (Kendler et al., 2007).

Here is a schematic of the most basic multivariate twin model, the bivariate Cholesky model:

Figure 5. Bivariate Cholesky model

Note. Among the curved correlation lines, solid lines are always 1, while dotted lines are 1 for monozygotic pairs and .50 for dizygotes.

For simplicity, it is common to show only half of this model; only the part for one of the two twins. The path coefficients from latent factors onto the traits are constrained to be equal for both twins in the pair, so the two halves of the model are basically identical. (This practice of showing only half the model is common across most twin models.) The paths causing Trait 2 to load on the variance components of Trait 1 (a_{21} , c_{21} , and e_{21}) are what makes a single bivariate model more informative than a pair of univariate models; these paths show the extent to which variance components are shared between traits.

While univariate models rely on just r_{MZ} and r_{DZ} , in the bivariate case there are several more correlations which are informative. Firstly, there are now two of r_{MZ} and r_{DZ} , one for each of the two traits. Additionally, there are correlations between the two traits. These correlations can be either within the same person (“within-twin cross-trait”; how well does a person’s score on Trait 1 predict their score on Trait 2) or they can be made across the twins in a pair (“cross-twin cross-trait”; how well does a person’s score on Trait 1 predict *their twin’s* score on Trait 2). It is these correlations across the traits that can be used for partitioning covariance.

If the covariance between two traits has nothing to do with neither genes nor shared environments, then a person’s score on Trait 1 should be completely *unpredictive* of their twin’s score on Trait 2. The factors causing Traits 1 and 2 to correlate are then only found in the unique environment, so a pair of twins should be no more similar in their exposures to these factors than any random pairing of people. On the other hand, if the factors causing the traits to correlate are rooted in genes, cross-twin cross-trait correlations will indeed be expected. In fact, if the covariance between the traits is *only* due to genes, then a person’s score on trait 1 should be equally predictive of their monozygotic twin’s score on trait 2 as it is for their own score on trait 2. After all, monozygotes are genetically identical. In other words, within-twin cross-trait correlations and cross-twin cross-trait correlations for monozygotes are expected to be the same if the covariance is all due to genes. By a similar logic, these two correlations are also expected to be equal if the covariance is all due to the

shared environment, since twins are assumed to be perfectly matched on that as well. The difference is that for the shared environment, the correlations are assumed to be equal for monozygotes as well as for dizygotes: when covariances between the traits are due to genes, then the cross-twin cross-trait correlations for dizygotes are expected to be only half of the within-twin cross-trait correlation, since they only share 50% of segregating genes. The extent to which Traits 1 and 2 are more correlated within the same person than they are across monozygotic twins reflects the influence of the unique environment.

This general logic for partitioning not just variances but also covariances into A, C, and E can be extended to models with more than two traits, as well as to models with different kinds of structures than the Cholesky structure in the figure above.

1.4.1. Heritability and causality

As a prelude to introducing the logic of the ACE- β approach to causal inference, I will discuss how causal effects impact the variance composition of a trait (i.e. the traits division of variance into A, C, and E).

Suppose the variables X and Y are completely uncorrelated. Then, we introduce a causal effect from X on Y. What will then happen to the variance composition of Y? The answer is that this depends on two things: The variance composition of X, and the size of the causal effect. More specifically, the new variance composition of Y becomes a weighted sum of the initial variance compositions of X and Y. The weight to give to X and Y's initial variance components in this sum depends on the strength of the causal effect: X's components are multiplied by the amount of variance in Y that is now explained by X, while Y's initial components are multiplied by the amount of variance in Y that remains unexplained by X after the causal effect.

I will make all this more concrete by putting numbers on it. Let us suppose that X and Y have the following variance compositions before the causal effect: X: [A = .1; C = .4; E = .5], and Y: [A = .5; C = .0; E = .5]. Both X and Y are 50% E, but for the remaining variance Y is all A while X is mostly C. Then, we introduce a causal effect from X to Y. For example, we can say that if someone scores a standard deviation above the mean on X, this causes their score on Y to go up by half a standard deviation, and so on. This can be simulated by simply adding X multiplied by .50 to Y, if both variables are standardized. Whatever constant X is multiplied with here, the variance in Y will then increase by the square of that constant (since the variables are uncorrelated; if X and Y were already correlated before the causal effect was introduced, the increase in Y's variance would be higher, as per the variance sum law). So if we multiply X by .50, Y's variance increases by .25, to 1.25, meaning that X now explains 20% of the variance in Y. This new variance in Y then has X's variance composition. To get Y's new variance composition we then make a weighted sum of the old components of X and Y, where weighting is determined by the proportion of variance explained in Y. So Y's old components are multiplied by $1/1.25 = .80$ and X's are multiplied by $.25/1.25 = .20$. This gives new a variance composition for Y of [A = .42; C = .08; E = .5]. Y then gets a brand new C component which is equal to 20% of the C-component in X; A is slightly shifted down, since X had less A than Y; and E remains as before, since both X and Y had the same E.

The main point here is that when there are causal effects from X to Y, then the variance components of X will all be carried over to Y, to an extent that is determined by the strength of this causal effect. This means that if X has variance components from all of A, C, and E, then a bivariate Cholesky analysis on X and Y should show overlap on all three of A, C, and E. If X has a substantial C-component and non of this is carried over to Y, then a correlation between X and Y seems more likely to be due to genetic and/or unique environmental confounding rather than causal effects.

Unfortunately, a reversal of this line of reasoning does not work. If the results of a bivariate Cholesky analysis are *fully consistent* with there being a causal effect, this is *not* conclusive evidence for a causal effect. Such results could still be partly or wholly due to confounding.

1.4.1.1. ACE- β models; The discordant twin design in an SEM framework

As is shown by e.g. the field of economics, you can often make a lot of theoretical and inferential headway if you just make some strong assumptions. The idea of ACE- β models (Kohler et al., 2011) is to make the following strong assumption for multivariate twin models: *All unique-environmental overlap in variance is due to causal effects*.

In essence, this is the same assumption that is made in discordant twin designs (McGue, Osler, & Christensen, 2010). Discordant twin designs take advantage of how twins are matched on rearing environments (regardless of zygosity) and genes (100% for MZ, 50% on average for DZ) in order to better infer causal effects. If differences on a variable *between twins in the same pair* predict within-pair differences on another variable, this is more strongly indicative of causal effects than what is typically achievable in non-experimental designs. When relationships between variables persist despite controlling for both genes and rearing environments, the only remaining alternative to a causal interpretation (in one or both directions) is then that there are confounds in the unique environment.

ACE- β models reformulate this discordant twin logic into an SEM framework through simple alterations of standard Cholesky models. In a standard Cholesky, the A, C and E-components of the first trait are allowed to also have paths onto the second trait. ACE- β models build from this by having an additional direct causal path from the first trait to the second. In order for ACE- β models to still be identifiable, the path from the E-component of the first trait onto the second trait is set to zero. In other words, it is assumed that unique environmental influences on the first trait do not also have effects on the second trait, except through the causal path. This formalizes the assumption of no E-confounding, which is required for making causal interpretations from discordant twin data.

When you make the assumption that all the overlap in E is due to a causal effect, it then becomes possible to calculate how much of the overlap in A and C that is due to this causal effect, and how much of that is not. Here is a figure comparing a regular bivariate Cholesky model with a bivariate ACE- β :

Figure 6. Bivariate Cholesky- and ACE- β models

The only difference is that the e_{21} path from E_1 to Trait 2 in the Cholesky has been removed, and that the causal path p_{21} from Trait 1 to Trait 2 has been added. An effect of this change is that estimates for a_{21} and c_{21} will change. While these coefficients reflect all kinds of overlap in the cholesky model (whether due to causal effects or confounding influences), they only reflect non-causal overlap in an ACE- β . All the overlap in A and C components between traits 1 and 2 that is due to causal effects is accounted for by the causal path.

1.4.1.2. Unique environmental confounding

Assuming that *all* shared E-variance⁵ between traits 1 and 2 is due to a causal effect is, as mentioned, a strong assumption. There are plausible mechanisms by which E-variance can be shared

⁵ Terminology might be confusing here. E is the unique environment, which can also be called the non-shared environment. So how can there be shared non-shared environmental variance? “Shared” refers to how the variance is shared between the two traits in the model; “Non-shared” means that the variance is not shared between the twins in a pair.

between two traits also in absence of causal effects. Experiences had by one twin but not the other can plausibly affect both traits at the same time. Additionally, for most pairs of traits that might be causally connected, there are many traits that could work as third-variable confounds. Such as intelligence or personality traits. Such traits are often in part shaped by the unique environment. All unique environmental influences on confounding traits should then constitute shared E-variance between the two traits in the analysis.

Thus, evidence for causal effects from ACE- β models such as ours should be met with some skepticism. On the other hand: there is something that is very common which *suppresses* overlaps in E-variance. Namely, *measurement error*. As discussed in the introductory section (about roses), measurement error is part of the unique environment. If error is random, all the variance it contributes to a trait will not be shared by the other trait. This makes it so that the overlap in E-components can appear much smaller than it would be without error. Which again can make it so that causal effects can be hidden from ACE- β models. The prevalence of measurement error is likely to be an important reason why substantial E-overlap, or even statistically significant E-overlap, is quite rare to see in real data.

References

- Chen, X., Kuja-Halkola, R., Rahman, I., Arpegård, J., Viktorin, A., Karlsson, R., ... & Magnusson, P. K. (2015). Dominant genetic variation and missing heritability for human complex traits: insights from twin versus genome-wide common SNP models. *The American Journal of Human Genetics*, 97(5), 708-714.
- Conley, D., Rauscher, E., Dawes, C., Magnusson, P. K., & Siegal, M. L. (2013). Heritability and the equal environments assumption: Evidence from multiple samples of misclassified twins. *Behavior Genetics*, 43(5), 415-426.
- Daw, J., Guo, G., & Harris, K. M. (2015). Nurture net of nature: Re-evaluating the role of shared environments in academic achievement and verbal intelligence. *Social science research*, 52, 422-439.
- Falconer, D. S. (1960). Introduction to quantitative genetics. *Introduction to quantitative genetics*.
- Hayden, E. C. (2013). Taboo genetics. *Nature*, 502(7469), 26.
- Kendler, K. S., Gardner, C. O., Gatz, M., & Pedersen, N. L. (2007). The sources of co-morbidity between major depression and generalized anxiety disorder in a Swedish national twin sample. *Psychol Med*, 37(3), 453-462.
- Lynch, K. E. (2017). Heritability and causal reasoning. *Biology & Philosophy*, 32(1), 25-49.
- McGue, M., Osler, M., & Christensen, K. (2010). Causal inference and observational research: The utility of twins. *Perspectives on Psychological Science*, 5(5), 546-556.
- Moore, D. S., & Shenk, D. (2017). The heritability fallacy. *Wiley Interdisciplinary Reviews: Cognitive Science*, 8(1-2), e1400.
- Neale, M. C. C. L., & Cardon, L. R. (2013). *Methodology for genetic studies of twins and families* (Vol. 67). Springer Science & Business Media.
- Plomin, R. (2013). Commentary: Missing heritability, polygenic scores, and gene-environment correlation. *J Child Psychol Psychiatry*, 54(10), 1147-1149.
- Plomin, R., DeFries, J. C., Knopik, V. S., & Neiderhiser, J. M. (2016). Top 10 replicated findings from behavioral genetics. *Perspectives on psychological science*, 11(1), 3-23.
- Scarr, S., & McCartney, K. (1983). How people make their own environments: A theory of genotype→ environment effects. *Child development*, 424-435.
- Wang, X., Guo, X., He, M., & Zhang, H. (2011). Statistical inference in mixed models and analysis of twin and family data. *Biometrics*, 67(3), 987-995.