Rethinking Nature and Nurture in Education

Jack Marley-Payne

The is the penultimate draft of this paper. For citations, please use the published version

Abstract: This paper evaluates how our understanding of natural talent affects questions of educational justice. We argue that education debates currently suffer from a naïve understanding of "nature vs. nurture" and present a more rigorous approach that allows us to see what is required for fair treatment of students. As it stands, there is controversy over the extent to which a student's achievements are a result of their natural talent, as opposed to the quality of their education. For those on the nature side of the debate, students who are performing poorly just are not cut out for academic life, and redoubling efforts at furthering their education is both a waste of resources for society and pointless torture for the student. For those on the nurture side, failing students are a sign of a failing education system. To make progress on educational justice, we must move away from attempts to simply assign proportional influence to natural talent and education respectively, and instead look at the broader landscape of individual reactions to a range of educational environments.

Keywords: Education justice; innate intelligence; heritability; nature and nurture; interaction effects

Introduction

At the centre of many disagreements regarding formal education is a version of the nature vs nurture debate. When it comes to the abilities that education programs aim to develop, there is controversy over the extent to which a student's achievements are a result of their natural talent, as opposed to the quality of their education.

The answer to this question has profound implications for educational practice, especially when it comes to low achieving students. For those on the nature side of the debate (henceforth *nativists*), students who are performing poorly just aren't cut out for academic life, and redoubling efforts at furthering their education is both a waste of resources for society and pointless torture for the student – instead such students should switch to a vocational training that better aligns with their natural abilities. For those on the nurture side (*environmentalists*), failing students are a sign of a failing education system. Therefore, not attempting to do better would be a great injustice to the students who are being let down by the current situation.

Environmentalism is the consensus position of many parts of the academic education establishment, and some may view the debate as already settled. The National Council of Teachers of Mathematics (NCTM), for example, distinguishes between "productive" and "unproductive" beliefs about student ability as follows: [1]

> **Unproductive beliefs**: Students possess different innate levels of ability in mathematics, and these cannot be changed by instruction. Certain groups or individuals have it while others do not.
>
> **Productive beliefs**: Mathematics ability is a function of opportunity, experience, and effort—not of innate intelligence. Mathematics teaching and learning cultivate mathematics abilities. All students are capable of participating and achieving in mathematics, and all deserve support to achieve at the highest levels. (NCTM 2014, p 63)[2]

Note, in particular, that the statement that mathematics ability is not a function of "innate intelligence" is an unequivocal rejection of nativism. Instead the differences in achievement are attributed to students being given different learning opportunities.

Despite this, outside of such circles, nativism persists and undoubtedly plays a role in shaping education policy and practice. Perhaps most significantly for students, it is likely that many teachers subscribe to an informal nativism, viewing some students as "bright" and others as lacking in ability.

Notably, influential philosophers of education, from the Ancient Greeks to the present day, have endorsed versions of nativism. In the Republic, Plato notoriously gave the analogy of people being born with either gold, silver, or bronze in their souls, to represent differing levels of innate intelligence. In the tradition of Rawls (1971), many modern political philosophers emphasize the importance of equality of *opportunity* in education: a system that ensures that those with the same talent applying the same effort will achieve the same academic outcomes.

---

[1] Labelling these "productive" and "unproductive" beliefs, suggests the emphasis is on what beliefs lead to effective teaching. The goal is for teacher attitudes to promote student *self-efficacy* (Bandura 1994). It is clear, though, that the authors also view the productive beliefs as true and unproductive as false. It would be highly problematic if they were instead pushing for the spread of "useful lies" in teacher training.

[2] For further support of this perspective see: Boaler & Staples (2008); Ellis (2018); Gutiérrez (2013); National Research Council (2009).

The assumption, though, is that people differ greatly in talents, so that equal effort and equal education do not necessarily lead to equal achievements. As Nagel (1973) starkly puts it: "When racial and sexual inequality have been reduced, we shall still be left with the great injustice of the smart and the dumb, who are so differently rewarded for comparable effort." In *Illusions of Equality*, Cooper (1982) describes and advocates for the kind of inegalitarian education system he believes nativism to justify.[3]

This issue skirts close to the well-trodden controversy over the relationship between genes and intelligence. The topic was debated extensively in the 1970's, primarily though the exchange between Arthur Jensen and Richard Lewontin, with the former arguing that there is a significant genetic component to intelligence and the latter forcefully disagreeing.[4] However, the issue has proven enduringly relevant, most notoriously due to prominent thinkers giving credence to the idea that different races have different levels of innate intelligence – and, specifically, that Black people are "genetically disadvantaged" (Hernstein & Murray 2010). Alongside this, though, the field of *behavioural genetics* has advanced significantly over recent decades, providing a sophisticated statistical interpretation of the claim that intelligence is due to a person's nature, that it is "heritable".[5]

It should be noted that our concerns are distinct from this specific debate: questions of educational achievement and educational equity do not precisely correspond to questions surrounding genes and intelligence. The goals of education go beyond intelligence, as measured by IQ score: even if one rejects the idea that intelligence is genetically determined, so education may play a role in developing it, it's clear that education also aims to develop a broader range of abilities than those assessed on an IQ test. Further, what counts as a student's nature in educational contexts may go beyond their genes. To move forward, it's necessary to better understand what's being disputed in this particular nature-nurture debate. Though much of what follows draws upon points made in debates over intelligence, the consequences differ in our context, with implications for education that are currently underappreciated.

---

[3] It should be noted that a number of philosophers of education have also pushed back against nativism (White 1974, Ryle 1974, Winch 1990).
[4] See, for example, Jensen (1970) and Lewontin (1974). For an overview of the broader debate, see the essays collected in Block and Dworkin (1976).
[5] See, e.g., Plomin et al (2008) for an introduction to the field.

A better understanding of the relationship between nature, nurture and academic ability changes how we see the problems facing the education system as it currently stands, and the possible paths forward. The focus of this paper is on the relevance of a more sophisticated understanding of *interaction effects* between nature and environment. We will see that traditional versions of both nativism and environmentalism presuppose an overly simplified view of how either nature or nurture influence academic achievement. However, the environmentalist contention that the current system may be failing low-achieving students is validated, and our improved theoretical tools will put us in a better position to address these failings.

Traits: Intelligence and Academic Ability

To begin, we must clarify what things, generally speaking, are eligible to be a product of nature or a product of nurture. These are commonly referred to as *traits* (or more esoterically "phenotypes"), and are properties of organisms such as height, weight, and intelligence. For different organisms, a given trait may take different *values*: different people have different heights, different weights, and different levels of academic ability. The value of such traits can be represented on a numerical scale – e.g. height in inches, weight in pounds, or, more controversially, intelligence in IQ score.[6]

We will refer to the trait we are concerned with as "academic ability", stipulating that this refers to the abilities that education aims to develop. We'll discuss what this entails in more detail below, but for now we can safely assume it involves the kind of skills that are used in core academic subjects. As mentioned above, such a trait is closely related to, though distinct from, intelligence. Neither trait is as straightforward to define or quantify as physical traits like height and weight, and it is worth pausing to consider their potentially problematic features. The claim that intelligence is a unified trait, suitable to be counted as a product of either nature or nurture has been subject to philosophical critique. An exchange between White (1974) and Ryle (1974) notes that conceptually speaking, intelligence involves a range of behavioral dispositions – for example, those involved in solving an equation intelligently or debating theology intelligently.

---

[6] This is easiest when the possible values of a trait form a well-ordered continuous scale, but it can be extended to other cases. Take the trait of "handedness": this can have three discrete values – right-handed, left-handed, ambidextrous. We can represent these values with the relevant quantity of natural numbers {1, 2, 3}. Or consider a trait like "freckle possession"; this has only two possible values: present or not present. These can be represented as {1, 0}.

A further problem is the inadequacy of IQ tests as a means of quantifying intelligence. Despite the claims of the test creators, one can train for the test in order to boost one's score (Te Nijenhuis et al. 2001). Effective training focuses on familiarizing oneself with the rules and quirks of the test, instead of enhancing more fundamental cognitive abilities. Therefore, IQ score is partly a reflection of test taking skill rather than broader property of intelligence.

A final complication is that intelligence is plausibly not an intrinsic property like height or blood type, but instead determined in part by external social factors such as the kind of cognitive abilities that are socially valued, and viewed as difficult and important. Intelligence is closely related to what Hacking (2007) refers to as *kinds of people* – for example, intellectuals and experts. Crucially, this can lead to *feedback loops* as our understanding of what constitutes that type of person informs our understanding of the property itself, so that our concept of intelligence changes to fit the image of university professors and television pundits.

These concerns carry over to academic ability. This trait also covers a range of areas: achievement in mathematics, reading, writing, and science; written and verbal skills; individual and group work – that is, completing academic tasks both while working individually and as part of a group. Also, the tools of measurement are similarly flawed – standardized tests are imperfect measures and test taking ability may be coached (Kulik et al. 1984). Finally, academic ability is again determined in part by external social factors – for example, what subject matters are considered "academic" and make it onto the curriculum.[7]

The response to these criticisms in the case of intelligence is instructive. Despite conceptual distinctions, empirical evidence shows that there are a range of cognitive behaviors that are correlated with each other and well predicted by IQ score (Sternburg et al. 2001). Therefore, the idea of intelligence is a useful theoretical construct, even if IQ score is an imperfect measure. Further, though the nature of intelligence may be contingent and influenced by social circumstances, it still may be relevant to understanding various social questions – indeed it may shed light on injustice.

Similar considerations apply to academic achievement – despite being disparate in nature, there is evidence of strong correlation between achievement level in different academic areas

[7] As an anonymous reviewer notes, in certain social conditions it may not make sense to ascribe academic ability all – consider a Paleolithic human, for example. As we will see below, developments in behavioural genetics show how we can make sense of a socially dependent trait being a product of nature (relative to an appropriate population).

(Shakeshaft et al. 2013). Further, though score on academic examinations is an imperfect measure, and the nature of academic achievement is liable to morph in response to societal changes, it may be useful for understanding education in our society and addressing questions of educational justice.

Genes and Nature

With the relevant trait identified, we need now to look at what it means for a trait's value to be due to nature or nurture. For example, is it due to nature or nurture that John weighs 150lb? To answer this question, we must first clarify what we mean by "nature" and what we mean by "nurture".

The intuitive answer to the first question is that "nature" refers to what a person has at birth, while "nurture" refers to the conditions in which a person is raised. This is often understood such that "nature" refers to a person's *genes*, while "nurture" refers to their *environment.* So, our previous question can be understood to mean: Is it due to John's genes or due to his environment that he weighs 150lb?

What constitutes a gene is itself a complex issue, so equating nature with genes does not eliminate ambiguity. Genes have been defined in terms of multiple features, which are in tension with each other. As Gayon (2016) puts it, according to classical genetics, "the gene is simultaneously a unit of function and transmission, a unit of recombination, and of mutation. Until the early 1950s, these concepts of the gene coincided. But when DNA was found to be the material basis of inheritance, this congruence dissolved."

With the discovery of DNA, the functional view took precedent, identifying genes with the sections of the DNA molecule that code for a specific protein – these proteins were thought to be what give people their traits. This neat picture has proved an over-simplification. Gene functioning is much more holistic than this picture suggests – there is not a one-to-one correlation between genes and traits created. The effects of an individual gene are dependent on the broader genetic architecture it is situated within. For example, the same gene that controls wing placement in flies, controls rib placement in mice (Ridley 2002, p. 31).

One reason for this is that there are parts of the DNA sequence that do not code for proteins and so do not counts as genes by the traditional definition but instead regulate the function of protein-coding genes. So called non-coding DNA creates RNA molecules that regulate the functioning of other genes directly, rather than via a protein molecule. Indeed only 1.5% of our DNA codes protein while 70% codes RNA. Additionally, even some protein coding genes contain sections of DNA, *introns*, that do not correspond to any part of the resulting protein. These introns were previously thought of as "junk DNA", irrelevant to development. However, they too have been shown to play are role in how genes function (Keller 2000, p. 58).

Further, the rapidly developing field of epigenetics shows how factors external to the DNA molecule play a role in processes that were previously thought to be the sole domain of genetic function. Two of the central phenomena here are DNA methylation, where smaller molecules attach to the DNA in specific patterns that modify genetic functioning and histone modification, which alters the structure of the protein scaffolding that holds the DNA molecule in a specific shape. Both of the phenomena possess three key features: they may have a significant effect on trait development; they may be caused by environmental stimuli such as stress but they persist after the stimuli has gone; they may be passed down from parent to child (Carey 2012).[8]

Finally, even further afield are processes that Jablonka (2012) refers to as "soma to soma transmission" – ways in which traits may be passed down across generations outside of gametic transmission (roughly, what is present in the sperm and egg). This includes substances passed through breast milk and placenta and even socially learned behaviors.

I think the right conclusion to draw here is that there is no single distinction that serves as the basis for all nature-nurture debates. Instead, there are a whole family of questions asking about the extent to which a trait is *malleable* at a given stage of development – be it conception, birth, or adolescence. No single one of these distinctions is privileged as being the only one worthy of scientific or philosophical investigation. We must choose which question is pertinent to our particular inquiry. For an evolutionary biologist, a crucial concern is what can be stably transmitted across multiple generations, so it may make sense for them to include DNA but exclude many epigenetic properties from their conception of nature. In contrast, for a developmental biologist, it may make sense to include everything present in the zygote, while in many practical circumstances we will include everything a person has at birth.

Part of our project, therefore, will be determining how to draw the nature-nurture distinction for questions in education. It is crucial to keep in mind that our conclusions will apply

[8] See Jablonka (2012) for an index of epigenetic effects found to be transferred across generations

only to the concerns of educators and education policy, despite the fact there may well be a range of contexts in which traits of academic achievement are also of interest – for example, developmental psychology or healthcare policy.

Giving Nature and Nurture their Due

The question we asked above was: Is a trait's value *due to* nature or *due to* nurture? In other words: Is it *due to* John's genes or *due to* his environment that he weighs 150lb? Articulating what it means for a person's weight to be due to their genes or their environment takes some work. This is causal language: in the words of Elliot Sober (1988) we are "apportioning causal responsibility" for a trait's value.

When asking whether something is due to A or B, the simplest way to answer is to look at what happens if A occurs without B and compare it what happens if B occurs without A.[9] Suppose Jean and Sara are both bakers. We want to know whether the baking of a particular cake was due to Jean or Sara. If Jean baked the cake by herself while Sara played on her phone, then hypothetically removing Sara does not affect the baking of the cake, while hypothetically removing Jean eliminates it. Thus, we can say that it was due to Jean that the cake was baked.

Sometimes this does not work, though. Suppose that in baking the cake, Jean measured out the ingredients while Sara mixed them together. If we remove Jean from the scenario, we have no ingredients available and so no cake; if we remove Sara, the ingredients are not combined and so there is no cake once again. Therefore, responsibility for the cake can be assigned to neither Jean nor Sara alone.

The nature-nurture question is a similarly tricky case since genes and environment are both *essential* to the development of any traits whatsoever. If a strand of human DNA is left in a petri dish, it will not develop at all: genes, or even a fertilized egg, require first an appropriate womb, and then food and a stable environment to develop. Similarly, you can take a stuffed animal (or any other object lacking human DNA) and give it a perfect human upbringing to no avail, in terms of human trait development.

If neither nature nor nurture alone is sufficient to produce a given trait, one may be tempted to simply say that it is a result of *both*, and view that as the end of the matter. Indeed, this has often been proposed as a "solution", purporting to end the debate once and for all.

[9] See Lewis (1973).

However, even if both have a role, we still might want to know the relative importance of their contributions.

Return to Jean and Sara: suppose they work together to bake a batch of 100 cakes. Responsibility for the creation of this batch goes to both of them. However, each baker works separately in their own kitchen, with Jean baking 65 cakes and Sara baking 35. In this case, we can safely say that 65% of the batch was due to Jean while 35% of the batch was due to Sara.

Unfortunately, this is not applicable to the nature-nurture issue. It's not the case that John's genes created 100lb out of his total weight and his environment created 50lb: as we said, if you remove either genes or environment, you get no weight whatsoever.

We can find a more apt analogy if we consider a case where Jean and Sara work together to bake a batch of 100 cakes, with Jean measuring all the ingredients and Sara mixing them – if you remove either you get no cakes at all. In such a case, we cannot assign specific cakes in the batch to Jean, and specific cakes to Sara. However, that doesn't mean there is nothing more to say about how to assign responsibility for the batch.

Suppose we learn that an *average* pair of bakers will produce a batch of 80 cakes in a day – 20 less than Jean and Sara's total. In this context, we can ask who is responsible for the increased yield. It turns out that Jean is just average at measuring ingredients while Sara is an unusually efficient ingredient mixer. If Jean had been replaced by another competent baker, they and Sara would still have produced 100 cakes; whereas, if Sara had been replaced by a competent baker, they and Jean would only have produced 80 cakes. In these circumstances, we can say it is Sara who is responsible for the batch being larger than average.

Translating to the case of nature and nurture, if John weighs 15lb more than average, is this due to his genes or his environment? To answer this using the method above, we would have to take someone with John's genes and raise them in an "average environment" and compare their weight to that of someone with "average genes" raised in John's environment. The problem here is that, as the scare quotes indicate, it's not clear how to make sense of a person with "average genes" or raised in an "average environment".

Population Variance

Note that even in the case of Jean and Sara's cake baking, we did some fudging in order to arrive at a determinate answer. Namely, we implicitly drew upon a background population of bakers when talking about an "average" baker. We need to specify whether we are talking about professional bakers, amateur bakers, or the whole human population (many of whom have no baking experience at all), since presumably average production will vary dramatically across these different populations.

We will stipulate that our previous figures applied to professional bakers – therefore, a pair of professional bakers will on average produce 80 cakes in a day; an average professional measurer paired with Sara mixing will produce 100 cakes, while an average professional mixer paired with Jean measuring will produce 80 cakes.

However, now assume that a pair of average *amateur* bakers working together will produce a batch of 50 cakes in a day. In this case we must assign responsibility between Jean and Sara for producing an additional 50 (rather than 20) cakes. Suppose that if Jean were paired with an average amateur measurer, they would have produced 80 cakes; therefore, Jean's presence increases cake production by 30. If Sara were paired with an average amateur mixer, they would have produced 70 cakes; therefore, Sara's presence increases cake production by 20.

In the context of a background population of amateur bakers, Jean and Sara both contribute to the increased production, with Jean's contribution being the greater one. In contrast, when we were concerned with professional bakers, the extra production could be attributed entirely to Sara.[10] What this tells us is that assigning causal responsibility in this way must be done *relative to a target population*. This should not be a surprise, as similar concerns crop up when analysing all manner of causal notions.[11]

We can talk about causal responsibility relative to a population more precisely. In our example, we have a population of baking duos (measurer and mixer) and an output value of interest (cakes baked in a day). This gives us a comprehensive background of information – the output value taken for each possible combination in the population. This can be used to answer questions about casual responsibility.

Our output value, cakes baked, varies across our population. We want to know how much of the variation is due to changing the value of our first population input (i.e. the measurer) and how much is due to changing the value of the second population input (i.e. the mixer). To

[10] Why might this be so? These figures would be explained by a story in which professional bakers were much better than amateur bakers at mixing ingredients but no better than them at measuring; Jean is as good as an average professional at mixing, while Sara is *better* than an average professional at measuring.
[11] See Schaffer (2005) for further discussion.

calculate this, we use a statistical tool called ANOVA (analysis of variance). Roughly speaking, this tool looks at the total variation in output and subdivides it into the variation that occurs when the first input is held fixed and the second input varies, and that which occurs when the second is fixed and the first varies. The relative size of these two components is used to assign responsibility for overall variation in the population to the two inputs proportionally.

The details on this mathematical function are not necessary for our discussion; we just need to know that the results are highly dependent on the background population. [12] It could say, for example, that in a population of professional bakers, 2/3 of the variation in production is due to the measurer, while 1/3 is due to the mixer; however, in a population of amateur bakers, 1/4 of the variation in production is due to the measurer, while 3/4 is due to the mixer.

This same approach has been employed to tackle nature-nurture questions. In this case we have a population of gene-environment pairs (AKA people, or other organisms) with different values for a given trait. The field of behavioural genetics uses statistical techniques based on this fundamental idea to apportion causal responsibility for variation of the trait's value within the population – for example, that 1/4 of variation in weight is due to environment and 3/4 due to genes. The value assigned to variation due to genes is referred to as a trait's *heritability*. It is important to note that heritability only ascribes responsibility to genes across the entire population; it does not say whether an *individual's* value is due to their genes or environment.[13]

A complication here is that calculating relative contributions of genes and environment using the methods sketched above requires a matrix of data with values for each genotype-environment combination. In humans, obtaining such data is impossible for both moral and practical reasons since it would require selectively breeding humans and placing them in controlled environments for experimental purposes. Alternative techniques have been developed, therefore, that allow for heritability estimates using the data that is available. Most prominent among these are twin studies using the "ACE model".[14] This model looks at a population of twins, both identical (monozygotic) and non-identical (dizygotic). It compares the variance between identical twins to non-identical twins to get an estimate for heritability.[15] These models

---

[12] For such details see, for example, Blalock (1960, ch. 16).
[13] Sobel (1988) argues that the approach can be extended to individuals.
[14] See Plomin et al. (2008, Appendix) for more details.
[15] Since identical twins share 100% of genetic material and non-identical twins share, on average 50%, the model assumes the difference in variance between the two groups is ½ of the total variance due to genes.

have found that both intelligence and academic achievement are highly heritable – usually attributing greater than 50% of variation to genetic factors.[16]

One can find fault with the empirical assumptions of the model – for example, whether one can safely generalize findings about twins to the population as a whole. Perhaps, one should not read too much into *precise* heritability figures for this reason. However, related studies – such as those looking at other kinds of relatives or adoptees – make the general conclusion that academic achievement is highly heritable hard to avoid. In any case, this will not be our concern. We will be looking at what implications a high heritability figure for academic achievement has for educational justice; whether such a result, as is sometimes claimed, implies that there is little to be done to help low achieving students.

Individuating Environments

Assigning causal responsibility to nature and environment for a person's traits requires that we have a reliable way of *individuating* each of the two inputs. That is, when we look at two different people, we must be able to say whether they have either the same nature or the same environment. The ACE model for estimating heritability builds in certain assumptions about how to do this, but these assumptions may not be appropriate for our purposes.[17]

On reflection, there's some ambiguity as to whether two people are raised in the same environment in the appropriate sense. For a start, it is impossible for two people to be raised in the *exact* same environment – even siblings will sleep in different beds and eat different meals off different plates. In the philosophical sense, two people cannot be raised in the same *token* environment, so we have to look at whether they are raised in the same *type* of environment.[18]

One issue here is how specific we are. The maximally specific option is that two environments must be *duplicates*, or indiscernible, to be of the same type – something that is conceptually possible, but not within our abilities to create.[19] We are going to want something more inclusive, so that being of the same environment type requires, say, getting the same level of nutrition, but not necessarily the precise same sequence of meals.

---

[16] See Shakeshaft et al. (2013). Johnson et al. (2009) note that twin studies in fact find significant heritability for virtually all human behavior investigated.
[17] Roughly, the model assumes that identical nature means identical DNA sequence while identical (shared) environment means being raised in the same family.
[18] This distinction is made, for example, in Davidson (1980).
[19] For an analysis of duplicates, see Lewis (1983).

A further problem is the extent to which the conditions for individuating environments are *indexical*. For example, suppose we want to know whether childhood footwear affects running ability as an adult. Assume that two children grew up wearing indiscernible shoes right down to the shoe size. This might seem the highest similarity possible. However, the one child has bigger feet than the other, and so for them the shoes are far too small while for the other they are a perfect fit. If the one child grows up a better runner than the other, it would be terrible reasoning to conclude that since they were raised in the same environment the difference must be genetic. For environments to be of the same type, they must meet certain *indexical* criteria, relating to the relationship between environment and individual. In this example, for footwear environments to be the same, the shoes must fit the children to comparable degrees.

This issue is highly relevant in the case of education. It might seem, for example, that going to the same school at the same time is the strongest criterion possible for sharing an educational environment. However, suppose that the school in question is a majority White school in the US and the first child is White while the second child is Black. Indexically, these environments are very different.

How we assign causal responsibility to genes and environment depends, therefore, not just on how we set boundaries on the population (which educational environments are ruled out as unpractical to administer, for example) but also on how we divide up the members of the population into types.

A related issue is that some individual-environment pairs are *incompatible,* in that it isn't possible for a given individual to be situated in a given environment. Individuals *respond* to their environments in various ways generating feedback loops. We may want to individuate environments in terms of a specific kind of feedback loop and, for individuals who don't engage in that kind of feedback, such environment types are not applicable.

Consider symbiotic relationships between plants or animals, such as the association between a green plant and a fungus known as mycorrhiza. The plant supplies molecules such as sugars to the fungus, and the fungus provides water and mineral nutrients to the plant.[20] When considering environment types for fungi, it might be necessary to specify features of this symbiotic process such as what nutrients are transferred between the two. When considering

[20] Ainsworth, (2008)

another type of entity, such as moss, these distinctions between environment are not applicable since there is no symbiotic process in the first place.

Analogous concerns apply with our motivating example since there is a great deal of interplay between an individual and their education environment, as discussed above with feedback loops involving intelligence and academic ability. Consider different teaching techniques for quiet students, various ways to encourage them to talk when they are reluctant to do so. This might include asking them to tell a small group first, asking them to read aloud what they have written down, or changing the topic to find something the student is comfortable talking about, and helping them connect their successes to the more intimidating setting. For a talkative student, the initial situation in which they are reluctant to talk does not occur, and so it is not possible for these kinds of back and forth between teacher and student to take place.[21]

The goal in raising these complications is not to come up with a comprehensive way of resolving all problems with individuating environments that arise from indexical factors and individual-environment incompatibility. Instead, it is to show how the notion of sameness of environment is somewhat vague and ambiguous in the current context, and must be treated accordingly.

A final question that must be addressed is what subject matters are relevant for us when individuating environments. Recall the purpose of asking the nature-nurture question about education: we are trying to decide the extent to which we can and should help underachieving students. The changes we are considering concern their formal education. Therefore, we want to individuate environments on the basis of whether they involve the same kind of formal education.

This means that we are ignoring a large amount that we usually think of as a person's environment – i.e. everything that goes on before and outside of school. If two students attend the same school but had a very different upbringing in their pre-school years, they may have the same *educational* environment, but not the same environment in a broader sense – certainly we cannot be sure that all differences between them in academic achievement are a result of their genes.

As we mentioned above, nature and nurture are assumed to be exhaustive in explaining traits. However, educational environments and genes clearly are not so. Does that mean such a

[21] For more on strategies encouraging student participation see, e.g., Turner & Patrick (2004).

restrictive sense of environment is impermissible? Not necessarily, since we can also plug the gap by expanding our notion of nature instead. As we discussed above, there is no single nature-nurture distinction that is universally applicable, and how we draw the line depends on the specific problem at hand. From the perspective of teachers and education decision-makers, it is irrelevant whether academic ability is a result of genes or some other pre-school factor. Instead, the pertinent question is the extent to which academic achievement is malleable at the point at which a child begins their formal education.[22]

Therefore, our population is going to be made up of <*pre-school child*, *formal education*> pairs – with the output being the academic achievement level a given pre-school child reaches after undergoing that formal education. For ease of expression, we will refer to a pre-school child as an "individual" and formal education as an "environment". Though this is a long way from the traditional gene-environment distinction, it has the same formal structure, and is a more faithful articulation of the vague nature-nurture talk surrounding debates in education.

Interaction Effects

Now we have specified the individuals, environments and trait that constitute our population, we can turn to the key issue: interaction effects. That is, how different individuals may respond differently to different kinds of educational setting. The significance of interaction for nature-nurture questions was famously suggested by Lewontin (1974) and has been at the centre of debates over the relationship between genetics and IQ ever since.[23] Less well-appreciated is the specific relevance of interaction effects for issues in education. To illustrate this, we will look at a simplified situation which considers just two individuals, A and B: A is a high achiever in a typical US school setting, and B is a low achiever in such a setting. Our concern is how outcomes might vary across different educational environments. We will use a series of graphs to illustrate some possibilities.

[Insert table of figures]

---

[22] It's important to note that this distinction is only irrelevant in the context of *education* policy, not in the context of policy more broadly. If, for example, nutritional quality during infancy has a significant effect on educational outcome, justice may demand instituting a policy that ensures all children receive high quality food (which is not to say justice doesn't demand this regardless).

[23] See Tabery (2014) for an authoritative history of how ideas about interaction have contributed to the nature-nurture debate.

Figure 1 represents the most extreme nativist view: it tells us that the two individuals' levels of academic achievement are the same, no matter the educational environment – and student A (the solid line) is a fixed amount above student B (the dotted line). There simply is no helping the low-achieving child. Of course, a view this extreme is highly implausible – educational environment is not completely epiphenomenal. For a start, crude modifications such as adding or removing a year of schooling undoubtedly affect academic achievement.[24]

Figure 2 represents the most extreme environmentalist view: it tells us that in any educational environment, the two students have the same level of achievement. This too is implausible: for any reasonable way of individuating environments, different students reach different levels of academic achievement in the same educational environment.

Figure 3 represents a more moderate, and more plausible, nativist position. It allows that education environment does affect outcomes. We can view the x-axis as representing quality of education, moving from lower quality to higher quality: the greater the value of the x coordinate, the greater the value of the y-coordinate for both high-achieving and low-achieving students. This view is more moderate in that it allows that a low achieving student can perform significantly better with an improved education. What makes it nativist is that the high achieving student's achievement *dominates* the low achieving student's: across all environments, the former does better than the latter. This is one reasonable elaboration of the idea that the former is *naturally* better than the latter.

Some variations on this position are displayed in figures 4 and 5, illustrating other cases in which one student's achievements dominate the other's. Significantly, they incorporate the idea of diminishing returrns – the higher a student's achievement level is in a given environment, the harder it is to improve it further. In figure 4, there is a fixed difference in achievement between the two students, while in figure 5, both approach the same limit (further, all sorts of in-between cases are conceivable).

Scenarios like these imply that whatever you do, the high achieving student will always do better than the lower achieving student, so this achievement gap is not in itself an educational failure or cause for concern. However, there are important differences between the variant cases in terms of what level of achievement gap is inevitable and what level of change a shift in education environment can bring about.

[24] See, e.g., Husén & Tuijnman (1991) or Ceci (1991).

Figure 6 presents a more plausible environmentalist position than figure 2, since it allows that different children have different achievement levels within the same education enviroment. What makes this non-nativist is that neither student dominates the other. In some environments, student A does better, and in other environments student B does. In this scenario, there is not a universal scale for educational environments from low quality to high quality – an environment that is highly effective for one student, may be ineffective for another. For example, some students might learn better with individual study from books, while others learn better from structured group discussion or practical activities.

This, therefore, displays the aforementioned *interaction effect* between individual and environment. In this situation, unlike the nativist scenarios, the fact that student A is outperforming student B is not inevitable – instead it means that the current educational environment is better suited to student A. Student B could reach similar levels in an environment that suited them.[25] When there is an interaction effect, a given change in environment will result in different changes in outcome for different individuals – e.g., an improvement for one and a decline for another.

It seems, then, that the key disagreement between the nativists and environmentalists is whether students that are currently high achieving *dominate* currently low achieving students across alternative educational environments. A problem for the nativist is that since this kind of domination requires that when two students are in the same environment the one always out-performs the other, it appears an ephemeral notion. This is because it depends on what counts as the *same environment* for the two students. As we saw above, with the example of shoe size, this is a somewhat ambiguous matter. Further, we also noted that there will be individual-envionment pairs that are incompatible, and if there is an environment which is incompatible with student A, but compatible with student B, then it is impossible for A to dominate B.

Is there a version of the nativist position that does not rely on domination? One might be tempted to think a situation represented by figure 7 is a possibility. In this graph, student A does not dominate student B, strictly speaking, but one is tempted to say that they almost do. In only a tiny fraction of the education environments does student B beat student A. It looks like student A does far better than student B on average and thus can be said to have superior natural talent. In

[25] In further exploring these kinds of scenarios, it should be noted that these pairs of reflected lines might well display the various kinds of diminishing returns displayed in the nativist case in figures 4-5.

other words, we use area under a student's achievement line (i.e. the function's integral) as a measure of their natural talent.

For this to make sense, though, there has to be a non-arbitrary way of mapping educational environments onto the x-axis. If we are working with a continuous space of environments, we need to impose a measure on it in order to perform the necessary operation. If we are working with discrete environment types, we need to decide how to divide them up into bins and weight them. For example, we must decide how much weight will be assigned will to predominantly book-learning environments, and how much to predominantly practical ones. How we answer this will affect the relative natural talent scores of "book-learners" and "practical-learners", respectively.[26]

One option, taking a cue from Sobel (1988), is to make weighting proportional to the frequency with which the environment types occur in schools currently.[27] So, if there are twice as many schools with book-learning environments as practical-learning environments, how an individual performs in book-learning environments has twice as much impact on their natural talent score as how they perform in practical-environments. In effect, this makes a student's natural academic ability equivalent to their *expected* academic achievement if they are enrolled in a randomly selected school.

Equity and Education

This expected value driven approach may make sense in a descriptive setting, such as if we are evaluating plant growth in the wild. It does not, however, make sense in a normative context, where we are looking at making improvements. As Lewontin (1974) notes, this kind of analysis does not tell us what would happen if there were an *intervention*. Suppose practical learning environments are currently rare, but that there are a large proportion of individuals who are practical learners – these would count as naturally low achievers using the measure under consideration. However, an intervention that made sure practical learners were provided with a practical learning environment would increase the achievement of these individuals. Indeed, after

[26] Note that we are talking about different types of learners hypothetically here and are not making any empirical commitments as to the existence of specific learning styles. We will discuss this issue below.
[27] In effect, this is what heritability estimates do, though they may weight results to control for various demographic factors.

this intervention, since a larger proportion of schools would now have practical learning environments, these students would now be classified as naturally high achievers!

The measure of natural talent under consideration leads to a pernicious and farcical cycle of reasoning. Students who learn best in an environment that is currently rarely employed will tend to be low achievers. Because most learning environments are bad for such students, these students will be classified as *naturally* low achievers. If they are naturally low achievers then they are not being failed by the current system, and therefore there is no need to expand access to the learning environments in which they would succeed. However, the idea that students could do much better (and better relative to their peers) in a different learning environment is precisely what we were concerned with when we asked whether we were failing students.

In the context of educational justice, what matters is a student's high point, not their average. In a fair education system, students should be provided with the optimal education environment, out of those we can reasonably be expected to provide. Crucially, a lot turns on what education environments we can "reasonably be expected to provide". For example, a student might perform better with a full-time private tutor than within a regular classroom environment, but this is not necessarily something fairness demands be provided. Certainly, there are not the resources available to give *every* child a fulltime private tutor.

From the nativist perspective, the private tutor example is especially instructive. As we have seen, nativism makes more sense when there is a uniform scale of quality of education and, presumably, increasing quality requires increasing expenditure. Spending more on any student's education would improve their academic achievement. However, on this view, high-achieving students dominate low-achieving students, so for a given level of education expenditure, the former will always outperform the latter. If one were to distribute resources evenly across all students, there would be significant differences in achievement levels.

Within such a scenario, one could potentially distribute resources such that all students ended up with the same level of academic achievement. However, doing so would require dedicating far more educational resources to the low achieving individuals than the high achieving individuals. This might seem unfair to the higher achieving students since they are not being given sufficient resources to flourish academically.

From a non-nativist perspective, however, things look very different. There is not a uniform scale for quality of education since different individuals learn better in different kinds of environment. Therefore, there is not a straightforward association between better education for an individual and higher education expenditure. Providing an education environment that fits with how an individual learns most effectively is not giving them preferential treatment – it is something we could do for every student. So, if we do not for a given student, it is plausible to say we are failing them. All this makes clear that a crucial empirical question for educational justice is in what ways and to what extent interaction effects occur.

Evidence for Interaction

Many nativists are skeptical of the significance of interaction effects when it comes to our understanding of human behavior, especially with regard to intelligence. One reason for this is that, within behavioral genetics, there is a statistical interpretation of interaction that can be measured when estimating heritability: the variable "GxE" (genes times environment). They argue that since heritability analyses have not revealed a significant value for GxE, the issue of interaction can be dismissed. However, the statistical measure only reflects the extent to which different environments *within the current population* are producing differential effects. This means that if the varieties of teaching environment that produce interaction effects are not currently in widespread use, they will not significantly affect the value of GxE.[28]

Nativists are generally unmoved by this point. Jensen accused his opponents of simply speculating that environmental changes no one has tried yet could potentially boost the intelligence of current low achievers (Jensen 1970). Sesardic similarly called the focus on interaction effects "the curious triumph of the *possible* over the *actual*" (Sesardic 2005, 84–85).[29]

Such critics are correct that merely noting the possibility of interaction effects does not on its own conclusively disprove nativism. To move forward, we need evidence that they do, or do not, in fact exist. Until recently, evidence on interaction effects was scarce, and so previous iterations of the debate often descended into a rather unsatisfying back and forth as to where the burden of proof lay.[30] In recent years, though, research into interaction effects has bloomed, and

[28] As Shakeshaft et al. (2013) note, the high level of heritability they find for academic achievement results in part from the highly homogenized education that members of the population (secondary school students in the UK) are receiving – and does not preclude the possibility that alternative teaching approaches would benefit students who are currently under-achieving. For more details on GxE and its role in debates over interaction effects, see Tabery (2014).
[29] Tabery (2014) provides a detailed breakdown of this debate.
[30] See Tabery (2014 pp. 59-60)

a number have been identified in a range of areas, which clarifies the terms of debate for education moving forward. As Tabery puts it:

> Whether it was research on crop yield in plants, meat production in livestock, or behavioral traits in humans, there has been plenty of empirical evidence for interaction, plenty of empirical evidence against interaction, and no warrant for assuming ahead of time whether any particular trait will or will not result from some particular interaction of nature and nurture. (Tabery 2014 p. 8)

Crucially, interaction has been found with regard to intelligence: the so-called Scarr-Rowe effect which finds an effect based on the income level of individuals (Turkheimer & Horn 2014). Further, biologists have presented an evolutionary argument that natural selection would favor producing offspring with differential susceptibility to environmental effects as a hedge against environmental risk. This provides a theoretical basis for why we should expect to find interaction effects for a range of behaviour (Belsky & Pluess 2009).

Therefore, there is good reason to think that there are interaction effects in education to be found. The goal now to is to identify those that can be leveraged in education interventions.

Evidence in Education

Unfortunately, research to date in education is not at the stage one might hope when addressing this question. It is worth noting that one of the ideas in education that is most suggestive of interaction has not been empirically validated. This is the hypothesis that people have different *sensory learning styles* – typically, visual, auditory, read–write or kinesthetic – and that they learn better when working within their preferred sensory mode (Cuevas 2015). However, research to date has found no reliable evidence that interaction effects due to such learning styles exist (Coffield et. al 2004; Pashler et. al. 2009).

Compounding the problem, both education research and practice has failed to respond appropriately to the empirical results. Despite the negative findings, many research articles and teacher training materials report the sensory learning styles hypothesis as a fact, and teachers are instructed to tailor their teaching methods accordingly (Cuevas 2015). In addition to promoting suboptimal teaching methods, this complicates the effort to identify legitimate interaction effects in education – one needs a body of research conducted with reliable empirical methods in order to have a starting point for future studies.

Despite these problems, research does point towards more promising areas for interaction effects in education. Pashler et al (2009), after dismissing sensory learning styles, identify a range of studies comparing structured and unstructured (or self-directed) learning environments. These are found to have differential effects based both on student's prior aptitude level (Janicki & Peterson 1980) and their belief level in their control over their own learning (Horak & Horak 1982). Additionally, it would seem worthwhile to look into whether interaction effects that have been found in genetics carry over to educational contexts. Natural places to look include the aforementioned interaction between income-level and genes with respect to intelligence, differential stress reactions (Caspi et al. 2003), and differential responses to discipline (Bakermans-Kranenburg & Van Ijzendoorn 2006).

Another factor to consider is the potential effects of demographic differences, such as those pertaining to race and gender (Ellis 2018). Within education research, many promising individual teaching interventions have been implemented to address this. For example, Gutiérrez (2013) outlines teaching techniques tailored specifically to minority urban students, Kisker et. al (2012) found approaches that worked with Native Alaskan students, while Razfar et. al. (2011) outline techniques effective for students learning in their second language. Unfortunately, to date, these kinds of education strategies have not been successfully scaled, nor have they been applied in a setting that allows for systematic quantitative analysis.

Conclusion

This paper has aimed to show how a better understanding the relationship between nature and nurture – individual and environment – can further our understanding of education, and educational justice in particular. It has been shown that interaction effects between individual and learning environment are a crucial factor in the extent to which low achieving studies are being failed by the current system. We have seen that if such effects exist, it does not make sense to place students on a single scale of natural academic ability. The empirical research to date makes it plausible to think that this is so; however, education research in particular has not done an adequate job in identifying interaction effects to target in educational interventions.

Hopefully, moving forward, we can move away from the crude nativist and environmentalist viewpoints that are currently endemic in debates in education. An improved understanding of the how individuals respond to their education can guide efforts to achieve a fairer education system.

References

Ainsworth, G. C. (2008). *Ainsworth & Bisby's dictionary of the fungi*. Cabi.

Bakermans-Kranenburg, M. J., & Van Ijzendoorn, M. H. (2006). Gene-environment interaction of the dopamine D4 receptor (DRD4) and observed maternal insensitivity predicting externalizing behavior in preschoolers. *Developmental Psychobiology: The Journal of the International Society for Developmental Psychobiology*, *48*(5), 406-409.

Bandura, A. (1994). Self-efficacy. In V. S. Ramachaudran (Ed.), *Encyclopedia of human behavior* (Vol. 4, pp. 71-81). New York: Academic Press. (Reprinted in H. Friedman [Ed.], *Encyclopedia of mental health*. San Diego: Academic Press, 1998).Block, N. J., & Dworkin, G. (1974). IQ, heritability and inequality, part 2. *Philosophy & public affairs*, 40-99.

Belsky, J., & Pluess, M. (2009). Beyond diathesis stress: differential susceptibility to environmental influences. *Psychological bulletin*, *135*(6), 885. Carey, N. (2012). *The epigenetics revolution: How modern biology is rewriting our understanding of genetics, disease, and inheritance*. Columbia University Press.

Blalock Jr, H. M. (1960). Social statistics. *McGraw Hill.*

Block, N. J., & Dworkin, G. E. (1976). The IQ controversy: Critical readings.

Boaler, J., & Staples, M. (2008). Creating mathematical futures through an equitable teaching approach: The case of Railside School. *Teachers College Record*, *110*(3), 608-645.

Caspi, A., Sugden, K., Moffitt, T. E., Taylor, A., Craig, I. W., Harrington, H., ... & Poulton, R. (2003). Influence of life stress on depression: moderation by a polymorphism in the 5-HTT gene. *Science*, *301*(5631), 386-389.

Ceci, S. J. (1991). How much does schooling influence general intelligence and its cognitive components? A reassessment of the evidence. *Developmental psychology*, *27*(5), 703.

Coffield, F., Moseley, D., Hall, E., & Learning and Skills Research Centre (Great Britain)(LSRC). (2004). Learning styles and pedagogy in post-16 learning: a systematic and critical review.

Cooper, D. E. (1982). Illusions of Equality.

Davidson, D. (1980), "Mental Events," in *Essays on Actions & Events*. Oxford: Clarendon Press.

Ellis (2018). Culturally Responsive Mathematics Teaching.

Gardner CD, Trepanowski JF, Del Gobbo LC, et al. (2018) Effect of Low-Fat vs Low-Carbohydrate Diet on 12-Month Weight Loss in Overweight Adults and the Association with Genotype Pattern or Insulin Secretion: The DIETFITS Randomized Clinical Trial. JAMA. 2018;319(7):667–679. doi:10.1001/jama.2018.0245

Garfinkel, A. (1982). Forms of explanation: Rethinking the questions in social theory. *Yale University Press.*

Gayon, J. (2016). From Mendel to epigenetics: History of genetics. *Comptes rendus biologies*, *339*(7-8), 225-230.

Gutiérrez, R. (2013). Why (Urban) Mathematics Teachers Need Political Knowledge. *Journal of Urban Mathematics Education*, *6*(2), 7-19.

Hacking, I. (2007, April). Kinds of people: Moving targets. In *Proceedings-British Academy* (Vol. 151, p. 285). OXFORD UNIVERSITY PRESS INC..

Herrnstein, R. J., & Murray, C. (2010). *The bell curve: Intelligence and class structure in American life*. Simon and Schuster.

Horak, V. M., & Horak, W. J. (1982). The influence of student locus of control and teaching method on mathematics achievement. *The Journal of Experimental Education*, *51*(1), 18-21.

Husén, T., & Tuijnman, A. (1991). The contribution of formal schooling to the increase in intellectual capital. *Educational Researcher*, *20*(7), 17-25.

Jablonka, E. (2012). Epigenetic variations in heredity and evolution. *Clinical Pharmacology & Therapeutics*, *92*(6), 683-688.

Janicki, T. C., & Peterson, P. L. (1981). Aptitude-treatment interaction effects of variations in direct instruction. *American Educational Research Journal*, *18*(1), 63-82.

Jensen, A. R. (1970). Race and the Genetics of Intelligence: A Reply to Lewontin. *Bulletin of the Atomic Scientists* 26 : 17–23

Johnson, W., Turkheimer, E., Gottesman, I. I., & Bouchard Jr, T. J. (2009). Beyond heritability: Twin studies in behavioral research. *Current directions in psychological science*, *18*(4), 217-220.

Keller, E. F. (2000). The century of the gene.

Keller, E. F. (2010). *The mirage of a space between nature and nurture*. Duke University Press.

Kisker, E. E., Lipka, J., Adams, B. L., Rickard, A., Andrew-Ihrke, D., Yanez, E. E., & Millard, A. (2012). The potential of a culturally based supplemental mathematics curriculum to improve the mathematics performance of Alaska Native and other students. *Journal for Research in Mathematics Education*, *43*(1), 75-113.

Kulik, J. A., Bangert-Drowns, R. L., & Kulik, C. L. C. (1984). Effectiveness of coaching for aptitude tests. *Psychological Bulletin*, *95*(2), 179.

Lewis, D. (1973). Causation. *Journal of Philosophy*, 70: 556–567.

Lewis, D. (1983). New work for a theory of universals. *Australasian journal of philosophy*, *61*(4), 343-377.

Lewontin, R. C. (1974). Annotation: the analysis of variance and the analysis of causes. *American journal of human genetics*, *26*(3), 400.

Nagel, T. (1973). Equal treatment and compensatory discrimination. *Philosophy & Public Affairs*, 348-363.

National Research Council. (2009). *Mathematics learning in early childhood: Paths toward excellence and equity*. National Academies Press.

NCTM (2014) Principles to Actions: Ensuring Mathematical Success for All. *NCTM, National Council of Teachers of Mathematics*

Pashler, H., McDaniel, M., Rohrer, D., & Bjork, R. (2008). Learning styles: Concepts and evidence. *Psychological science in the public interest*, *9*(3), 105-119.

Plomin, R., DeFries, J. C., & McClearn, G. E. (2008). *Behavioral genetics*. Macmillan.

Rawls, J. (1971). *A theory of justice*. Harvard university press.

Razfar, A., Licón Khisty, L., & Chval, K. (2011). Re-mediating second language acquisition: A sociocultural perspective for language development. *Mind, Culture, and Activity*, *18*(3), 195-215.

Ridley, M. (2003). Nature via nurture: Genes, experience, and what makes us human. *New York*.

Ryle, G. (1974). Intelligence and the Logic of the Nature-Nurture Issue Reply to JP White. *Journal of Philosophy of Education*, *8*(1), 52-60.

Schaffer, J. (2005). Contrastive causation. *The Philosophical Review*, *114*(3), 327-358.

Sesardic, N. (2005). *Making sense of heritability*. Cambridge University Press.

Shakeshaft, N. G., Trzaskowski, M., McMillan, A., Rimfeld, K., Krapohl, E., Haworth, C. M., ... & Plomin, R. (2013). Strong genetic influence on a UK nationwide test of educational achievement at the end of compulsory education at age 16. *PloS one*, *8*(12), e80341.

Sober, E. (1988). Apportioning causal responsibility. *The Journal of philosophy*, *85*(6), 303-318.

Sternberg, R. J., Grigorenko, E. L., & Bundy, D. A. (2001). The predictive value of IQ. *Merrill-Palmer Quarterly (1982-)*, 1-41.

Tabery, J. (2014). *Beyond versus: The struggle to understand the interaction of nature and nurture*. MIT Press.

Te Nijenhuis, J., Voskuijl, O. F., & Schijve, N. B. (2001). Practice and coaching on IQ tests: Quite a lot of g. *International Journal of Selection and Assessment*, *9*(4), 302-308.

Turkheimer, E., & Horn, E. E. (2014). Interactions between socioeconomic status and components of variation in cognitive ability. In *Behavior genetics of cognition across the lifespan* (pp. 41-68). Springer, New York, NY.

Turner, J. C., & Patrick, H. (2004). Motivational influences on student participation in classroom learning activities. *Teachers College Record*, *106*(9), 1759-1785.

White, J. P. (1974). Intelligence and the Logic of the Nature-Nurture Issue. *Journal of Philosophy of Education*, *8*(1).

Winch, C. (1990). *Language, ability and educational achievement*. Routledge.

Figures

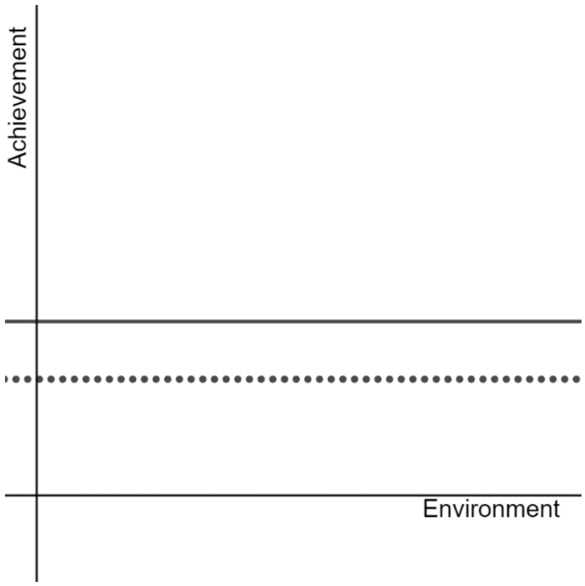

figure 1

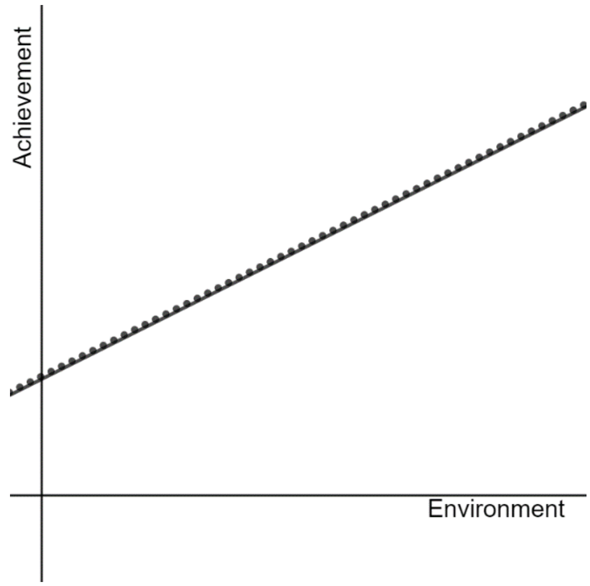

figure 2

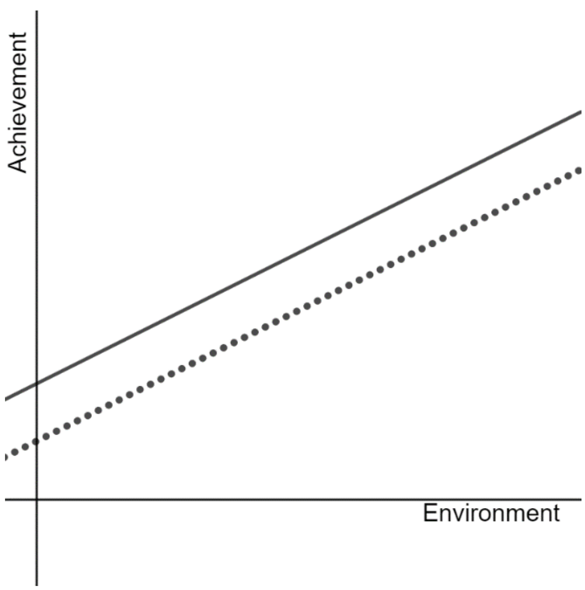

figure 3

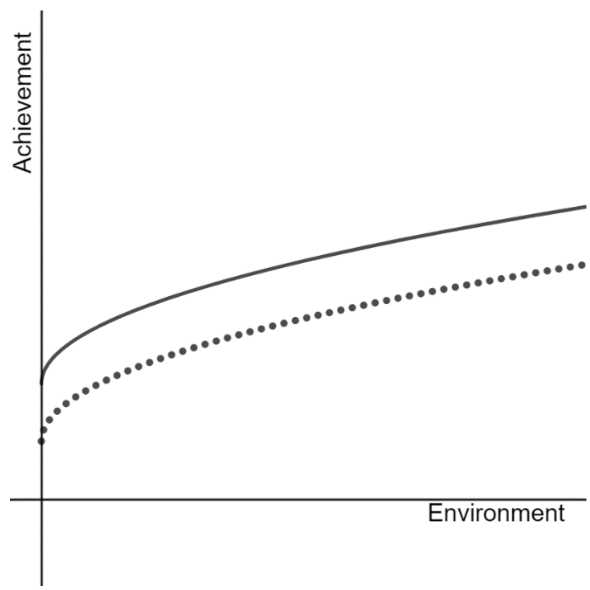

figure 4

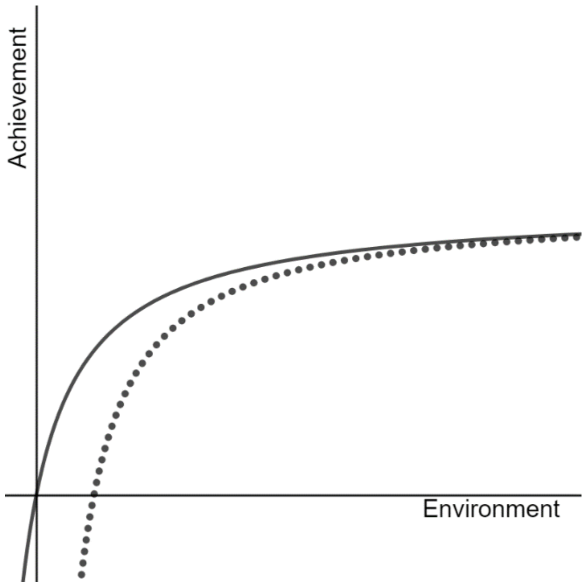

figure 5

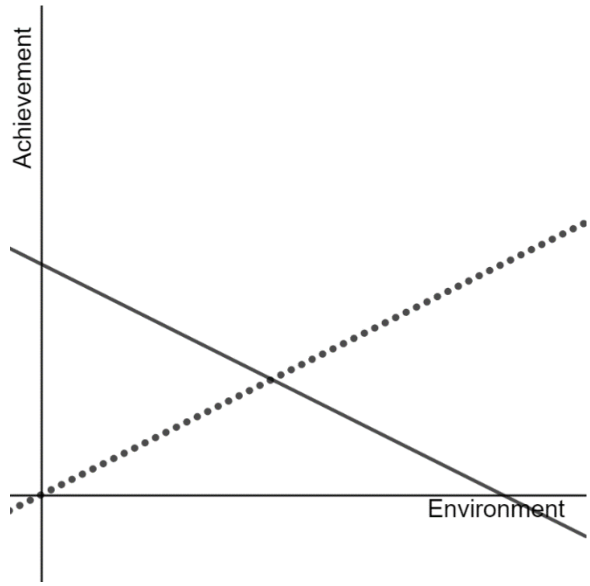

figure 6

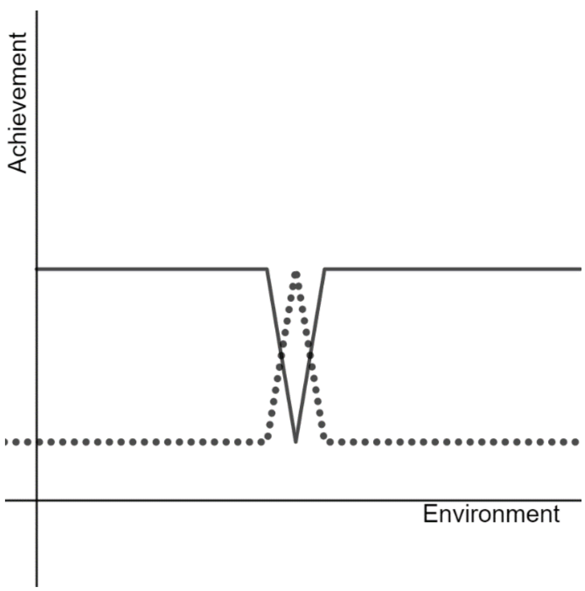

figure 7