# Racial Classification Without Race: Edwards' Fallacy

Adam Hochman

**Abstract**

A. W. F. Edwards famously named "Lewontin's fallacy" after Richard Lewontin, the geneticist who showed that most human genetic diversity can be found within any given racialized group. "Lewontin's fallacy" is the assumption that uncorrelated genetic data would be sufficient to classify genotypes into conventional "racial" groups. In this chapter, I argue that Lewontin does not commit the fallacy named after him, and that it is not a genuine fallacy. Furthermore, I argue that when Edwards assumes that stable classification is all it takes to vindicate racial naturalism, he commits a fallacy of his own: *Edwards' fallacy*. The ability to create a classificatory system, and then reliably sort things within that system, is not enough to make that system scientifically respectable. I show that Edwards' fallacy is rife in debates about the existence of human biological races.

**Introduction**

By demonstrating that most human genetic diversity resides within any given racialized group, Richard Lewontin (1972) provided one of the most influential arguments against the existence of human biological races. Yet according to A. W. F. Edwards (2003), Lewontin was guilty of fallacious reasoning. Edwards argues that Lewontin analyzed genetic data under the false assumption that a locus-

by-locus analysis was sufficient to reveal all of the information necessary for racial classification. Edwards calls this "Lewontin's fallacy."

In this chapter, I argue that Lewontin (1972) does not actually commit "Lewontin's fallacy." That is, Lewontin does not make the assumption Edwards believes he makes (or at least, he does not make it in the 1972 article where Edwards claims he does). Furthermore, I show that it is not a false assumption after all. It has been demonstrated that, given the right sampling scheme, a single-locus analysis is sufficient to classify people into conventional racialized groups (Witherspoon et al., 2007).

I argue that while Lewontin does not make a false assumption, Edwards does. Edwards assumes that the ability to sort genotypes into conventional racialized groups demonstrates that race is real. This is a mistake. The ability to classify genotypes into conventional racial categories is neither necessary nor sufficient to vindicate racial naturalism. Lewontin did not commit a fallacy, but Edwards did.

*Edwards' fallacy* is to assume that the ability to create a stable classification system entails that the classification system is scientifically meaningful. This is false. It is possible to create a stable classification system that does not pick out valid scientific categories. Such classifications must be picking up on some regularity—whether it be in the world or, say, artificially produced by a biased sampling scheme—but not all regularities are scientifically meaningful. Edwards' fallacy is common in the debate about the existence of biological races in humans. It is possible to classify people according to conventional racial categories even though race is not a valid biological category.

**"Lewontin's Fallacy" is not a Fallacy**

Fallacies come in two flavors: formal, and informal. A formal fallacy is invalid because of its form. An informal fallacy is invalid because of its content. 'Lewontin's fallacy' is an informal fallacy. When it comes to informal fallacies, there is always the possibility that they will turn out not to be fallacies at all. It might seem as though there is something objectionable about the content of an argument, but it

could turn out that the argument was fine all along. This is the case, I will argue, with Lewontin's so-called fallacy.

What was Lewontin's apparently fallacious argument? To answer this question, we first need to take into account his findings. Lewontin showed that approximately 85.4% of all human genetic diversity resides within any given racialized group. Furthermore, he found that another 8.3% is accounted for by variation within racialized groups. He concluded that racial classification is "of virtually no genetic or taxonomic significance" and that "no justification can be offered for its continuance" (Lewontin, 1972, p.397). The argument takes a very simple form. There is an implicit premise: races are major biological groups. There is an explicit premise: human genetic variation is minor. Then there is a conclusion: there are no human races. What's wrong with that?

According to Edwards, Lewontin's

> conclusions are based on the old statistical fallacy of analysing data on the assumption that it contains no information beyond that revealed on a locus-by-locus analysis, and then drawing conclusions solely on the results of such an analysis. The 'taxonomic significance' of genetic data in fact often arises from correlations amongst the different loci, for it is these that may contain the information which enables a stable classification to be uncovered. (Edwards, 2003, p.799)

Edwards is not disputing Lewontin's findings, which have been repeatedly verified (Barbujani et al., 1997; Rosenberg et al., 2002), but rather his conclusions. "There is nothing wrong with Lewontin's statistical analysis of variation," he explains, "only with the belief that it is relevant to classification" (Edwards, 2003, p.800). It is possible, Edwards is suggesting, to classify people racially even if we are all as genetically similar as Lewontin has shown, as long as we correlate gene frequencies.

However, Lewontin does not argue that it is impossible to classify people into racialized groups. Rather, he argues that "our perception of relatively large differences between human races and subgroups, as

compared to the variation within these groups, is indeed a biased perception and that, based on randomly chosen genetic differences, human races and populations are remarkably similar to each other, with the largest part by far of human variation being accounted for by the differences between individuals" (Lewontin, 1972, p.397). The claim is that conventional racial classification is misleading, not that it is impossible.

Whether it is possible to classify people into conventional racial categories using genetic data is not Lewontin's question. As he clearly states, "the question I am asking is, 'How much of human diversity between populations is accounted for by more or less conventional racial classification?'" (1972, p.386). As such, Edwards misinterprets Lewontin's argument. Edwards interprets the argument, it would seem, as follows. There is an implicit premise: single-locus statistics would be sufficient to categorize people into 'races.' There is an explicit premise: it is impossible to categorize people into 'races' using single-locus statistics. Then there is a conclusion: it is impossible to categorize people into 'races.'

This is not Lewontin's argument. As noted above, Lewontin actually argues the following: races are major biological groups; human genetic variation is minor; therefore, there are no human races. It is a problem, obviously, that Edwards gets the argument wrong. However, that does not mean that Edwards hasn't named a genuine fallacy; it just means that has named a potential fallacy after someone who did not commit it. Nevertheless, the claim that there is a fallacy at all needs to be shown, not assumed.

Is there a "Lewontin's fallacy," even if Lewontin did not commit it? To begin, note that it is uncontroversial that multivariate statistics, of the kind Edwards champions, is useful for classification. Peter Taylor (2011) offers the nice example of classifying males and females according to height and hip circumference. Males and females overlap on these measures, so that it is very difficult to guess who is male and who is female based on either measure alone. However, when the measures are correlated, males and females tend to cluster apart.

The new race naturalists argue that the same happens with 'race.' Correlate gene frequencies, and it is possible to group genotypes into more-or-less conventional racial categories. This does seem to suggest that Lewontin's fallacy is a genuine fallacy, even if Lewontin himself did not commit it. If it were only possible to classify genotypes into conventional racial groups when using multivariate, multi-locus analysis, then it would be fallacious to assume that single-locus analysis is sufficient to classify people racially.

As I noted above, because informal fallacies are not determined by form, but by content, there is always a risk that they will turn out not to be fallacies after all. If it were actually possible to classify people into conventional racial categories using uncorrelated genetic data, then "Lewontin's fallacy" would not be a genuine fallacy.

An important study by David Witherspoon and colleagues (2007) appears to settle the matter. They found that, when 10 polymorphic loci are analyzed, two individuals from the same geographically and historically separated population are more genetically different from each other than from individuals from a different population nearly one-third of the time. However, when individuals from geographically and historically separated populations are sampled, and many thousands of loci are analyzed, two individuals from different populations are always more genetically diverse than two individuals from the same population (Witherspoon et al., 2007).

Witherspoon and colleagues' study demonstrates that there is no such thing as Lewontin's fallacy. They show that, if enough genes are analyzed (and those genes are sampled from geographically and historically separated populations) it is possible to classify genotypes into conventional racial categories *without* correlating gene frequencies. And as they note, "The fact that, given enough [uncorrelated] genetic data, individuals can be correctly assigned to their populations of origin is compatible with the observation that most human genetic variation is found within populations, not between them" (2007, p.358).

Recall that for Edwards, "The 'taxonomic significance' of genetic data in fact often arises from correlations amongst the different loci, for it is these that may contain the information which enables a stable classification to be uncovered" (Edwards, 2003, p.799). The key word here is "often." Often multivariate statistical analysis is needed for stable classification. Often, but not always. By calling Lewontin's (1972) assumption that single-locus analysis would be sufficient to classify people into racialized groups—an assumption which Lewontin never actually makes—fallacious, Edwards hedges his bets. It turns out that he made a losing bet. Lewontin did not commit Lewontin's fallacy, and it isn't even a fallacy.

**Edwards' Fallacy**

Lewontin did not commit a fallacy, but Edwards—I argue in this section—did. Edwards assumes that racial naturalism is vindicated because it is possible to cluster genotypes into more-or-less conventional racial categories using multi-locus analysis. However, not all stable classification systems are scientifically valid. This is not a new point—I am merely giving it a name. As Koffi Maglo, Tesfaye Mersha and Lisa Martin observe, "in a rational classification of biological organisms, the computational possibility to determine group membership… does not imply that these groups are meaningful according to biological systematic and evolutionary classification criteria" (Maglo et al., 2016, p.8). I call the assumption that the ability to create a stable classification system vindicates that system as scientifically respectable 'Edwards' fallacy.'

To see why this is a fallacy, consider the following examples. It would be possible to create a stable classification system of fruit, based on the color of its skin when ripe. This system would lump strawberries with pomegranates, limes with watermelons, star fruit with quince. This classification system would be fairly stable, but it would not be of any taxonomic significance whatsoever. Similarly, amateur bird watchers classify birds based on phenotypic features that can be reliably identified, but which are often not indicative of any special taxonomic status. If the ability to design a stable classification system were sufficient to vindicate that system as scientifically meaningful, then it would

be legitimate to classify fruit based on color, and any identifiable bird variety would instantly gain a place within ornithological taxonomy—which is already quite a messy corner of science (Zink, 2004).

Stable classification is simply not sufficient for scientific taxonomy. Edwards' assumption that it is sufficient is fallacious. The argument for using multi-locus analysis for racial classification is that it is a more powerful classificatory tool than single-locus analysis. However, a more powerful classificatory tool is not necessarily helpful unless we know what we are classifying. Recall Taylor's example of classifying people based on sex by using multi-variate analysis—correlating height and hip circumference. We know that the classification is correct, but that is only because we have an independent theory of biological sex.

The same cannot be said about "race." How do the new race naturalists "know" that the clusters produced by multi-locus analysis vindicate racial taxonomy? Because they reflect back to us, roughly, *conventional* racial categories. They reflect back folk racial taxonomy. Race naturalists then exclaim that we were right all along. This cannot be the right way to determine whether there are human biological races. As Sally Haslanger observes, "we can determine whether there "really are" races only if the term 'race' has a specified meaning" (2008, p.60). We need a theory, or at least of definition of race before we can tell whether these worldwide clustering studies vindicate racial naturalism.

Edwards does not provide such a theory. While the term 'race' appears eight times in Edwards' paper, six of those instances are in quotations. The other two times are describing Lewontin's 1972 paper, and Lewontin himself uses the term 'race' freely throughout that paper until he concludes that we should eliminate "race" from our biological ontology. My point is that Edwards never discusses race in any detail; he never tells us what 'race' actually is.

Edwards appears much more comfortable with the term 'population' than with the language of 'race.' However, Lewontin's conclusion is about 'race,' rather than 'population.' Unless 'race' and 'population' are synonyms (we will return to this possibility) Edwards is changing the topic. But make

no mistake—Edwards article is about 'race.' This is crystal clear when the quotations about "race" are read in context. Yet in order to know whether racial naturalism is correct, we need a theory of race.

Even if racial naturalism were correct (I will be arguing that it is not below), Edwards' fallacy would still be a genuine fallacy, because it points to a general problem. It is an informal fallacy, because it depends on facts about the world and about how science works, but unlike Lewontin's so-called fallacy, it is unlikely to be refuted. Valid scientific taxonomies need more than mere stability. They need to be meaningful in the context of the relevant science. Unlike "Lewontin's fallacy," Edwards' fallacy is a genuine fallacy.

**Edwards' Fallacy in the Race Debate I: Forensic Anthropology**

So far, I hope to have convinced the reader that Lewontin does not commit the "fallacy" bearing his name, and that it is not a fallacy after all. Further, I hope to have shown that *Edwards* commits a fallacy. However, I have not yet argued that Edwards' fallacy is particularly harmful for racial naturalism. The ability to create a stable classification system that resembles conventional racial classification does not vindicate racial naturalism, but nor does it show that racial naturalism is false. In the following sections, I want to demonstrate that Edwards' fallacy is indicative of serious problems for racial naturalism.

Edwards fallacy is not only committed by those discussing genetics, but also in other fields where "race" is debated, such as forensic anthropology. In this section I will discuss Edwards' fallacy as it is committed in discussions surrounding "race" and forensic anthropology, before moving on to how it is committed in racial interpretations of genetic clustering studies in the section which follows. The logic in both instances is the same.

Forensic anthropology offers a useful analogy to genetic anthropology because both fields make use of multivariate statistics. Forensic anthropologists correlate, for instance, skull measurements when

identifying human remains. That is, they use the statistical toolkit that Edwards admonishes Lewontin for ignoring.

Some have argued that forensic anthropology vindicates racial naturalism. For example, Anthony Walsh and Ilhong Yun write that,

> Forensic anthropologists estimate that about 90 percent of skeletal remains can be racially identified by skull morphology alone (Sauer, 1992), which is one very useful application of racial classification. Yet surveys of physical anthropologists find that between 41 percent and 69 percent answered "no" to the question: "There are biological races within the species *Homo sapiens*" (Lieberman & Kirk, 2002). A significant majority of physical anthropologists thus function within a classificatory system they allegedly do not believe in." (Walsh & Yun, 2011, p.1287)

The evidence for racial naturalism coming out of forensic anthropology is so overwhelming, suggest Walsh and Yun, that the fact that a high percentage of anthropologists do not believe in race is treated with suspicion. They only "allegedly" don't believe in race. The surveys were anonymous, so the idea is not that individual anthropologists were afraid of being associated with a view deemed to be politically correct, or something of the sort. The idea is rather that the anthropologists may have such a "distaste for race" that they lied about whether race is real (Walsh & Yun, 2011, p.1286).

Neven Sesardic (2010) makes the same argument, also citing Norman Sauer in support. But there is a problem with this interpretation of Sauer's work. Sauer was arguing that forensic anthropology does *not* support racial naturalism, despite the fact that it is possible to classify skulls into conventional racial categories. In my reply to Sesardic, I pointed this out (Hochman, 2013b). In his response, Sesardic (2013) held his ground. Sauer was just being politically correct: he really believes in race, you just have to read between the lines. In the last instalment of this debate, I argued that Sauer was indeed being genuine (Hochman, 2016). He is not a part of the international political correctness conspiracy against racial naturalism.

In fact, Sauer anticipates the identification of Edwards' fallacy. He writes that

> Some of the confusion about this issue may stem from an assumption that to identify a specimen as having ancestors in Africa or Europe, for example, is tantamount to race identification and a verification of geographic races. No one who argues against the race concept denies that human variation exists or claims that this variation is not systematic. In fact, it is systematic variation that allows anyone to estimate, with varying degrees of specificity, a person's place of ancestry from their physical features. However, to identify a person as having ancestors from, say, Northern Europe does not identify a biological race of Northern Europeans. (Sauer, 1992, p.110)

Sauer is saying that it is consistent to accept that there is systematic biogeographic variation in our species that allows for accurate classification into conventional "racial" groups, while at the same time to reject racial naturalism.

If this is hard to see, a study by Stephen Ousley, Richard Jantz, and Donna Freid (2009) should help. It perfectly demonstrates why Edwards' Fallacy is so seductive for race naturalists, and why—when uncovered—so damaging to racial naturalism. Using multivariate methods (correlating skull measurements) Ousley and colleagues were able to classify human remains—with 84% accuracy—into the following categories: American black male, American white male, Japanese male, and Native American male. These labels match conventional racial categories. On the face of it, this seems like strong support for racial naturalism. Sesardic (2010), for instance, cites it as evidence that race is a valid biological category.

The problem with a racial interpretation of this study becomes clear when we consider the rest of the findings. Ousley and colleagues were able to classify the skulls of Chinese males, Japanese males, and Vietnamese males with 80% accuracy. These groups belong to the same so-called 'race,' but their skulls can be classified with almost as high an accuracy as those from conventional racial groups can be. The

skulls of Arikara and Sioux Females were classified with 87% accuracy, and those of Nagasaki and Tohoku Males with 94% accuracy—this is even higher accuracy than the conventional "racial" groupings.

If a race were any group that could be classified accurately (say, over 75% of the time) then the findings I describe above have shown the following: there is a black, a white, a Chinese and a Vietnamese race, and there are multiple Japanese and Native American races. One could not conclude that there were Asian and Native American "sub-races," because there is no justification for subordinating those groupings. Rather than offering support for racial naturalism, Ousley and colleagues run a reductio argument. "There are so many possible distinctive biological races," they observe, "that the concept is virtually meaningless" (2009, p.74).

If one is predisposed towards believing that there are human biological races, the fact that it is possible to classify human skulls, with a fairly high degree of accuracy, into conventional racial categories might give rise to an 'aha moment.' "I *knew* race was real," one might think to oneself. Yet accurate classification cannot be all it takes for a group to be a race. If accurate classification is all that counts, then there was a race of white males born between 1840-1890, and another race of white males born between 1930-1980, because it is possible to distinguish between these groups with 96% accuracy (Ousley et al., 2009).

Without a theory of race, race naturalists are shooting in the dark. Discussion about the reality of race tends to focus on whether scientific methods can be used to sort people into conventional racial categories. However, from a naturalistic perspective, this does not matter. Scientific racial categories could diverge from conventional ones. The fact that it is possible to group skulls into conventional racial categories tells us little of any interest about racial naturalism.

The ability to match human remains to conventional racial categories is of interest to forensic anthropologists, because they need classification systems that match those ordinary folk use to describe,

for instance, their loved ones. However, in their day-to-day work, they are not in the business of determining the truth or falsity of racial naturalism. They do not even need a theory of race to do their work: they just need to match human remains to whatever categories are used to describe people. That does not mean their work is irrelevant to the race debate. Indeed, Ousley and colleagues show that if a race is any classifiable group, racial naturalism is subject to a reductio argument.

The Ousley et al. study demonstrates why Edwards' fallacy is a serious problem for racial naturalism. Any version of racial naturalism worth defending entails that race has some special taxonomic status. What Ousley and colleagues show is that, if taxonomic status has to do with accurate classification, then there is nothing scientifically privileged about conventional racial taxonomy. Fine-grained classifications are just as accurate, in some cases more accurate, than the coarse-grained classifications associated with 'race.' Human biological variation does not adhere to conventional racial taxonomy.

Recall that Walsh and Yun claim that the ability of forensic anthropologists to classify skulls into conventional racial groups is "one very useful application of racial classification" (2011, p.1287). This is muddled logic. Forensic anthropologists are able to classify human remains according to "race" because "racial" categories are also geographical categories, and the little amount of human biological diversity that exists is distributed fairly smoothly along geographical lines. We could come up with any number of different ways of classifying people according to geographical origins that differ from conventional "racial" taxonomy, and forensic anthropologists would—if the Ousley et al. study is anything to go by—likely be able to sort remains into those categories. The fact that forensic anthropologists can sort between the skulls of people from different racialized groups shows us that forensic anthropology is useful, not—as Walsh and Yun claim—that racial classification is.

**Edwards' Fallacy in the Race Debate II: Genetic Anthropology**

If we turn from skin and bones anthropology to genetic anthropology—specifically to the worldwide genetic clustering studies that have acted as tinder to the race debate—we can see exactly the same

phenomenon taking place. Genetic anthropologists are able to classify samples from around the globe into (more-or-less) conventional racial categories (Rosenberg et al., 2002; Tang et al., 2005). While the mapping onto conventional racial categories in the Rosenberg study is imperfect (Atkin, 2017), race naturalists have read it, in particular, as offering strong evidence in favor of racial naturalism (Edwards, 2003; Sesardic, 2010; Levin, 2002; Leroi, 2005). However, as I have argued, the ability to classify something does not make it a valid scientific category, and to claim otherwise is to commit Edwards' fallacy. The new genetic clustering studies do not actually decide the debate in favor of racial naturalism, and under close inspection, they actually offer much greater support for anti-realism about race.

Edwards' fallacy is so common in discussions about genetic clusters and 'race' that there is a debate premised on the fuzzy logic that leads to it. Quayshawn Spencer calls this the *philosophical race debate*. Spencer writes that "A new racial naturalist in the philosophical race debate is a person who uses genetic clusters of human populations to argue that race, as understood by current ordinary folk in the U.S., is biologically real" (2014, p.41). The philosophical race debate contrasts with the *biological race debate*: "A new racial naturalist in the biological race debate is a person who uses genetic clusters of human populations to argue that humans can be divided into biological races according to a concept of race in use among professional biologists" (2014, p.41).

As Spencer explains, "The philosophical race debate is popular mostly among philosophers of race, but it also attracts philosophers of biology, biologists, and anthropologists, as well as other humanists and social scientists" (2014, p.41). Conversely, there are many philosophers who participate in the biological race debate, including myself. This choice of focus is not simply a matter of preference or interest. I am not a "philosophical" race debater because that debate is premised on Edwards' fallacy.

"Philosophical" race debaters are interested in whether conventional races are "biologically real," whereas biological race debaters are interested in whether race is real according to "a concept of race in use among professional biologists." 'Biological reality' is something entirely different for the two

race debaters. Unlike biological race debaters, "philosophical" race debaters are not actually interested in whether race has a genuine place biological taxonomy. They are interested in whether it is possible to classify people into conventional racial categories using scientific methodologies. However, to claim that the ability to sort genotypes into conventional racial categories makes race "biologically real" is to commit Edwards' fallacy.

The philosophical race debate should therefore be ended. There are all sorts of important race debates to be had, but this is not one of them. The philosophical race debate can be reduced to an interesting question for the biological race debater: if there are biological races, do they match folk races? Conducting a separate debate on this question, with its own logic and standards for "biological reality," makes no good sense.

The ground rules for the biological race debate make much better sense. Defend a scientific definition of 'race,' or argue that some term is the scientific synonym for 'race,' and try to show that there are (or are not) human races according to that definition. In the biological race debate, it does not matter whether the resulting racial taxonomy mirrors a conventional racial taxonomy. Epistemic authority is given to scientific methodologies, and there is every possibility that actual biological races—if there are any—would not be the same as those recognized in folk racial taxonomy.

There have been attempts to defend racial naturalism using a range of definitions of race (e.g. Andreasen, 2004; Kitcher, 1999; Pigliucci & Kaplan, 2003; Sesardic, 2010). Only one approach seems to have any real traction: the genetic population approach (Gannett, 2013; Millstein, 2015). This approach goes back at least to Theodosius Dobzhansky, and his definition of races "as populations of species that differ in the frequencies of one or more genetic variants, gene alleles or chromosomal structures" (Dobzhansky, 1937, p.138). This appears to be the definition Edwards endorses, given his use of 'population' as a synonym of 'race.'

Dobzhansky's definition has been incredibly influential. It was, for example, the definition endorsed in the UNESCO statements on race of the 1950s (UNESCO, 1952). However, it was also controversial from its inception. As Marcus Feldman and Lewontin observe, "The problem with this definition is that every geographical population of every species in the world is a "geographical race" because no two populations have identical allelic frequencies for polymorphic genes, so geographical race becomes synonymous with population" (2008, p.96). This leads to the reductio argument that Ousley, Jantz, and Freid (2009) run against race: if any identifiable population were a race, then 'race' loses its recognizable meaning; we might as well just speak of 'populations' instead.

To be fair to Dobzhansky, he did not actually believe that all populations are races. He performed some conceptual hair-splitting to avoid that implication. Dobzhansky believed that all populations were "racially distinct" (1962, p.280). "But it does not follow," he explained, "that any racially distinct populations must be given racial (or subspecific) labels. Discovery of races is a biological problem, naming races is a nomenclatorial problem. There is nothing arbitrary about whether race differences do or do not exist, but whether races should or should not be named, and if they should, how many should be recognized, is a matter of convenience and hence of judgment" (1962, p.280).

There we have it. Conventional racial classification is arbitrary, and "is a matter of convenience." Despite all of the advances in genetics, and all of the hype around worldwide genetic clustering studies, this is where we are today: with race naturalists appealing to the same deflated version of racial naturalism that Dobzhansky offered almost a century ago. It is so deflated, in fact, that it can be hard to tell the race naturalists from the race skeptics. In his article *The Race Concept: A Defense* Michael Levin observes that "No substantive disagreement separates splitters from lumpers, because any set of objects, including human beings, can be carved up in many ways, into a large number of small bits or a small number of large bits. None is more correct than, and all are consistent with, any other" (2002, p.30). When defenses of racial naturalism sound much more like attacks, you know something has gone wrong. If conventional racial classification has no scientific privilege over any other form of

classification, why keep trying to justify it? To claim that it "is a matter of convenience" just does not cut it.

Levin is no exception in holding the following views simultaneously: that conventional racial classification is arbitrary, and that racial naturalism is true. This can also be seen, for example, in the writings of Armand Leroi (2005), and of Sesardic, especially in his reply to me (Sesardic, 2013). The best way to make sense of these seemingly incompatible positions is through the lens of Edwards' fallacy. If the ability to classify genotypes into conventional racial classifications were sufficient to give race, in Edwards' words, "taxonomic significance," then it would not matter that there are other possible taxonomies of equal taxonomic significance.

The picture we see coming out of genetic anthropology mirrors that in forensic anthropology. The ability to classify genotypes and phenotypes according to conventional "race" is taken as evidence in favor of the reality of race. Yet the ability to sort genotypes, phenotypes—anything—into categories is not sufficient to make those categories scientifically respectable. To assume that it does is to commit Edwards' fallacy. Just as with forensic anthropology, genetic anthropology has shown that conventional racial classification is only one way to divide the species, and it is not privileged over other possible classification systems.

This can even be seen in the famous worldwide clustering study by Rosenberg et al. (2002), which is commonly appealed to by race naturalists. They estimated worldwide genetic structure at five different grains of analysis, dividing worldwide populations into two clusters at the lowest grain, and six at the highest. They explain (in p1. of the supplementary information section) that they only stop at six because multiple clustering solutions appeared when they attempted higher resolutions—a consequence of working with complex data sets. Indeed, it is generally accepted that "There are clearly no objective reasons for stopping at any particular level of taxonomic splitting … [and] the level at which we stop our classifications is completely arbitrary" (Cavalli-Sforza et al., 1994, p.220). As Philip Kitcher puts it, when it comes to 'race,' "there is a genuine issue about *level* or *fineness of grain*" (2007, p.305).

I've called this the 'grain-of-resolution problem' (Hochman, 2013a), and the problem is this: there seems to be no way of determining which is the "racial" grain of analysis, and thus which clusters represent "races," without appealing to convention. In relation to racial interpretations of the Rosenberg study, this problem is not confined to the worldwide analysis. Rosenberg and colleagues also perform within-continent analyses. For instance, while there is a distinct African cluster on all but the lowest grain analysis in the Rosenberg study, there are four fairly distinct African clusters in the within-continent analysis. On a racial reading of this study, is there one African "race," or are there four?

For Spencer, the philosophical race debater "can solve Hochman's 'grain-of-resolution problem' with ease because the racial level of genetic clusters of human populations is just the level that corresponds to what ordinary folk in the U.S. mean by 'race'" (2014, p.41). The population naturalist, on the other hand, solves it by saying that the racial grain of analysis is just the grain it is convenient to call the racial grain. In the end, both race debaters essentially appeal to convention, which is not at all compelling.

Far from vindicating racial naturalism, the Rosenberg study provides greater support for anti-realism about race. First, as I've shown, it highlights the seemingly intractable grain-of-resolution problem. It also demonstrates, confirming Lewontin's (1972) findings, that almost all human genetic diversity— they estimate 93-95%—resides within continental populations (a problem if "races" are meant to be major biological groups). Furthermore, in a follow-up article, they clarify that "allele frequency differences generally increase gradually with geographic distance" (Rosenberg et al., 2005, p.661). This is what anti-realists about race have been arguing for years: "There are no races, there are only clines" (Livingstone, 1962).

As Rosenberg and colleagues explain, "small discontinuities occur as geographic barriers are crossed, allowing clusters to be produced" (2005, p.661). They argue that the clusters "arise from genuine features of the underlying pattern of human genetic variation, rather than as artifacts of uneven sampling along continuous gradients of allele frequencies" (2005, p.661). This is an important point for Spencer,

who wants the clusters to be "an actual partition of human populations," rather than a product of biased sampling (2014, p.41). However, I am not convinced Rosenberg et al. have discovered an "actual partition."

For instance, Rosenberg et al. were working with genotypes from a disproportionally small number of African populations. Only six of the 52 populations they analyzed were sub-Saharan African. Yet the average nucleotide diversity within Africans almost doubles that found within Asians and Europeans (Yu et al., 2002). Practical factors (and possibly bias) means that the sampling schemes in genetic clustering studies are far from perfect. Tishkoff and colleagues used the same base sample as Rosenberg et al., but they also included new genotypes from individuals from the following populations: African, African American, Indian, Indigenous Australian, and Yemenite. Their study produce radically different clusters, showing that, indeed, "you-get-out-what-you-put-in" (Winther & Kaplan, 2013, p.66).

When Tishkoff and colleagues set the STRUCTURE program to find six clusters, four of them were located *within* Africa, even though it was a worldwide study. Where Rosenberg and colleagues stop at six clusters, Tishkoff and colleagues stop at 13. Most of the 13 clusters can also be found *within* Africa. Furthermore, most of the African genotypes have mixed-cluster membership; hardly the simple picture we find in the Rosenberg et al. study. If clusters were races, the Tishkoff study would seem to indicate that most extant human races are African.

Even if Tishkoff et al. used a disproportionally high number African genotypes (Spencer, 2014), it was not to the same extent that Rosenberg et al. used a disproportionally low number (Hochman, 2014). My point is that while it is possible, depending on the sampling scheme, to classify people into conventional racial categories—such as 'black African'—it is not an accurate way of capturing human biological diversity or the structure of human populations. In fact, it simplifies it to the point of massive distortion.

The new genetic clustering studies can only be interpreted as supporting racial naturalism through the lens of Edwards' fallacy. The findings coming out of genetic anthropology are actually devastating for racial naturalism. The amount of genetic diversity in our species is relatively quite small; that diversity is almost entirely smooth in its distribution across geographic space; any attempt to number the "races" is arbitrary; and there is greater genetic diversity within Africa than between Africa and other continents.

Racial naturalism can only be saved by deflating it to such an extent that it becomes trivial. We should reject the redefinition of race as 'a population that we choose to call a race.' It is not only a radical change in the meaning of race, so that 'race' does not refer only to a kind of major biological group, but also potentially to the inhabitants neighboring villages (Gannett, 2013). It also trivializes the concept of race. If neighboring villagers are racially distinct, and count as races if it is "convenient" to call them races, then of course race is real. Racial naturalism would be true, but at the cost of its utter triviality.

If we want to define race by looking for a synonym in biology—not a bad idea, generally speaking—'subspecies' is a much better fit than 'population.' While 'subspecies' has long been a controversial taxonomic level (Wilson & Brown, 1953), it would be the most obvious candidate for a scientific synonym for 'race,' as it is the level below species. Most current race naturalists seem to hold Dobzhansky's view that "races" are populations, but 'population' is much more like the old Linnaean category 'variety.' Linnaeus did not consider 'variety' to be a scientific category, and later category of 'race,' which was intended to be scientific, was supposed to contrast with it (Doron, 2012). So 'race' is much better understood as synonymous with 'subspecies, rather than 'population.'

Race naturalists are abandoning the idea that races are subspecies, because when the usual methodologies for detecting subspecies are used on humans, no subspecies are found (Templeton, 2013). But the fact that there are no human subspecies is not a good reason to reject the definition of 'race' as the way to talk about subspecies in humans—which, contra Spencer (2014), is quite common (Hochman, 2014). To define race as such is to give it a substantive definition. The hypothesis that there

are human subspecies/races is on-the-surface plausible, and, if true, would tell us something interesting about human biological diversity. Defining races as populations, as Edwards (2003) seems to do, tells us nothing of any real interest about human biodiversity.

**Conclusion**

Despite a few critical responses to the idea of 'Lewontin's fallacy' (Hochman, 2013a, 2013b; Maglo et al., 2016; Fujimura et al., 2014), Ian Hacking is right to remark that "A. W. F. Edwards's 2003 theoretical refutation of Lewontin, attending to correlations among traits and genetic markers, is now widely judged to be correct" (Hacking, 2005, p.85). I hope to have shown that "attending to correlations" does not actually constitute a "theoretical refutation of Lewontin." Attending to correlations can be useful for answering certain questions, but not the question that Lewontin was asking.

"Lewontin's fallacy" is meant to be the assumption that uncorrelated genetic data would be sufficient to classify genotypes into conventional racial groups. He was supposed to have committed this fallacy in his famous 1972 article, where he showed that most human genetic diversity can be found within any given racialized group. I see no evidence whatsoever that Lewontin makes the assumption that Edwards calls fallacious. Lewontin wasn't trying to classify people into racial categories; he was taking conventional racial categories and seeing how much genetic diversity resided within, as compared to between, its members.

Moreover, the assumption that uncorrelated genetic data is sufficient to classify genotypes into conventional racial categories is true (Witherspoon et al., 2007). Not only does Lewontin not commit "Lewontin's fallacy," it is not actually a fallacy.

In assuming that stable classification is all it takes to vindicate racial naturalism, Edwards commits a fallacy of his own. The ability to create a classificatory system, and then reliably sort things within that system, is not enough to make that system scientifically respectable.

Edwards' fallacy is rife in the debate about the existence of human biological races. Yes, it is possible, given certain sampling schemes, to sort skulls and genotypes into conventional racial categories. However, that does not make conventional racial taxonomy scientifically privileged. There are all kinds of ways that our species can be divided. Traditional racial taxonomies distort the true nature of human biological diversity. Human biological diversity is relatively low, and it is predominantly smooth in its distribution across geographical space. The idea of race brings to mind fairly distinct human lineages, with relatively high levels of biological differentiation, distributed in more-or-less discrete bio-packages. Human biological diversity is just not like that. Biological race is an illusion.

## References

Andreasen, R. (2004). The Cladistic Race Concept: A Defense. *Biology and Philosophy*, 19 (3), pp. 425–442.

Atkin, A. (2017). Race, definition, and science. In: N. Zack (ed.). *The Oxford Handbook of Philosophy and Race*. New York: Oxford University Press, pp. 139–149.

Barbujani, G., Magagni, A., Minch, E. and Cavalli-Sforza, L.L. (1997). An Apportionment of Human DNA Diversity. *Proceedings of the National Academy of Sciences of the United States of America*, 94 (9), pp. 4516–4519.

Cavalli-Sforza, L.L., Menozzi, P. and Piazza, A. (1994). *The history and geography of human genes*. New Jersey: Princeton University Press.

Dobzhansky, T. (1962). Comment. *Current Anthropology*, 3 (3), pp. 279–280.

Dobzhansky, T. (1937). *Genetics and the origin of species*. New York: Columbia University Press.

Doron, C.-O. (2012). Race and Genealogy: Buffon and the Formation of the Concept of "Race". *HUMANA.MENTE Journal of Philosophical Studies*, 22, pp. 75–109.

Edwards, A.W.F. (2003). Human Genetic Diversity: Lewontin's Fallacy. *Bioessays*, 25 (8), pp. 798–801.

Feldman, M.W. and Lewontin, R.C. (2008). Race, ancestry, and medicine. In: B. Koenig, S. S.-J. Lee, and S. S. Richardson (eds.). *Revisiting Race in a Genomic Age*. New Brunswick, NJ: Rutgers University Press, pp. 89–101.

Fujimura, J.H., Bolnick, D.A., Rajagopalan, R., Kaufman, J.S., Lewontin, R.C., Duster, T., Ossorio, P. and Marks, J. (2014). Clines Without Classes: How to Make Sense of Human Variation. *Sociological Theory*, 32 (3), pp. 208–227.

Gannett, L. (2013). Theodosius Dobzhansky and the Genetic Race Concept. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 44, pp. 250–261.

Hacking, I. (2005). Why Race Still Matters. *Daedalus*, 134 (1), pp. 102–116.

Haslanger, S. (2008). A social constructionist analysis of race. In: B. A. Koenig, S. Soo-Jin Lee, and S. S. Richardson (eds.). *Revisiting Race in a Genomic Age*. New Brunswick, NJ: Rutgers University Press, pp. 56–69.

Hochman, A. (2013a). Against the New Racial Naturalism. *The Journal of Philosophy*, 110 (6), pp. 331–351.

Hochman, A. (2016). Race: Deflate or Pop? *Studies in History and Philosophy of Biological and Biomedical Sciences*, 57, pp. 60–68.

Hochman, A. (2013b). Racial Discrimination: How Not to Do It. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 44 (3), pp. 278–286.

Hochman, A. (2014). Unnaturalised Racial Naturalism. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 46, pp. 79–87.

Kitcher, P. (2007). Does 'Race' Have a Future? *Philosophy & Public Affairs*, 35 (4), pp. 293–317.

Kitcher, P. (1999). Race, ethnicity, biology, culture. In: L. Harris (ed.). *Racism*. New York: Humanity Books, pp. 87–120.

Leroi, A.L. (2005). A Family Tree in Every Gene. *Journal of Genetics*, 84 (1), pp. 3–6.

Levin, M. (2002). The Race Concept: A Defense. *Behavior and Philosophy*, 30, pp. 21–42.

Lewontin, R. (1972). The apportionment of human diversity. In: T. Dobzhansky, M. K. Hecht, and W. C. Steere (eds.). *Evolutionary Biology*. New York: Meredith Publishing Company, pp. 381–398.

Lieberman, L. and Kirk, R.C. (2002). The 1999 Status of the Race Concept in Physical Anthropology: Two Studies Converge. *American Journal of Physical Anthropology, Suppl*, 34, pp. 102–102.

Livingstone, F.B. (1962). On the Non-Existence of Human Races. *Current Anthropology*, 3 (3), pp. 279–281.

Maglo, K.N., Mersha, T.B. and Martin, L.J. (2016). Population Genomics and the Statistical Values of Race: An Interdisciplinary Perspective on the Biological Classification of Human Populations and Implications for Clinical Genetic Epidemiological Research. *Frontiers in Genetics*, 7 (22), pp. 1–13.

Millstein, R.L. (2015). Thinking About Populations and Races in Time. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 52, pp. 5–11.

Ousley, S., Jantz, R. and Freid, D. (2009). Understanding Race and Human Variation: Why Forensic Anthropologists are Good at Identifying Race. *American Journal of Physical Anthropology*, 139 (1), pp. 68–76.

Pigliucci, M. and Kaplan, J.M. (2003). On the Concept of Biological Race and its Applicability to Humans. *Philosophy of Science*, 70 (5), pp. 1161–1172.

Rosenberg, N.A., Mahajan, S., Ramachandran, S., Zhao, C., Pritchard, J.K. and Feldman, M.W. (2005). Clines, Clusters, and the Effect of Study Design on the Inference of Human Population Structure. *PLoS Genetics*, 1 (6), pp. 660–671.

Rosenberg, N.A., Pritchard, J.K., Weber, J.L., Cann, H.M., Kidd, K.K., Zhivotovsky, L.A. and Feldman, M.W. (2002). Genetic Structure of Human Populations. *Science*, 298 (5602), pp. 2381–2385.

Sauer, N.J. (1992). Forensic Anthropology and the Concept of Race: If Races Don't Exist, Why Are Forensic Anthropologists So Good at Identifying Them? *Social Science & Medicine*, 34 (2), pp. 107–111.

Sesardic, N. (2013). Confusions About Race: A New Installment. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 44 (3), pp. 287–293.

Sesardic, N. (2010). Race: A Social Destruction of a Biological Concept. *Biology and Philosophy*, 25 (2), pp. 143–162.

Spencer, Q. (2014). The Unnatural Racial Naturalism. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 46, pp. 38–43.

Tang, H., Quertermous, T., Rodriguez, B., Kardia, S.L.R., Zhu, X., Brown, A., Pankow, J.S., Province, M.A., Hunt, S.C. and Boerwinkle, E. (2005). Genetic Structure, Self-Identified Race/Ethnicity, and Confounding in Case-Control Association Studies. *The American Journal of Human Genetics*, 76 (2), pp. 268–275.

Taylor, P. (2011). Rehabilitating a Biological Notion of Race? A Response to Sesardic. *Biology and Philosophy*, 26 (3), pp. 469–473.

Templeton, A. (2013). Biological Races in Humans. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 44 (3), pp. 262–271.

UNESCO (1952). *The race concept: Results of an inquiry*. Paris: United Nations Educational, Scientific and Cultural Organization.

Walsh, A. and Yun, I. (2011). Race and Criminology in the Age of Genomic Science. *Social Science Quarterly*, 92 (5), pp. 1279–1296.

Wilson, E.O. and Brown, W. (1953). The Subspecies Concept and its Taxonomic Application. *Systematic Zoology*, 2 (3), pp. 97–111.

Winther, R.G. and Kaplan, J.M. (2013). Ontologies and Politics of Biogenomic 'Race'. *Theoria*, 60 (3), pp. 54–80.

Witherspoon, D.J., Wooding, S., Rogers, A.R., Marchani, E.E., Watkins, W.S., Batzer, M.A. and Jorde, L.B. (2007). Genetic Similarities Within and Between Human Populations. *Genetics*, 176 (1), pp. 351–359.

Yu, N., Chen, F.C., Ota, S., Jorde, L.B., Pamilo, P., Patthy, L., Ramsay, M., Jenkins, T., Shyue, S.K. and Li, W.H. (2002). Larger Genetic Differences Within Africans than Between Africans and Eurasians. *Genetics*, 161 (1), pp. 269–274.

Zink, R.M. (2004). The Role of Subspecies in Obscuring Avian Biological Diversity and Misleading Conservation Policy. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 271 (1539), pp. 561–564.