

A Population Genetic Signal of Polygenic Adaptation

Jeremy J. Berg^{1,2,3*}, Graham Coop^{2,3*}

1 Graduate Group in Population Biology, University of California, Davis, Davis, California, United States of America, **2** Center for Population Biology, University of California, Davis, Davis, California, United States of America, **3** Department of Evolution and Ecology, University of California, Davis, Davis, California, United States of America



Abstract

Adaptation in response to selection on polygenic phenotypes may occur via subtle allele frequencies shifts at many loci. Current population genomic techniques are not well posed to identify such signals. In the past decade, detailed knowledge about the specific loci underlying polygenic traits has begun to emerge from genome-wide association studies (GWAS). Here we combine this knowledge from GWAS with robust population genetic modeling to identify traits that may have been influenced by local adaptation. We exploit the fact that GWAS provide an estimate of the additive effect size of many loci to estimate the mean additive genetic value for a given phenotype across many populations as simple weighted sums of allele frequencies. We use a general model of neutral genetic value drift for an arbitrary number of populations with an arbitrary relatedness structure. Based on this model, we develop methods for detecting unusually strong correlations between genetic values and specific environmental variables, as well as a generalization of Q_{ST}/F_{ST} comparisons to test for over-dispersion of genetic values among populations. Finally we lay out a framework to identify the individual populations or groups of populations that contribute to the signal of overdispersion. These tests have considerably greater power than their single locus equivalents due to the fact that they look for positive covariance between like effect alleles, and also significantly outperform methods that do not account for population structure. We apply our tests to the Human Genome Diversity Panel (HGDP) dataset using GWAS data for height, skin pigmentation, type 2 diabetes, body mass index, and two inflammatory bowel disease datasets. This analysis uncovers a number of putative signals of local adaptation, and we discuss the biological interpretation and caveats of these results.

Citation: Berg JJ, Coop G (2014) A Population Genetic Signal of Polygenic Adaptation. *PLoS Genet* 10(8): e1004412. doi:10.1371/journal.pgen.1004412

Editor: Marcus W. Feldman, Stanford University, United States of America

Received: August 3, 2013; **Accepted:** April 17, 2014; **Published:** August 7, 2014

Copyright: © 2014 Berg, Coop. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: This material is based upon work supported by the National Science Foundation under Grant No. 1262645 (GC) and 1148897 (JJB) and by the National Institute of General Medical Sciences of the National Institutes of Health under award numbers NIH RO1GM83098 and RO1GM107374 (GC). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* Email: jjberg@ucdavis.edu (JJB); gmcoop@ucdavis.edu (GC)

Introduction

Population and quantitative genetics were in large part seeded by Fisher's insight [1] that the inheritance and evolution of quantitative characters could be explained by small contributions from many independent Mendelian loci [2]. While still theoretically aligned [3], these two fields have often been divergent in empirical practice. Evolutionary quantitative geneticists have historically focused either on mapping the genetic basis of relatively simple traits [4], or in the absence of any such knowledge, on understanding the evolutionary dynamics of phenotypes in response to selection over relatively short time-scales [5]. Population geneticists, on the other hand, have usually focused on understanding the subtle signals left in genetic data by selection over longer time scales [6–8], usually at the expense of a clear relationship between these patterns of genetic diversity and evolution at the phenotypic level.

Recent advances in population genetics have also allowed for the genome-wide identification of individual recent selective events either by identifying unusually large allele frequency differences among populations and environments or by detecting the effects of these events on linked diversity [9]. Such approaches are nonetheless limited because they rely on identifying individual loci that look unusual, and thus are only capable of identifying selection on traits where an individual allele has a large and/or

sustained effect on fitness. When selection acts on a phenotype that is underwritten by a large number of loci, the response at any given locus is expected to be modest, and the signal instead manifests as a coordinated shift in allele frequency across many loci, with the phenotype increasing alleles all on average shifting in the same direction [10–14]. Because this signal is so weak at the level of the individual locus, it may be impossible to identify against the genome-wide background without a very specific annotation of which sites are the target of selection on a given trait [15,16].

The advent of well-powered genome wide association studies with large sample sizes [17] has allowed for just this sort of annotation, enabling the mapping of many small effect alleles associated with phenotypic variation down to the scale of linkage disequilibrium in the population. The development and application of these methods in human populations has identified thousands of loci associated with a wide array of traits, largely confirming the polygenic view of phenotypic variation [18].

Although the field of human medical genetics has been the largest and most rapid to pursue such approaches, evolutionary geneticists studying non-human model organisms have also carried out GWAS for a wide array of fitness-associated traits, and the development of further resources is ongoing [19–21]. In human populations, the cumulative contribution of these loci to the additive variance so far only explain a fraction of the narrow sense

Author Summary

The process of adaptation is of fundamental importance in evolutionary biology. Within the last few decades, genotyping technologies and new statistical methods have given evolutionary biologists the ability to identify individual regions of the genome that are likely to have been important in this process. When adaptation occurs in traits that are underwritten by many genes, however, the genetic signals left behind are more diffuse, and no individual region of the genome is likely to show strong signatures of selection. Identifying this signature therefore requires a detailed annotation of sites associated with a particular phenotype. Here we develop and implement a suite of statistical methods to integrate this sort of annotation from genome wide association studies with allele frequency data from many populations, providing a powerful way to identify the signal of adaptation in polygenic traits. We apply our methods to test for the impact of selection on human height, skin pigmentation, body mass index, type 2 diabetes risk, and inflammatory bowel disease risk. We find relatively strong signals for height and skin pigmentation, moderate signals for inflammatory bowel disease, and comparatively little evidence for body mass index and type 2 diabetes risk.

heritability for a given trait (usually less than 15%), a phenomenon known as the missing heritability problem [22,23]. Nonetheless, these GWAS hits represent a rich source of information about the loci underlying phenotypic variation.

Many investigators have begun to test whether the loci uncovered by these studies tend to be enriched for signals of selection, in the hopes of learning more about how adaptation has shaped phenotypic diversity and disease risk [24–27]. The tests applied are generally still predicated on the idea of identifying individual loci that look unusual, such that a positive signal of selection is only observed if some subset of the GWAS loci have experienced strong enough selection to make them individually distinguishable from the genomic background. As noted above, it is unlikely that such a signature will exist, or at least be easy to detect, if adaptation is truly polygenic, and thus many selective events will not be identified by this approach.

Here we develop and implement a general method based on simple quantitative and population genetic principals, using allele frequency data at GWAS loci to test for a signal of selection on the phenotypes they underwrite while accounting for the hierarchical structure among populations induced by shared history and genetic drift. Our work is most closely related to the recent work of Turchin et al [28], Fraser [29] and Corona et al [30], who look for co-ordinated shifts in allele frequencies of GWAS alleles for particular traits. Our approach constitutes an improvement over the methods implemented in these studies as it provides a high powered and theoretically grounded approach to investigate selection in an arbitrary number of populations with an arbitrary relatedness structure.

Using the set of GWAS effect size estimates and genome wide allele frequency data, we estimate the mean genetic value [31,32] for the trait of interest in a diverse array of human populations. These genetic values may often be poor predictors of the actual phenotypes for reasons we address below and in the Discussion. We therefore make no strong claims about their ability to predict present day observed phenotypes. We instead focus on population genetic modeling of the joint distribution of genetic values, which

provides a robust way of investigating how selection may have impacted the underlying loci.

We develop a framework to describe how genetic values covary across populations based on a flexible model of genetic drift and population history. In Figure 1 we show a schematic diagram of our approach to aid the reader. Using this null model, we implement simple test statistics based on transformations of the genetic values that remove this covariance among populations. We judge the significance of the departure from neutrality by comparing to a null distribution of test statistics constructed from well matched sets of control SNPs. Specifically, we test for local adaptation by asking whether the transformed genetic values show excessive correlations with environmental or geographic variables. We also develop and implement a less powerful but more general test, which asks whether the genetic values are over-dispersed among populations compared to our null model of drift. We show that this overdispersion test, which is closely related to Q_{ST} [33,34] and a series of approaches from the population genetics literature [35–39], gains considerable power to detect selection over single locus tests by looking for unexpected covariance among loci in the deviation they take from neutral expectations. Lastly, we develop an extension of our model that allows us to identify individual populations or groups of populations whose genetic values deviate from their neutral expectations given the values observed for related populations, and thus have likely been impacted by selection. While we develop these methods in the context of GWAS data, we also relate them to recent methodological developments in the quantitative genetics of measured phenotypes (as opposed to allele frequencies) [40,41], highlighting the useful connection between these approaches. An implementation of the methods described here in the form of a collection of R scripts is available at <https://github.com/jjberg2/PolygenicAdaptationCode>.

Results

Estimating Genetic Values with GWAS Data

Consider a trait of interest where L loci (e.g. biallelic SNPs) have been identified from a genome-wide association study. We arbitrarily label the phenotype increasing allele A_1 and the alternate allele A_2 at each locus. These loci have additive effect size estimates $\alpha_1, \dots, \alpha_L$, where α_ℓ is the average increase in an individual's phenotype from replacing an A_2 allele with an A_1 allele at locus ℓ . We have allele frequency data for M populations at our L SNPs, and denote by $p_{m\ell}$ the observed sample frequency of allele A_1 at the ℓ^{th} locus in the m^{th} population. From these, we estimate the mean genetic value in the m^{th} population as

$$Z_m = 2 \sum_{\ell=1}^L \alpha_\ell p_{m\ell} \quad (1)$$

and we take \vec{Z} to be the vector containing the mean genetic values for all M populations.

A Model of Genetic Value Drift

We are chiefly interested in developing a framework for testing the hypothesis that the joint distribution of \vec{Z} is driven by neutral processes alone, with rejection of this hypothesis implying a role for selection. We first describe a general model for the expected joint distribution of estimated genetic values (\vec{Z}) across populations under neutrality, accounting for genetic drift and shared population history.

A simple approximation to a model of genetic drift is that the current frequency of an allele in a population is normally

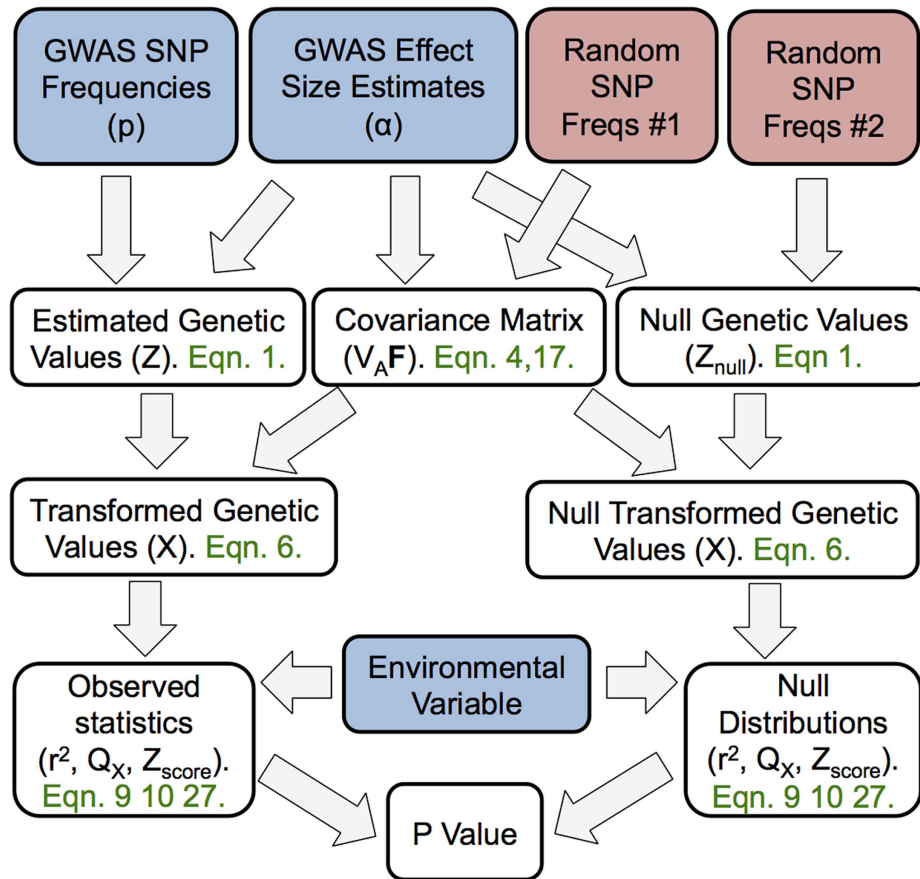


Figure 1. A schematic representation of the flow of our method. The boxes colored blue are items provided by the investigator (GWAS SNP effect sizes, the frequency of the GWAS SNPs across populations, and an environmental variable). The boxes colored red make use of random SNPs sampled to match the GWAS set as described in “Choosing null SNPs” in the methods section. For each box featuring a calculated quantity a set of equation numbers are provided for the relevant calculation. The Z score uses the untransformed genetic values, rather than the transformed genetic values, but this relationship is not depicted in the figure for the sake of readability. doi:10.1371/journal.pgen.1004412.g001

distributed around some ancestral frequency (ϵ). Under a Wright-Fisher model of genetic drift, the variance of this distribution is approximately $f\epsilon(1-\epsilon)$, where f is a property of the population shared by all loci, reflecting the compounded effect of many generations of binomially sampling [42]. Note also that for small values, f is approximately equal to the inbreeding coefficient of the present day population relative to the defined ancestral population, and thus has an interpretation as the correlation between two randomly chosen alleles relative to the ancestral population [42].

We can expand this framework to describe the joint distribution of allele frequencies across an arbitrary number of populations for an arbitrary demographic history by assuming that the vector of allele frequencies in M populations follows a multivariate normal distribution

$$\vec{p} \sim MVN(\epsilon \vec{1}, \epsilon(1-\epsilon)\mathbf{F}), \quad (2)$$

where \mathbf{F} is an M by M positive definite matrix describing the correlation structure of allele frequencies across populations relative to the mean/ancestral frequency. Note again that for small values it is also approximately the matrix of inbreeding coefficients (on the diagonal) and kinship coefficients (on the off-diagonals) describing relatedness among populations [38,43].

This flexible model was introduced, to our knowledge, by [44] (see [45] for a review), and has subsequently been used as a computationally tractable model for population history inference [42,46], and as a null model for signals of selection [38,39,47,48]. So long as the multivariate normal assumption of drift holds reasonably well, this framework can summarize arbitrary population histories, including tree-like structures with substantial gene flow between populations [46], or even those which lack any coherent tree-like component, such as isolation by distance models [49,50].

Recall that our estimated genetic values (\vec{Z}) are merely a sum of sample allele frequencies weighted by effect size. If the underlying allele frequencies are well explained by the multivariate normal model described above, then the distribution of \vec{Z} is a weighted sum of multivariate normals, such that this distribution is itself multivariate normal

$$\vec{Z} \sim MVN(\mu \vec{1}, 2V_A \mathbf{F}) \quad (3)$$

where $\mu = \frac{2}{L} \sum_{\ell=1}^L \alpha_{\ell} \epsilon_{\ell}$ and $V_A = 2 \sum_{\ell=1}^L \alpha_{\ell}^2 \epsilon_{\ell} (1-\epsilon_{\ell})$ are respectively the expected genetic value and additive genetic variance of the ancestral (global) population. The covariance matrix

describing the distribution of \vec{Z} therefore differs from that describing the distribution of frequencies at individual loci only by a scaling factor that can be interpreted as two times the contribution of the associated loci to the additive genetic variance present in a hypothetical population with allele frequencies equal to the grand mean of the sampled populations.

The assumption that the drift of allele frequencies around their shared mean is normally distributed (2) may be problematic if there is substantial drift. However, even if that is the case, the estimated genetic values may still be assumed to follow a multivariate normal distribution by appealing to the central limit theorem, as each estimated genetic value is a sum over many loci. We show in the Results that this assumption often holds in practice.

It is useful here to note that the relationship between the model for drift at the individual locus level, and at the genetic value level, gives an insight into where most of the information and statistical power for our methods will come from. Each locus adds a contribution $2\alpha_\ell(\vec{p}_\ell - \epsilon_\ell \vec{1})$ to the vector of deviations of the genetic values from the global mean. If the allele frequencies are unaffected by selection then the frequency deviation of an allele at locus ℓ in population m ($p_{m,\ell} - \epsilon_\ell$) will be uncorrelated in magnitude or sign with both the effect at locus ℓ (α_ℓ) and the allele frequency deviation taken by other unlinked loci. Thus the expected departure of the genetic value of a population from the mean is zero, and the noise around this should be well described by our multivariate normal model.

The tests described below will give positive results when these observations are violated. The effect of selection is to induce a non-independence of the allele frequency deviation ($\vec{p}_\ell - \epsilon_\ell \vec{1}$) across loci, determined by the sign and magnitude of the effect sizes [10–14] and as we demonstrate below, all of our methods rely principally on identifying this non-independence. This observation has important considerations for the false positive profile of our methods. Specifically, false positives will arise only if the GWAS ascertainment procedure induces a correlation between the estimated effect size of an allele (α_ℓ) and the deviation that this allele takes across populations ($\vec{p}_\ell - \epsilon_\ell \vec{1}$). This should not be the case if the GWAS is performed in a single population which is well mixed compared to the populations considered in the test. False positives can occur when a GWAS is performed in a structured population and fails to account for the fact that the phenotype of interest is correlated with ancestry in this population. We address this case in greater depth in the Discussion.

These observation also allows us to exclude certain sources of statistical error as a cause of false positives. For example, simple error in the estimation of α_ℓ , or failing to include all loci affecting a trait cannot cause false positives, because this error has no systematic effect on $\vec{p}_\ell - \epsilon_\ell \vec{1}$ across loci. Similarly, if the trait of interest truly is neutral, variation in the true effects of an allele across populations or over time or space (which can arise from epistatic interactions among loci, or from gene by environment interactions) will not drive false positives, again because no systematic trends in population deviations will arise. This sort of heterogeneity can, however, reduce statistical power, as well as make straightforward interpretation of positive results difficult, points which we address further below.

Fitting the Model and Standardizing the Estimated Genetic Values

As described above, we obtain the vector \vec{Z} by summing allele frequencies across loci while weighting by effect size. We do not get to observe the ancestral genetic value of the sample (μ), so we

assume that this is simply equal to the mean genetic value across populations ($\mu = \frac{1}{M} \sum_m Z_m$). This assumption costs us a degree of freedom, and so we must work with a vector \vec{Z} , which is the vector of estimated genetic values for the first $M - 1$ populations, centered at the mean of the M (see Methods for details). Note that this procedure will be the norm for the rest of this paper, and thus we will always work with vectors of length $M - 1$ that are obtained by subtracting the mean of the M vector and dropping the last component. The information about the dropped population is retained in the mean of the $M - 1$ length vectors, and thus the choice of which population to drop is arbitrary and does not affect the inference.

To estimate the null covariance structure of the $M - 1$ populations we sample a large number K random unlinked SNPs. In our procedure, the K SNPs are sampled so as to match certain properties of the L GWAS SNPs (the specific matching procedure is described in more depth below and in the Methods section). Setting ϵ_k to be the mean sample allele frequency across populations at the k^{th} SNP, we standardize the sample allele frequency in population m as $(p_{mk} - \epsilon_k) / (\epsilon_k(1 - \epsilon_k))$. We then calculate the sample covariance matrix (\mathbf{F}) of these standardized frequencies, accounting for the $M - 1$ rank of the matrix (see Methods). We estimate the scaling factor of this matrix \mathbf{F} as

$$2V_A = 4 \sum_{\ell=1}^L \alpha_\ell^2 \epsilon_\ell (1 - \epsilon_\ell). \tag{4}$$

We now have an estimated genetic value for each population, and a simple null model describing their expected covariance due to shared population history. Under this multivariate normal framework, we can transform the vector of mean centered genetic values (\vec{Z}) so as to remove this covariance. First, we note that the Cholesky decomposition of the \mathbf{F} matrix is

$$\mathbf{F} = \mathbf{C}\mathbf{C}^T \tag{5}$$

where \mathbf{C} is a lower triangular matrix, and \mathbf{C}^T is its transpose. Informally, this can be thought of as taking the square root of \mathbf{F} , and so \mathbf{C} can loosely be thought of as analogous to the standard deviation matrix.

Using this matrix \mathbf{C} we can transform our estimated genetic values as:

$$\vec{X} = \frac{1}{\sqrt{2V_A}} \mathbf{C}^{-1} \vec{Z}. \tag{6}$$

If $\vec{Z} \sim MVN(\vec{0}, 2V_A \mathbf{F})$ then $\vec{X} \sim MVN(0, \mathbb{I})$, where \mathbb{I} is the identity matrix. Therefore, under the assumptions of our model, these standardized genetic values should be independent and identically distributed $\sim N(0, 1)$ random variates [39].

It is worth spending a moment to consider what this transformation has done to the allele frequencies at the loci underlying the estimated genetic values. As our original genetic values are written as a weighted sum of allele frequencies, our transformed genetic values can be written as a weighted sum of transformed allele frequencies (which have passed through the same transform). We can write

$$\bar{X} = \frac{1}{\sqrt{2V_A}} \mathbf{C}^{-1} \bar{Z}' = \frac{2}{\sqrt{2V_A}} \sum_{\ell} \alpha_{\ell} \mathbf{C}^{-1} (\bar{p}_{\ell} - \epsilon_{\ell} \bar{\mathbf{I}}) \quad (7)$$

and so we can define the vector of transformed allele frequencies at locus ℓ to be

$$\bar{p}'_{\ell} = \mathbf{C}^{-1} (\bar{p}_{\ell} - \epsilon_{\ell} \bar{\mathbf{I}}). \quad (8)$$

This set of transformed frequencies exist within a set of transformed populations, which by definition have zero covariance with one another under the null, and are related by a star-like population tree with branches of equal length.

As such, we can proceed with simple, straightforward and familiar statistical approaches to test for the impact of spatially varying selection on the estimated genetic values. Below we describe three simple methods for identifying the signature of polygenic adaptation, which arise naturally from this observation.

Environmental Correlations

We first test if the genetic values are unusually correlated with an environmental variable across populations compared to our null model. A significant correlation is consistent with the hypothesis that the populations are locally adapted, via the phenotype, to local conditions that are correlated with the environmental variable. However, the link from correlation to causation must be supported by alternate forms of evidence, and in the lack of such evidence, a positive result from our environmental correlation tests may be consistent with many explanations.

Assume we have a vector \bar{Y} , containing measurements of a specific environmental variable of interest in each of the M populations. We mean-center this vector and put it through a transform identical to that which we applied to the estimated genetic values in (7). This gives us a vector \bar{Y}' , which is in the same frame of reference as the transformed genetic values.

There are many possible models to describe the relationship between a trait of interest and a particular environmental variable that may act as a selective agent. We first consider a simple linear model, where we model the distribution of transformed genetic values (\bar{X}) as a linear effect of the transformed environmental variables (\bar{Y}')

$$\bar{X} \sim \beta \bar{Y}' + \bar{\epsilon} \quad (9)$$

where $\bar{\epsilon}$ under our null is a set of normal, independent and identically distributed random variates (i.e. residuals), and β can simply be estimated as $\frac{Cov(\bar{X}, \bar{Y}')}{Var(\bar{Y}')}$. We can also calculate the associated squared Pearson correlation coefficient (r^2) as a measure of the fraction of variance explained by our variable of choice, as well as the non-parametric Spearman's rank correlation $\rho(\bar{X}, \bar{Y}')$, which is robust to outliers that can mislead the linear model. We note that we could equivalently pose this linear model as a mixed effects model, with a random effect covariance matrix $2V_A \mathbf{F}$. However, as we know both V_A and \mathbf{F} , we would not have to estimate any of the random effect parameters, reducing it to a fixed effect model as in (9) [51].

In the Methods (section "The Linear Model at the Individual Locus Level") we show that the linear environmental model

applied to our transformed genetic values has a natural interpretation in terms of the underlying individual loci. Therefore, exploring the environmental correlations of estimated genetic values nicely summarizes information in a sensible way at the underlying loci identified by the GWAS.

In order to assess the significance of these measures, we implement an empirical null hypothesis testing framework, using β , r^2 , and ρ as test statistics. We sample many sets of L SNPs randomly from the genome, again applying a matching procedure discussed below and in the Methods. With each set of L SNPs we construct a vector \bar{Z}_{null} , which represents a single draw from the genome-wide null distribution for a trait with the given ascertainment profile. We then perform an identical set of transformations and analyses on each \bar{Z}_{null} , thus obtaining an empirical genome-wide null distribution for all test statistics.

Excess Variance Test

As an alternative to testing the hypothesis of an effect by a specific environmental variable, one might simply test whether the estimated genetic values exhibit more variance among populations than expected due to drift. Here we develop a simple test of this hypothesis.

As \bar{X} is composed of $M - 1$ independent, identically distributed standard normal random variables, a natural choice of test statistic is given by

$$Q_X = \bar{X}^T \bar{X} = \frac{\bar{Z}'^T \mathbf{F}^{-1} \bar{Z}'}{2V_A}. \quad (10)$$

This Q_X statistic represents a standardized measure of the among population variance in estimated genetic values that is not explained by drift and shared history. It is also worth noting that by comparing the rightmost form in (10) to the multivariate normal likelihood function, we find that Q_X is proportional to the negative log likelihood of the estimated genetic values under the neutral null model, and is thus the natural measurement of the model's ability to explain their distribution. Multivariate normal theory predicts that this statistic should follow a χ^2 distribution with $M - 1$ degrees of freedom under the null hypothesis. Nonetheless, we use a similar approach to that described for the linear model, generating the empirical null distribution by resampling SNPs genome-wide. As discussed below, we find that in practice the empirical null distribution tends to be very closely matched by the theoretically predicted χ^2_{M-1} distribution.

Values of this statistic that are in the upper tail correspond to an excess of variance among populations. This excess of variance is consistent with the differential action of natural selection on the phenotype among populations (e.g. due to local adaptation). Values in the lower tail correspond a paucity of variance, and thus potentially to widespread stabilizing selection, with many populations selected for the same optimum. In this paper we report upper tail p-values from the empirical null distribution of Q_X both for our power simulations and empirical results. A two tailed test would be appropriate in cases where stabilizing selection is also of interest, however such signals are likely to be difficult to spot with GWAS data because the we are missing the large effect, low frequency alleles most likely to reveal a signal of stabilizing selection.

The relationship of Q_X to previous tests. Our Q_X statistic is closely related to Q_{ST} , the phenotypic analog of F_{ST} , which measures the fraction of the genetic variance that is among populations relative to the total genetic variance [33,34,52]. Q_{ST} is

typically estimated in traditional local adaptation studies via careful measurement of phenotypes from related individuals in multiple populations in a common garden setting. If the loci underlying the trait act in a purely additive manner and are experiencing only neutral genetic drift, then $\mathbb{E}[Q_{ST}] = \mathbb{E}[F_{ST}]$ [53,54].

If both quantities are well estimated, and we also assume that there is no hierarchical structure among the populations, then $\frac{(M-1)Q_{ST}}{F_{ST}}$ is known to have a χ^2_{M-1} distribution under a wide range of models [55–57]. This statistic is thus a natural phenotypic extension of Lewontin and Krakauer’s F_{ST} based-test (LK test) [35].

To see the close correspondence between Q_X and Q_{ST} , consider the case of a starlike population tree with branches of equal length (i.e. $f_{mm} = F_{ST}$ and $f_{m \neq n} = 0$). Under this demographic model, we have

$$\begin{aligned}
 Q_X &= \frac{(\bar{Z} - \mu)^T \mathbf{F}^{-1} (\bar{Z} - \mu)}{2V_A} = \frac{(Z_1 - \mu)^2}{2V_A F_{ST}} + \dots + \frac{(Z_{M-1} - \mu)^2}{2V_A F_{ST}} \\
 &= \frac{(M-1)\text{Var}(\bar{Z})}{2V_A F_{ST}} \tag{11} \\
 &= \frac{(M-1)\hat{Q}_{ST}}{F_{ST}}
 \end{aligned}$$

where \hat{Q}_{ST} is an estimated value for Q_{ST} obtained from our estimated genetic values. This relationship between Q_X and Q_{ST} breaks down when some pairs of populations do not have zero covariance in allele frequencies under the null, in which case the χ^2 distribution of the LK test also breaks down [36,37]. Bonhomme and colleagues [38] recently proposed an extension to the LK test that accounts for a population tree, thereby recovering the χ^2 distribution (see also [39], which relaxes the tree-like assumption), and our Q_X statistic is a natural extension of this enhanced statistic to the problem of detecting coordinated selection at multiple loci. This test is also nearly identical to that developed independently by Ovaskainen and colleagues for application to direct phenotypic measurements [40].

Writing Q_X in terms of allele frequencies. Given that our estimated genetic values are simple linear sums of allele frequencies, it is natural to ask how Q_X can be written in terms of these frequencies. Again, restricting ourselves to the case where \mathbf{F} is diagonal, (i.e. $f_{mm} = F_{ST}$ and $f_{m \neq n} = 0$), we can express Q_X as

$$Q_X = \frac{1}{V_A F_{ST}} \sum_{m=1}^{M-1} \sum_{\ell \neq m} \alpha_\ell \alpha_{\ell'} (p_{m\ell} - \epsilon_\ell)(p_{m\ell'} - \epsilon_{\ell'}), \tag{12}$$

which can be rewritten as

$$Q_X = \frac{M-1}{F_{ST}} \left(\frac{\sum_{\ell} \alpha_\ell^2 \text{Var}(\bar{p}_\ell)}{\sum_{\ell} \alpha_\ell^2 \epsilon_\ell (1 - \epsilon_\ell)} + \frac{\sum_{\ell \neq \ell'} \alpha_\ell \alpha_{\ell'} \text{Cov}(\bar{p}_\ell, \bar{p}_{\ell'})}{\sum_{\ell} \alpha_\ell^2 \epsilon_\ell (1 - \epsilon_\ell)} \right). \tag{13}$$

The numerator of the first term inside the parentheses is the weighted sum of the variance among populations over all GWAS loci, scaled by the contribution of those loci to the additive genetic variance in the total population. As such this first term is similar to F_{ST} calculated for our GWAS loci, but instead of just averaging

the among population and total variances equally across loci in the numerator and denominator, these quantities are weighted by the squared effect size at each locus. This weighting nicely captures the relative importance of different loci to the trait of interest.

The second term in (13) is less familiar; the numerator is the weighted sum of the covariance of allele frequencies between all pairs of GWAS loci, and the denominator is again the contribution of those loci to the additive genetic variance in the total population. This term is thus a measure of the correlation among loci in the deviation they take from the ancestral value, or the across population component of linkage disequilibrium. For a more in depth discussion of this relationship in the context of Q_{ST} , see [10–14].

As noted above (8), when \mathbf{F} is non-diagonal, our transformed genetic values can be written as a weighted sum of transformed allele frequencies. Consequently, we can obtain a similar expression to (13) when population structure exists, but now expressed in terms of the covariance of a set of transformed allele frequencies in populations that have no covariance with each other under the null hypothesis. Specifically, when the covariance is non-diagonal we can write:

$$\begin{aligned}
 Q_X &= \\
 (M-1) &\frac{\sum_{\ell} \alpha_\ell^2 \text{Var}(\bar{p}'_\ell)}{\sum_{\ell} \alpha_\ell^2 \epsilon_\ell (1 - \epsilon_\ell)} + (M-1) \frac{\sum_{\ell \neq \ell'} \alpha_\ell \alpha_{\ell'} \text{Cov}(\bar{p}'_\ell, \bar{p}'_{\ell'})}{\sum_{\ell} \alpha_\ell^2 \epsilon_\ell (1 - \epsilon_\ell)}. \tag{14}
 \end{aligned}$$

We refer to the first term in this decomposition as the standardized F_{ST} -like component and the second term as the standardized LD-like component. Under the neutral null hypothesis, the expectation of the second term is equal to zero, as drifting loci are equally likely to covary in either direction. With differential selection among populations, however, we expect loci underlying a trait not only to vary more than we would expect under a neutral model, but also to covary in a consistent way across populations. Models of local adaptation predict that it is this covariance among alleles that is primarily responsible for differentiation at the phenotypic level [10–14], and we therefore expect the Q_X statistic to offer considerably increased power as compared to measuring average F_{ST} or identifying F_{ST} outliers. We use simulations to demonstrate this fact below, and also demonstrate the perhaps surprising result that for a broad parameter range the standardized LD-like component exhibits almost no loss of power when used as a test statistic.

Identifying Outlier Populations

Having detected a putative signal of selection for a given trait, one may wish to identify individual regions and populations which contribute to the signal. Here we rely on our multivariate normal model of relatedness among populations, along with well understood methods for generating conditional multivariate normal distributions, in order to investigate specific hypotheses about individual populations or groups of populations. Using standard results from multivariate normal theory, we can generate the expected joint conditional distribution of genetic values for an arbitrary set of populations given the observed genetic values in some other set of populations. These conditional distributions allow for a convenient way to ask whether the estimated genetic values observed in certain populations or groups of populations differ significantly from the values we would expect them to take under the neutral model given the values observed in related populations.

Specifically, we exclude a population or set of populations, and then calculate the expected mean and variance of genetic values in

these excluded populations given the values observed in the remaining populations, and the covariance matrix relating them. Using this conditional mean and variance, we calculate a Z-score to describe how well fit the estimated genetic values of the excluded populations are by our model of drift, conditional on the values in the remaining populations. In simple terms, the observation of an extreme Z-score for a particular population or group of populations may be seen as evidence that that group has experienced directional selection on the trait of interest (or a correlated one) that was not experienced by the related populations on which we condition the analyses. The approach cannot uniquely determine the target of selection, however. For example, conditioning on populations that have themselves been influenced by directional selection may lead to large Z-scores for the population being tested, even if that population has been evolving neutrally. We refer the reader to the Methods section for a mathematical explication of these approaches.

Datasets

We conducted power simulations and an empirical application of our methods based on the Human Genome Diversity Panel (HGDP) population genomic dataset [58], and a number of GWAS SNP sets. To ensure that we made the fullest possible use of the information in the HGDP data, we took advantage of a genome wide allele frequency dataset of ~ 3 million SNPs imputed from the Phase II HapMap into the 52 populations of the HGDP. These SNPs were imputed as part of the HGDP phasing procedure in [59]; see our Methods section for a recap of the details. We applied our method to test for signals of selection in six human GWAS datasets identifying SNPs associated with height, skin pigmentation, body mass index (BMI), type 2 diabetes (T2D), Crohn's Disease (CD) and Ulcerative Colitis (UC).

Choosing null SNPs. Various components of our procedure involve sampling random sets of SNPs from across the genome. While we control for biases in our test statistics introduced by population structure through our \mathbf{F} matrix, we are also concerned that subtle ascertainment effects of the GWAS process could lead to biased test statistics, even under neutral conditions. We control for this possibility by sampling null SNPs so as to match the joint distribution of certain properties of the ascertained GWAS SNPs. Specifically, we chose our random SNPs to match the GWAS SNPs in each study in terms of their minor allele frequency (MAF) in the ascertainment population and the imputation status of the allele in our population genomic dataset (i.e. whether the allele was imputed or present in the original HGDP genotyping panel). In addition, we were concerned that GWAS SNPs might be preferentially found close to genes and in low recombination regions, the latter due to better tagging, and as such may be subject to a high rate of drift due to background selection, leading to higher levels of differentiation at these sites [60]. Therefore, in addition to MAF and imputation status, we also matched our random SNPs to an estimate of the background selection environment experienced by the GWAS SNPs, as measured by B value [61], which is a function of both the density of functional sites and recombination rate calibrated to match the reduction in genetic diversity due to background selection. We detail the specifics of the binning scheme for matching the discretized distributions of GWAS and random SNPs in the Methods.

Power Simulations

To assess the power of our methods in comparison to other possible approaches, we conducted a series of power simulations. There are two possible approaches to simulate the effect of selection on large scale allele frequency data of the type for which

our methods are designed. The first is to simulate under some approximate model of the evolutionary history (e.g. full forward simulation under the Wright-Fisher model with selection). The second is to perturb real data in such a way that approximates the effect of selection. We choose to pursue the latter, both because it is more computationally tractable, and because it allows us to compare the power of our different approaches for populations with evolutionary histories of the same complexity as the real data we analyze. Each of our simulations will thus consist of sampling 1000 sets of SNPs matched to the height dataset (in much the same way we sample SNPs to construct the null distributions of our test statistics), and then adding slight shifts in frequency in various ways to mimic the effect of selection.

Below we first describe the set of alternative statistics to which we compare our methods. We then describe the manner in which we add perturbations to mimic selection, and lastly describe a number of variations on this theme which we pursued in order to better demonstrate how the power of our statistics changes as we vary parameters of the trait of interest, evolutionary process, or the ascertainment.

Statistics tested. For our first set of simulation experiments we compared two of our statistics, (r^2 and Q_X) against their naive counterparts, which are not adjusted for population structure (*naive* r^2 and Q_{ST}). We also include the adjusted F_{ST} -like and LD-like components of Q_X as their behavior over certain parameter ranges is particularly illuminating. For Q_{ST} , Q_X , and its components, we count a given simulation as producing a positive result if the statistic lies in the upper 5% tail of the null distribution, whereas for the environmental correlation statistics (r^2 and *naive* r^2) we use a two-tailed 5% test. We also compared our tests to a single locus enrichment test, where we tested for an enrichment in the number of SNPs that individually show a correlation with the environmental variable. We considered this test to produce a positive result if the number of individual loci in the 5% tail of the null distribution for individual locus r^2 was itself in the 5% tail using a binomial test. We do not include our alternative linear model statistics β and ρ in these plots for the sake of figure legibility, but they generally had very similar power to that of r^2 . While slightly more powerful versions of the r^2 enrichment test that better account for sampling noise are available [47], note that our tests could be extended similarly as well, so the comparison is fair.

Simulating selection. We base our initial power simulations on empirical data altered to have an increasing effect of directional selection along a latitudinal gradient. In order to mimic the effect of selection, we generate a new set of allele frequencies ($p_{s,m\ell}$) by taking the original frequency ($p_{m\ell}$) and adding a small shift according to

$$p_{s,m\ell} = p_{m\ell} + p_{m\ell}(1 - p_{m\ell})\alpha_{\ell}\delta Y_m \quad (15)$$

where α_{ℓ} is the effect size assigned to locus ℓ , and Y_m is the mean centered absolute latitude of the population. We use 1000 simulations at $\delta=0$ to form null distribution for each of our test statistics, and from this established the 5% significance level. We then increment δ and give the power of each statistic as the fraction of simulations whose test statistic falls beyond this cutoff. While this approach to simulating selection is obviously naive to the way selection actually operates, it captures many of the important effects on the loci underlying a given trait. Namely, loci will have greater shifts if they experience extreme environments, have large effects on the phenotype, or are at intermediate frequencies. Because we add these shifts to allele frequencies sampled from real, putatively neutral loci, the effect of drift on

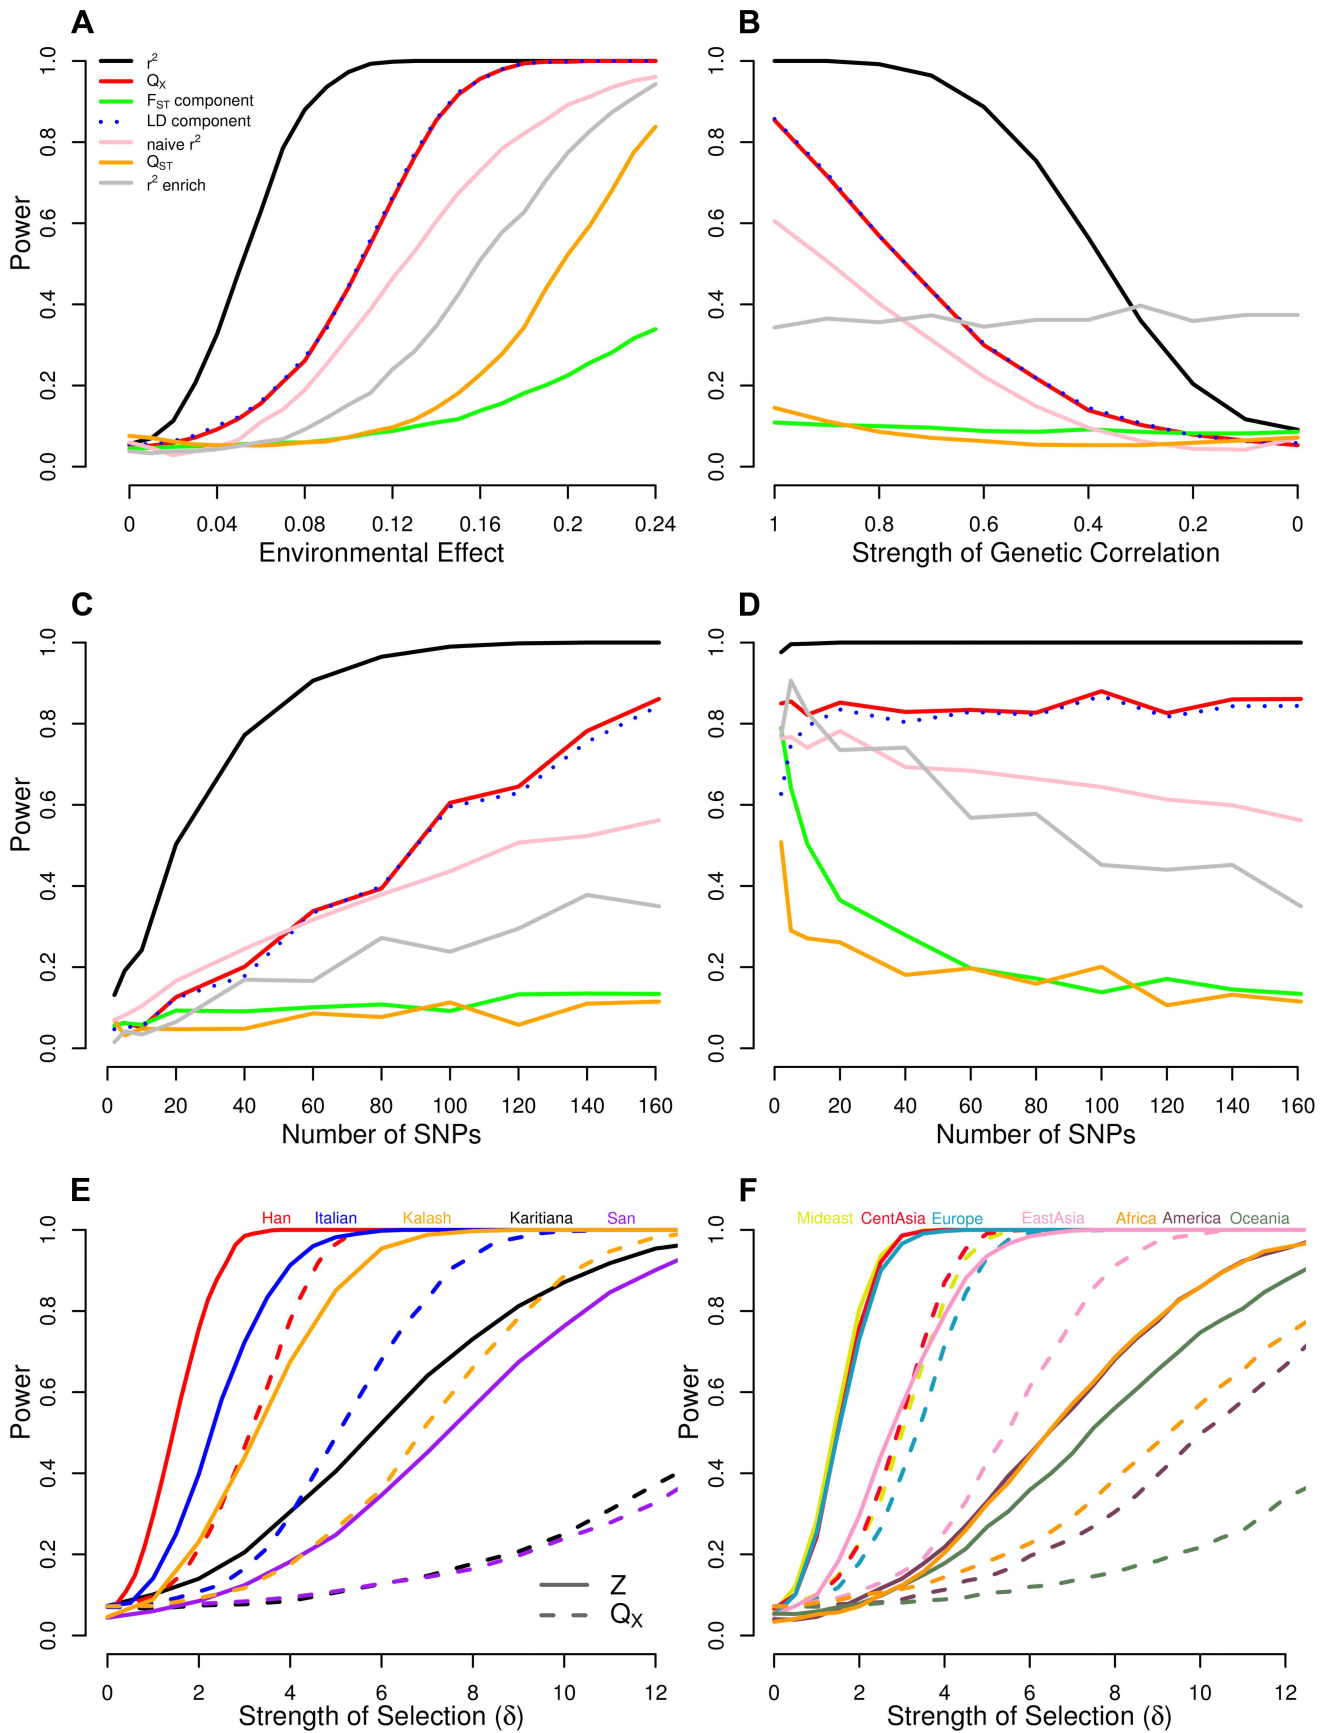


Figure 2. Power of our statistics as compared to alternative approaches. (A) across a range of selection gradients (δ) of latitude, and when we hold δ constant at 0.14 and (B) decrease ϕ , the genetic correlation between the trait of interest and the selected trait, (C) vary the number of loci, and (D) vary the number of loci while holding the fraction of variance explained constant. Bottom panels show power of the Z-test and Q_X approaches to detect selection affecting (E) a single population, and (F) multiple populations in a given region. See main text for simulation details.

doi:10.1371/journal.pgen.1004412.g002

their joint distribution is already present, and thus does not need to be simulated. The results of these simulations are shown in Figure 2A.

Our population structure adjusted statistics clearly outperform tests that do not account for structure, as well as the single locus outlier based test. Particularly noteworthy is the fact that the power of a test relying on Q_X and that using only the LD-like component are essentially identical over the entire range of simulation, while the F_{ST} -like component achieves only about 20% power by the point at which the former statistics have reached 100%. This reinforces the observation from previous studies of Q_{ST} that for polygenic traits, nearly all of the differentiation at the trait level arises as a consequence of across population covariance among the underlying loci, and not as a result of substantial differentiation at the loci themselves [14]. While our environment-genetic value correlation tests considerably outperform Q_X , this is somewhat artificial as it assumes that we know the environmental variable responsible for our allele frequency shift. In reality, the power of the environmental variable test will depend on the investigator's ability to accurately identify the causal variable (or one closely correlated with it) in the particular system under study, and thus in some cases Q_X may have higher power in practice. Panels A and B from Figure 2 with SNPs matched to each of the other traits we investigate can be found in Figures S1–S5.

Pleiotropy and correlated selection. We next considered the fact that many of the loci uncovered by GWAS are may be relatively pleiotropic, and thus may simultaneously respond to selection on multiple different traits. To explore how our methods perform in the presence of undetected pleiotropy, we consider the realization that from the perspective of allele frequency change there is only one effect that matters, and that is the effect on fitness. We therefore chose a simple and general approach to capture a flavor of this situation. We simulate the effect of selection as above (15), but give each locus an effect on fitness (α'_ℓ) that may be only partially correlated with the observed effect sizes for the trait of interest (with the unaccounted for effect on fitness coming via pleiotropic relationships to any number of unaccounted for phenotypes). For simplicity we assume that α_ℓ and α'_ℓ have a bivariate normal distribution around zero with equal variance and correlation parameter ϕ . We then simulate α'_ℓ from its conditional distribution given α_ℓ (i.e. $\alpha'_\ell|\alpha_\ell \sim N(\phi\alpha_\ell, (1-\phi^2)Var(\alpha))$). For each SNP ℓ in (15) we replaced α_ℓ by its effect α'_ℓ on the unobserved phenotype, but then perform our tests using the α_ℓ measured for the trait of interest. Here ϕ can be thought of as the genetic correlation between our phenotype and fitness if this simple multivariate form held true for all of the loci contributing to the trait. The extremes of $\phi=1$ and $\phi=0$ respectively represent the cases where selection acts only on the focal trait and that were all the underlying loci are affected by selection, but not due to their relationship with the focal trait. These simulations can also informally be seen as modeling the case where the GWAS estimated effect sizes are imperfectly correlated with the true effect sizes that selection sees, for example due to measurement error in the GWAS.

In Figure 2B we hold the value of δ constant at 0.14 and vary the genetic correlation ϕ from one down to zero. Predictably, our GWAS genetic value based statistics lose power as the the focal trait becomes less correlated with fitness but do retain reasonable power out to quite low genetic correlations (e.g. our r^2 outperforms the single locus metrics until $\phi < 0.3$). In contrast, counting the number of SNPs that are significantly correlated with a given environmental variable remains equally powerful across all genetic correlations. This is because the single locus environmental correlation tests treat each locus separately with no regards to whether there is agreement across alleles with the same direction of effect size. This may be a desirable property of the environmental outliers enrichment approach, as it does not rely on a close relationship between the effect sizes and the way that selection acts on the loci. On the other hand, this is also problematic, as such tests may often be detecting selection on only very weakly pleiotropically related phenotypes. Our approaches, however, are more suited to determining whether the genetic basis of a trait of interest, or one that is genetically correlated it, has been affected by differentiating selection.

Ascertainment and genetic architecture. We next investigated the relationship between statistical power, the number of loci associated with the trait, and the amount of variance explained by those loci. Our simulations were motivated by the fact that the number of loci identified by a given GWAS, and the fraction of variance explained by those loci, will depend on both the design of the study (e.g. sample size) and the genetic architecture of the trait. To illustrate the impact these factors have on the power of our methods, we performed two experiments in which we again held δ constant at 0.14. In the first, for each of the 1000 sets of 161 loci chosen above to mimic the height data ascertainment, we randomly sampled n loci, without regard to effect sizes, and recalculated the null distribution and power for these reduced sets, allowing n to range from 2 to 161. The results of these simulations are shown in Figure 2C. This corresponds to imagining that fewer loci had been ascertained by the initial GWAS, and estimating the power our methods would have with this reduced set of loci. As we down sample our loci without regard to effect sizes, the horizontal axis of Figure 2C is proportional to the phenotypic variance explained, e.g. the simulations in which only 80 loci are subsampled correspond to having a dataset which explains only 50% of the variance explained in those for which all 161 were used.

The second experiment is nearly identical to the first, except that before adding an effect of selection to the subsampled loci, we linearly rescale the effect sizes such that V_A is held constant at the value calculated for the full set of 161 loci. The results of these simulations are shown in Figure 2D. These simulations correspond to imagining that we have explained an equivalent amount of phenotypic variance, but the number of loci over which this variation is partitioned varies.

Our results (Figure 2C and 2D) demonstrate that even if only a small number of loci associated with the phenotype have been identified, our tests offer higher power than single locus-based tests. Moreover, for statistics that appropriately deal with both covariance among loci and among populations (r^2 and Q_X), power is generally a constant function of variance explained by the

underlying loci, regardless of the number of loci over which it is partitioned. Notably, most the power of Q_X comes from the LD-like component, especially when the number of loci is large. Statistics that rely on an average of single locus metrics (the F_{ST} -like component of Q_X), and those that rely on outliers (r^2 enrichment) all lose power as the variance explained is partitioned over more loci, as the effect of selection at each locus is weaker. Somewhat surprisingly, the versions of our tests that fail to adequately control for population structure (*naïve* r^2 and Q_{ST}) also lose power as the phenotypic variance is spread among more loci. We believe this reflects the fact that they are being systematically misled by LD among SNPs due to population structure, a problem which is compounded as more loci are included in the test. Overall these results suggest that accounting for population structure and using the LD between like effect alleles is key to detecting selection on polygenic phenotypes.

Localizing signatures of selection. Lastly, we investigated the power of our conditional Z-scores to identify signals of selection that are specific to particular populations or geographic regions, and contrast this with the power of the global Q_X statistic to detect the same signal. We again perform two experiments. In the first, we choose a single population whose allele frequencies to perturb, and leave all other populations unchanged. In other words, an effect of selection is mimicked according to (15), but with Y_m set equal to one for a single population, and zero for all others. We then increment δ to see how power changes as the effect of selection becomes more pronounced. In Figure 2E we display the results of these simulations for five populations chosen to capture the range of power profiles for the populations we consider in our empirical applications. In the last experiment, we chose a group of populations to which to apply the allele frequency shift, again consistent with (15), but now with Y_m set equal to 1 for all populations in an entire region, and zero elsewhere. In Figure 2F, we show the results of these simulations, with each of the seven geographic/genetic clusters identified by Rosenberg et al (2002) [62], chosen in turn as the affected region.

These simulations demonstrate that the conditional Z test can detect subtler frequency shifts than the global Q_X test, provided one knows which population(s) to test *a priori*. They also show how unusual frequency patterns indicative of selection are easier to detect in populations for which the dataset contains closely related populations that are unaffected (e.g. compare the Han and Italian to the San and Karitiana at the individual population level, or Europe, the Middle East and Central Asia to Africa, America, and Oceania at the regional level). Lastly, note that the horizontal axes in Figure 2E and 2F are equivalent in the sense that for a given value of δ , alleles in (say) the Italian population have been shifted by the same amount in the Italian specific simulations in Figure 2E as in the Europe-wide simulation in Figure 2F, indicating that the HGDP dataset, power is similar in efforts to detect local, population specific events, as well as broader scale, regional level events.

Empirical Applications

We estimated genetic values for each of six traits from the subset of GWAS SNPs that were present in the HGDP dataset, as described above. We discuss the analysis of each dataset in detail below, and address general points first. For each dataset, we constructed the covariance matrix from a sample of approximately 20,000 appropriately matched SNPs, and the null distributions of our test statistics from a sample of 10,000 sets of null genetic values, which were also constructed according to a similar matching procedure (as described in the Methods).

In an effort to be descriptive and unbiased in our decisions about which environmental variables to test, we tested each trait for an effect of the major climate variables considered by Hancock et al (2008) [63] in their analysis of adaptation to climate at the level of individual SNPs. We followed their general procedure by running principal components (PC) analysis for both seasons on a matrix containing six major climate variables, as well as latitude and longitude (following Hancock et al's rationale that these two geographic variables may capture certain elements of the long term climatic environment experienced by human populations). The percent of the variance explained by these PCs and their weighting (eigenvectors) of the different environmental variables are given in Table 1. We view these analyses largely as a descriptive data exploration enterprise across a relatively small number of phenotypes and distinct environmental variables, and do not impose a multiple testing penalty against our significance measures. A multiple testing penalization or false discovery rate approach may be needed when testing a large number phenotypes and/or environmental variables.

We also applied our Q_X test to identify traits whose underlying loci showed consistent patterns of unusual differentiation across populations, with results reported in Table 2. In Figure 3 we show for each GWAS set the observed value of Q_X and its empirical null distribution calculated using SNPs matched to the GWAS loci as described above. We also plot the expected null distribution of the Q_X statistic ($\sim \chi^2_{51}$). The expected null distribution closely matches the empirical distribution in all cases, suggesting that our multivariate normal framework provides a good null model for the data (although we will use the empirical null distribution to obtain measures of statistical significance).

For each GWAS SNP set we also separate our Q_X statistic into its F_{ST} -like and LD-like terms, as described in (14). In Figure 4 we plot the null distributions of these two components for the height dataset as histograms, with the observed value marked by red arrows (Figures S6–S10 give these plots for the other five traits we examined). In accordance with the expectation from our power simulations, the signal of selection on height is driven entirely by covariance among loci in their deviations from neutrality, and not by the deviations themselves being unusually large.

Lastly, we pursue a number or regionally restricted analyses. For each trait and for each of the seven geographic/genetic clusters described by Rosenberg et al (2002) [62], we compute a region specific Q_X statistic to get a sense for the extent to which global signals we detect can be explained by variation among populations with these regions, and to highlight particular populations and traits which may merit further examination as more association data becomes available. The results are reported in Table 3. We also apply our conditional Z-score approach at two levels of population structure: first at the level of Rosenberg's geographic/genetic clusters, testing each cluster in turn for how differentiated it is from the rest of the world, and second at the level of individual populations. The regional level Z-scores are useful for identifying signals of selection acting over broad regional scale or on deeper evolutionary timescales, while the population specific Z-scores are useful for identifying very recent selection that has only impacted a single population. We generally employ these regional statistics as a heuristic tool to localize signatures of selection uncovered in global analyses, or in cases where there is no globally interesting signal, to highlight populations or regions which may merit further examination as more association data becomes available. The result of these analyses are depicted in Figure 5, as well as Tables S3–S14.

Height. We first analyzed the set of 180 height associated loci identified by Lango Allen and colleagues [64], which explain

Table 1. The contribution of each geo-climatic variable to each of our four principal components, scaled such that the absolute value of the entries in each column sum to one (up to rounding error).

Geo-Climatic Variable	SUMPC1	SUMPC2	WINPC1	WINPC2
Latitude	-0.16	-0.10	-0.17	-0.01
Longitude	0.02	0.12	0.04	0.05
Maximum Temp	0.24	-0.08	0.17	-0.03
Minimum Temp	0.24	0.07	0.17	0.08
Mean Temp	0.25	-0.03	0.17	0.03
Precipitation Rate	-0.01	0.16	0.07	0.32
Relative Humidity	-0.06	0.21	-0.06	0.34
Short Wave Radiation Flux	-0.03	-0.22	0.15	-0.13
Percent of Variance Explained	38%	35%	58%	20%

We also show for each principal component the percent of the total variance across all eight variables that is explained by the PC.
doi:10.1371/journal.pgen.1004412.t001

about 10.5% of the total variance for height in the mapping population, or about 15% of heritability [65]. This dataset is an ideal first test for our methods because it contains the largest number of associations identified for a single phenotype to date, maximizing our power gain over single locus methods (Figure 2). In addition, Turchin and colleagues [28] have already identified a signal of pervasive weak selection at these same loci in European populations, and thus we should expect our methods to replicate this observation.

Of the 180 loci identified by Lango Allen and colleagues, 161 were present in our HGDP dataset. We used these 161 loci in conjunction with the allele frequency data from the HGDP dataset to estimate genetic values for height in the 52 HGDP populations. Although the genetic values are correlated with the observed heights in these populations, they are unsurprisingly imperfect predictions (see Figure S11 and Table S1, which compares our estimated genetic values to observed sex-average heights for the subset of HGDP populations with a close proxy in the dataset of Gustafsson and Lindfors (2009) [66]).

We find a signal of excessive correlation with winter PC2 (Figure 6 and Table 2), but find no strong correlations with any other climatic variables. Our Q_X test also strongly rejects the neutral hypothesis, suggesting that our estimated genetic values are overly dispersed compared to the null model of neutral genetic drift and shared population history (Figure 3 and Table 2). These results are consistent with directional selection acting in

concert on alleles influencing height to drive differentiation among populations at the level of the phenotype.

We followed up on these results by conducting regional level analyses, which indicate that our signal of excess variance arises primarily from extreme differentiation among populations within Europe (Table 3). Analyses using the conditional multivariate normal model indicate that this signal is driven largely by divergence between the French and Sardinian populations, in line with Turchin et al's (2012) previous observation of a North-South gradient of height associated loci in Europe. We also find weaker signals of over-dispersion in other regions, but the globally significant Q_X statistic can be erased by removing either the French or the Sardinian population from the analysis, suggesting that the signal is primarily driven by differentiation among those two populations.

Skin pigmentation. We next analyzed data from a recent GWAS for skin pigmentation in an African-European admixed population of Cape Verdeans [67], which identified four loci of major effect that explain approximately 35% of the variance in skin pigmentation in that population after controlling for admixture proportion. Beleza et al (2013) report effect sizes in units of modified melanin (MM) index, which is calculated as $100 \times \log(1/\% \text{melanin reflectance at } 650 \text{ nm})$, i.e. a higher MM index corresponds to darker skin, and a lower value to lighter skin.

We used these four loci to calculate a genetic skin pigmentation score in each of the HGDP populations. As expected, we identified

Table 2. Climate Correlations and Q_X statistics for all six phenotypes in the global analysis.

Phenotype	SUMPC1	SUMPC2	WINPC1	WINPC2	Latitude	Q_X
Height	-0.03 (0.21)	10^{-5} (0.99)	-0.008 (0.52)	0.086 (0.035)	0.009 (0.50)	86.9 (0.002)
Skin Pigmentation	0.061 (0.073)	0.003 (0.69)	0.048 (0.13)	-0.008 (0.51)	-0.085 (0.038)	79.1 (0.006)
Body Mass Index	-0.034 (0.19)	0.001 (0.82)	-0.022 (0.31)	0.044 (0.14)	0.031 (0.22)	67.2 (0.087)
Type 2 Diabetes	0.014 (0.40)	0.012 (0.45)	0.025 (0.27)	-0.006 (0.573)	-0.05 (0.11)	39.3 (0.902)
Crohn's Disease	0.07 (0.062)	-0.099 (0.022)	0.0001 (0.94)	-0.09 (0.039)	0.01 (0.55)	47.1 (0.68)
Ulcerative Colitis	0.03 (0.21)	-0.087 (0.034)	0.004 (0.67)	-0.049 (0.12)	0.01 (0.43)	48.58 (0.61)

We report $\text{sign}(\beta)r^2$, for the correlation statistics, such that they have an interpretation as the fraction of variance explained by the environmental variable, after removing that which is explained by the relatedness structure, with sign indicating the direction of the correlation. P-values are two-tailed for r^2 and upper tail for Q_X . Values for β and ρ are reported in Tables S15 and S16.
doi:10.1371/journal.pgen.1004412.t002

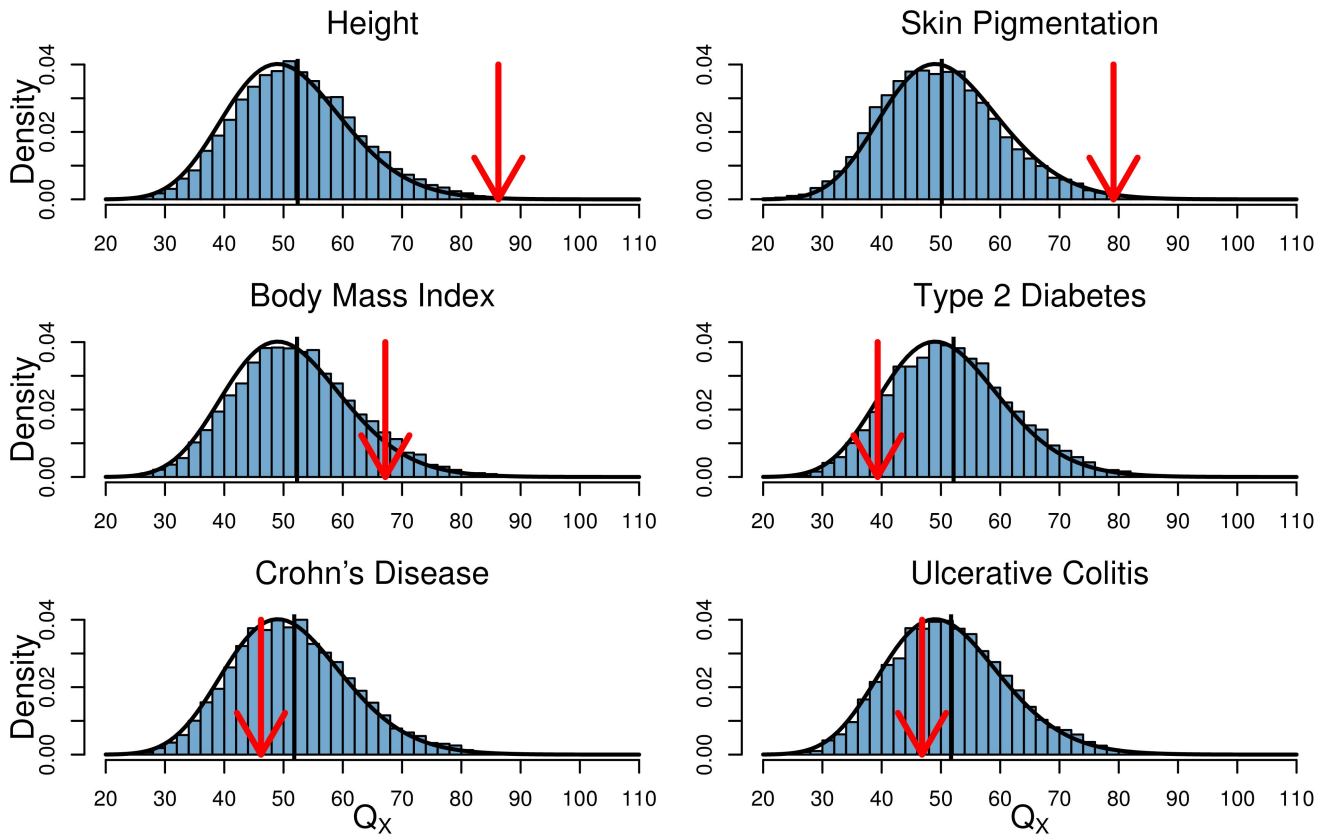


Figure 3. Histogram of the empirical null distribution of Q_X for each trait, obtained from genome-wide resampling of well matched SNPs. The mean of each distribution is marked with a vertical black bar and the observed value is marked by a red arrow. The expected χ^2_{M-1} density is shown as a black curve.
doi:10.1371/journal.pgen.1004412.g003




Figure 4. The two components of Q_X for the height dataset, as described by the left and right terms in (14). The null distribution of each statistic is shown as a histogram. The mean value is shown as a black bar, and the observed value as a red arrow.
doi:10.1371/journal.pgen.1004412.g004

Table 3. Q_X statistics and their empirical p-values for each of our six traits in each of the seven geographic regions delimited by [62].

	Europe (7)	Middle East (3)	Central Asia (8)	East Asia (16)	Americas (4)	Oceania (1)	Africa (6)
Height	32.6 ($<10^{-4}$)	7.3 (0.07)	15.5 (0.05)	18.2 (0.33)	4.2 (0.43)	0.007 (0.94)	5.4 (0.53)
Skin Pigmentation	9.7 (0.22)	9.6 (0.026)	23.4 (0.002)	13.8 (0.62)	1.3 (0.89)	0.38 (0.57)	16.2 (0.011)
Body Mass Index	9.1 (0.24)	1.6 (0.66)	9.3 (0.32)	28.4 (0.03)	13.1 (0.016)	1.2 (0.31)	1.9 (0.94)
Type 2 Diabetes	2.0 (0.96)	0.90 (0.83)	8.1 (0.43)	7.5 (0.96)	8.0 (0.13)	2.5 (0.15)	2.5 (0.88)
Crohn's Disease	6.6 (0.47)	0.87 (0.84)	7.56 (0.48)	15.5 (0.52)	1.3 (0.88)	2.5 (0.13)	2.6 (0.82)
Ulcerative Colitis	8.4 (0.30)	2.6 (0.48)	10.9 (0.21)	9.2 (0.907)	0.43 (0.986)	2.6 (0.12)	3.5 (0.77)

The theoretical expected value of the statistic under neutrality for each region is equal to $M - 1$, where M is the number of populations in the region. We report the value of $M - 1$ next to each region for reference.
doi:10.1371/journal.pgen.1004412.t003

a strong signal of excess variance among populations, as well as a strong correlation with latitude (Figure 6 and Table 2), again consistent with directional selection having acted on the phenotype of skin pigmentation to drive divergence among populations. Note, however, that this signal was driven entirely by the fact that populations of western Eurasian descent have a lower genetic skin pigmentation score than populations of African descent. Using only the markers from [67], light skinned populations in East Asian and the Americas have a genetic skin pigmentation score that is almost as high (dark) as that of most African populations, an effect that is clearly visible when we plot the measured skin pigmentation and skin reflectance of HGDP populations [68,69] against their genetic values (see Figures S12 and S13). The correlation with latitude is thus weaker than one might expect, given the known phenotypic distribution of skin pigmentation in human populations [68,70]. To illustrate this point further, we reran the analysis on a subsample of the HGDP consisting of populations from Europe, the Middle East, Central Asia, and Africa. In this subsample, the correlation with latitude, and signal of excess variance, was notably stronger ($r^2=0.2$, $p=0.019$; $Q_X=60.1$, $p=8 \times 10^{-4}$).

This poor fit to observed skin pigmentation is due to the fact that we have failed to capture all of the loci that contribute to variation in skin pigmentation across the range of populations sampled, likely due to the partial convergent evolution of light skin pigmentation in Western and Eastern Eurasian populations [71]. Including other loci putatively involved in skin pigmentation (OCA2 and KITLG) [72,73] decreases the estimated genetic pigmentation score of the other Eurasian populations (Figures S12 and S13 and Table S2), but we do not include these in our main analyses as they differ in ascertainment (and the role of KITLG in human pigmentation variation has been contested by [67]).

Within Africa, the San population has a decidedly lower genetic skin pigmentation score than any other HGDP African population. This is potentially in accordance with the observation that the San are more lightly pigmented than other African populations represented by the HGDP [68] and the observation that other putative light skin pigmentation alleles have higher frequency in the San than other African populations [71]. Although there is still much work to be done on the genetic basis of skin pigment variation within Africa, in this dataset a regional analysis of the six African populations alone identifies a marginally significant correlation with latitude ($r^2=0.62$, $p=0.0612$), and a signal of excess variance among populations ($Q_X=16.19$, $p=0.01$), suggesting a possible role for selection in the shaping of modern pigimentary variation within the continent of Africa.

Body mass index. We next investigate two traits related to metabolic phenotypes (BMI and Type 2 diabetes), as there is a long history of adaptive hypotheses put forward to explain phenotypic variation among populations, with little conclusive evidence emerging thus far. We first focus on the set of 32 BMI associated SNPs identified by Speliotes and colleagues [74] in their Table 1, which explain approximately 1.45% of the total variance for BMI, or about 2–4% of the additive genetic variance. Of these 32 associated SNPs, 28 were present in the HGDP dataset, which we used to calculate a genetic BMI score for each HGDP population. We identified no significant signal of selection at the global level (Table 2).

Our regional level analysis indicated that the mean genetic BMI score is significantly lower than expected in East Asia ($Z=-2.48$, $p=0.01$; see also Figure 5 and Table S7), while marginal Q_X statistics identify excess intraregional variation within East Asia and the Americas (Table 3). While these results are intriguing, given the small fraction of the additive genetic variance

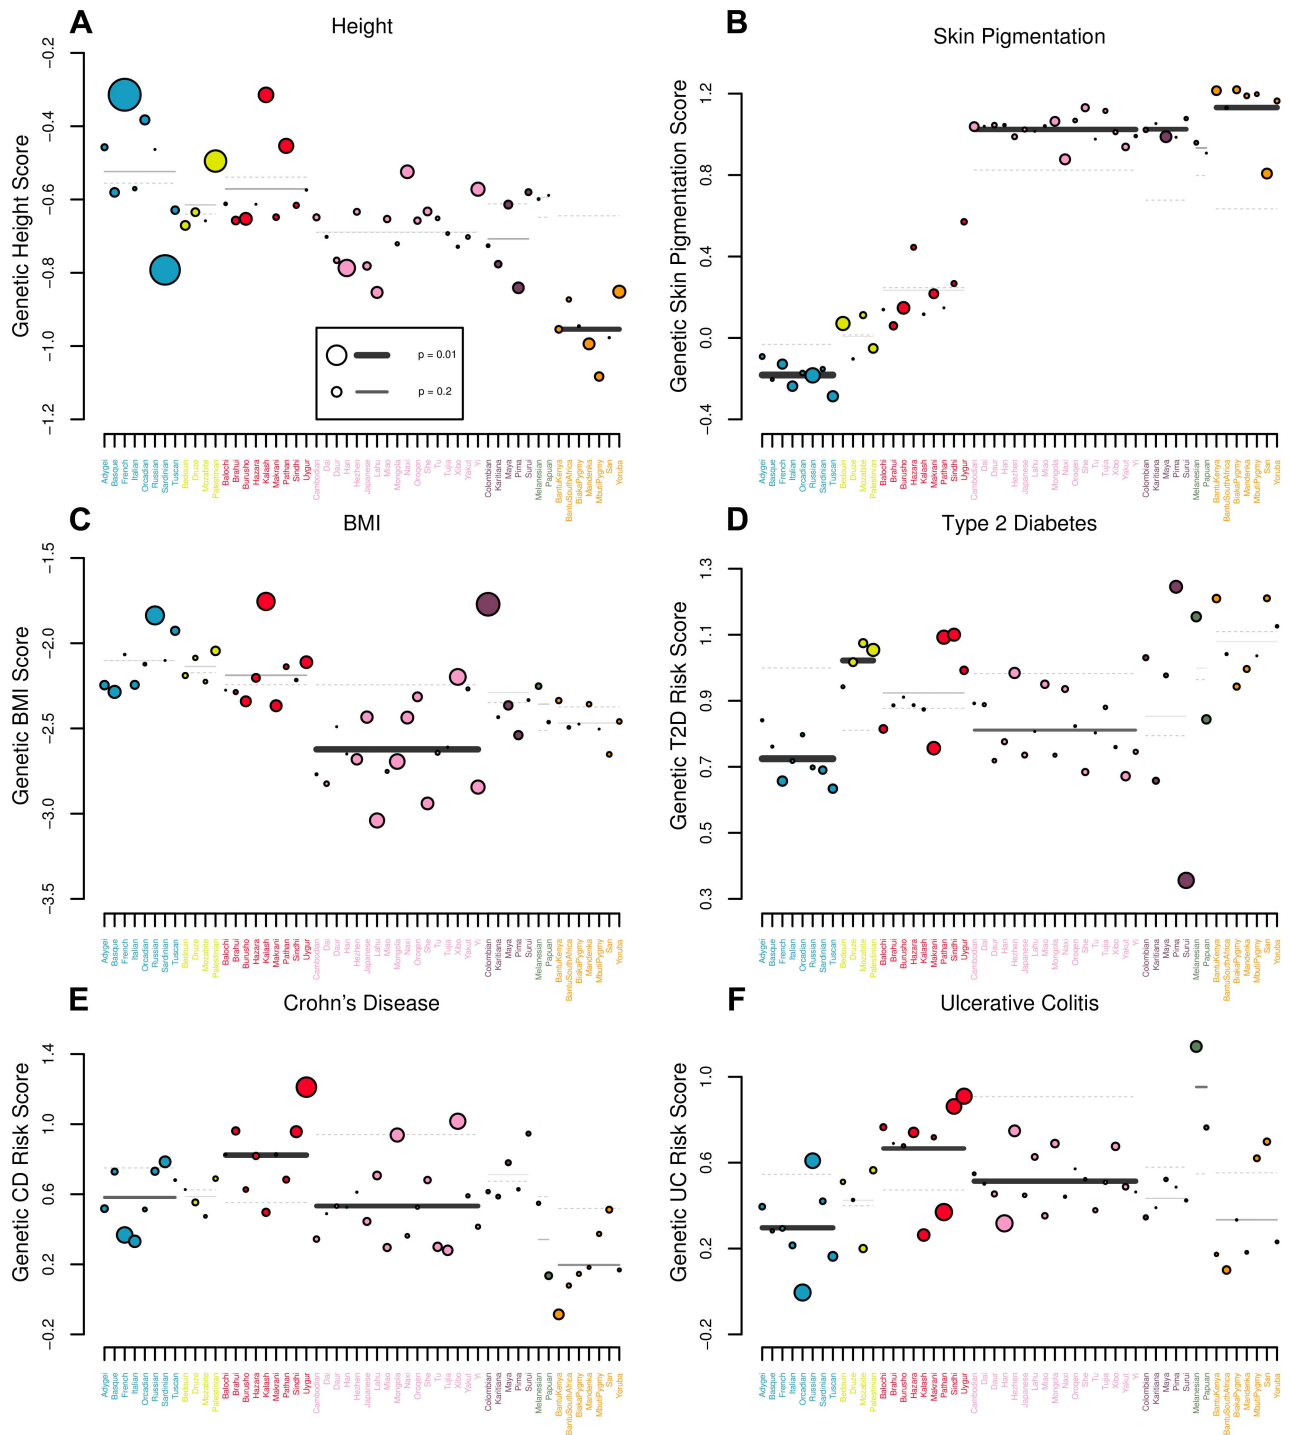


Figure 5. Visual representation of outlier analysis at the regional and individual population level for (A) height, (B) skin pigmentation, (C) body mass index, (D) type 2 diabetes, (E) Crohn’s disease and (F) ulcerative colitis. For each geographic region we plot the expectation of the regional average, given the observed values in the rest of the dataset as a grey dashed line. The true regional average is plotted as a solid bar, with darkness and thickness proportional to the regional Z score. For each population we plot the observed value as a colored circle, with circle size proportional to the population specific Z score. For example, in (A), one can see that estimated genetic height is systematically lower than expected across Africa. Similarly, estimated genetic height is significantly higher (lower) in the French (Sardinian) population than expected, given the values observed for all other populations in the dataset.
doi:10.1371/journal.pgen.1004412.g005

explained by the ascertained SNPs and the lack of a globally significant signal or a clear ecological pattern or explanation, it is difficult to draw strong conclusions from them. For this reason

BMI and other related traits will warrant reexamination as more association results arise and methods for analyzing association results from multiple correlated traits are developed.




Figure 6. Estimated genetic height (A) and skin pigmentation score (B) plotted against winter PC2 and absolute latitude respectively. Both correlations are significant against the genome wide background after controlling for population structure (Table 2). doi:10.1371/journal.pgen.1004412.g006

Type 2 diabetes. We next investigated the 65 loci reported by Morris and colleagues [75] as associated with T2D, which explain 5.7% of the total variance for T2D susceptibility, or about 8–9% of the additive genetic variance. Of these 65 SNPs, 61 were present in the HGDP dataset. We used effect sizes from the stage 1 meta-analysis, and where a range of allele frequencies are reported (due to differing sample frequencies among cohorts), we simply used the average. Where multiple SNPs were reported per locus we used the lead SNP from the combined meta-analysis. Also note that Morris and colleagues report effects in terms odds ratios (OR), which can be converted into additive effects by taking the logarithm (the same is true of the IBD data from [26], analyzed below).

The distribution of genetic T2D risk scores showed no significant correlations with any of the five eco-geographic axes we tested, and was in fact fairly underdispersed worldwide relative to the null expectation due to population structure (Table 2), suggesting we have little to no evidence that differential selection has influenced the distribution of T2D risk across human populations.

We note that our regional level analyses do find that European populations have a significantly lower T2D risk score than expected due to drift ($Z = -2.79$, $p = 0.005$), while Middle Eastern populations have significantly higher risk score than expected ($Z = 2.37$, $p = 0.018$). However, we are skeptical that this represents a meaningful signal of selection for two reasons. The first is that we have probed the data quite deeply despite seeing no evidence for adaptive differentiation at the global level. Second, expanding to the next closest region, an analysis in which we treated regional membership as the linear predictor was unable to find significant differentiation between Central Asia and Europe ($r^2 = 0.13$, $p = 0.16$; $Q_X = 12.0$, $p = 0.75$) or between Central Asia and the Middle East ($r^2 = 0.15$, $p = 0.19$; $Q_X = 9.8$, $p = 0.63$). Our results are therefore consistent at most with a recent, but fairly weak selective event which influenced only European and Middle Eastern populations, but we do not feel our results count as strong evidence for this hypothesis.

A number of investigators have claimed that individual European GWAS loci for Type 2 Diabetes show signals of selection [63,76–78], a fact that may be seen as support for the idea that genetic variation for T2D risk has been shaped by local adaptation, potentially consistent with a variation on the thrifty genotype hypothesis [79]. However, our result suggest that local adaptation has not had a large role in shaping the present day world-wide distribution of T2D susceptibility alleles (as mapped to date in Europe). One explanation of this discrepancy is that it is biologically unrealistic that the phenotype of T2D susceptibility would exhibit strong adaptive differentiation. Rather, local adaptation may have shaped some pleiotropically related phenotype (which shares only some of the loci involved). However, as seen in Figure 2, our methods have better power than single locus statistics so long as there is a reasonable correlation ($\phi > 0.3$) between the focal phenotype and the one under selection. As such, the intersection of our results with previous studies support the idea that local adaptation has had little direct influence on the genetic basis of T2D or closely correlated phenotypes, but that a handful of individual SNPs associated with T2D may have experienced adaptive differentiation as a result of their function in some other phenotype.

IBD. Finally, we analyzed the set of associations reported for Crohn's Disease (CD) and Ulcerative Colitis (UC) [26]. Because CD and UC are closely connected phenotypes that share much of their genetic etiology, Jostins and colleagues used a likelihood ratio test of four different models (CD only, UC only, both CD and UC with equal effects on each, both CD and UC with independent effects) to distinguish which SNPs were associated with either or both phenotypes, and to assign effect sizes to SNPs (see their supplementary methods section 1d). We take these classifications at face value, resulting in two partially overlapping lists of 140 and 135 SNPs associated with CD and UC, which explain 13.6% and 7.5% of disease susceptibility variance respectively. Of these, there are 95 SNPs for CD and 89 SNPs for UC were present in our HGDP dataset, and these remaining SNPs on which our analyses are based explain 9% and 5.1% of the total variance. For now, we

treat these sets of loci independently, and leave the development of methods that appropriately deal with correlated traits for future work.

We used these sets of SNPs to calculate genetic risk scores for CD and UC across the 52 HGDP populations. Both CD and UC showed strong negative correlations with summer PC2 (Figure 7), while CD also showed a significant correlation with winter PC1, and a marginally significant correlation with summer PC1 (Table 2).

We did not observe any significant Q_X statistics for either trait, either at the global or the regional level, suggesting that our environmental correlation signals most likely arise from subtle differences between regions, as opposed to divergence among closely related populations. Indeed, we find moderate signals of regional level divergence in Europe (UC: $Z = -2.08, p = 0.04$), Central Asia (CD: $Z = 2.21, p = 0.03$), and East Asia (CD: $Z = -1.90, p = 0.06$ and UC: $Z = -2.12, p = 0.03$; see also Figure 5 and Tables S13 and S14).

Discussion

In this paper we have developed a powerful framework for identifying the influence of local adaptation on the genetic loci underlying variation in polygenic phenotypes. Below we discuss two major issues related to the application of such methods, namely the effect of the GWAS ascertainment scheme on our inference, and the interpretation of positive results.

Ascertainment and Population Structure

Among the most significant potential pitfalls of our analysis (and the most likely cause of a false positive) is the fact that the loci used to test for the effect of selection on a given phenotype have been obtained through a GWAS ascertainment procedure, which can introduce false signals of selection if potential confounds are not properly controlled. We condition on simple features of the ascertainment process via our allele matching procedure, but deeper issues may arise from artifactual associations that result from the effects of population structure in the GWAS ascertainment panel. Given the importance of addressing this issue to the broader GWAS community, a range of well developed methods exist for doing GWAS in structured populations, and we refer the reader to the existing literature for a full discussion [80–86]. Here, we focus on two related issues. First, the propensity of population structure in the GWAS ascertainment panel to generate false positives in our selection analysis, and second, the difficulties introduced by the sophisticated statistical approaches employed to deal with this issue when GWAS are done in strongly structured populations.

The problem of population structure arises generally when there is a correlation in the ascertainment panel between phenotype and ancestry such that SNPs that are ancestry informative will appear to be associated with the trait, even when no causal relationship exists [81]. This phenomenon can occur regardless of whether the correlation between ancestry and phenotype is caused by genetic or environmental effects. To make matters worse, multiple false positive associations will tend to line up with same axis of population structure. If the populations being tested with our methods lie at least partially along the same axis of structure present in the GWAS ascertainment panel, then the ascertainment process will serve to generate the very signal of positive covariance among like effect alleles that our methods rely on to detect the signal of selection.

The primary takeaway from this observation is that the more diverse the array of individuals sampled for a given GWAS are

with respect to ancestry, the greater the possibility that failing to control for population structure will generate false associations (or bias effect sizes) and hence false positives for our method.

What bearing do these complications have on our empirical results? The GWAS datasets we used can be divided into those conducted within populations of European descent and the skin pigmentation dataset (which used an admixed population). We will first discuss our analysis of the former.

The European GWAS loci we used were found in relatively homogeneous populations, in studies with rigorous standards for replication and control for population structure. Therefore, we are reasonably confident that these loci are true positives. Couple this with the fact that they were ascertained in populations that are fairly homogenous relative to the global scale of our analyses, and it is unlikely that population structure in the ascertainment panels is driving our positive signals. One might worry that we could still generate false signals by including European populations in our analysis, however many of the signals we see are driven by patterns outside of Europe (where the influence of structure within Europe should be much lessened). For height, where we do see a strong signal from within Europe, we use a set of loci that have been independently verified using a family based design that is immune to the effects of population structure [28].

We further note that for a number of GWAS datasets, including some of those analyzed here, studies of non-European populations have replicated many of the loci identified in European populations [87–93], and for many diseases, the failure of some SNPs to replicate, as well as discrepancies in effect size estimate, are likely due to simple considerations of statistical power and differences in patterns of LD across populations [94,95]. This suggests that, at least for GWAS done in relatively homogenous human populations, structure is unlikely to be a major confounding factor.

The issue of population structure may be more profound for our style of approach when GWAS are conducted using individuals from more strongly structured populations. In some cases it is desirable to conduct GWAS in such populations as locally adaptive alleles will be present at intermediate frequencies in these broader samples, whereas they may be nearly fixed in more homogeneous samples. A range of methods have been developed to adjust for population structure in these setting [96–98]. While generally effective in their goal, these methods present their own issues for our selection analysis. Consider the extreme case, such as that of Atwell et al (2010) [19], who carried out a GWAS in *Arabidopsis thaliana* for 107 phenotypes across an array of 183 inbred lines of diverse geographical and ecological origin. Atwell and colleagues used the genome-wide mixed model program EMMA [83,96,97] to control for the complex structure present in their ascertainment panel. This practice helps ensure that many of the identified associations are likely to be real, but also means that the loci found are likely to have unusual frequencies patterns across the species range. This follows from the fact that the loci identified as associated with the trait must stand out as being correlated with the trait in a way not predicted by the individual kinship matrix (as used by EMMA and other mixed model approaches). Our approach is predicated on the fact that we can use genome-wide patterns of kinship to adjust for population structure, but this correction is exactly the null model that loci significantly associated with phenotypes by mixed models have overcome. For this reason, both the theoretical χ^2 distribution of the Q_X statistic, as well as the empirical null distributions we construct from resampling, may be inappropriate.

The Cape Verde skin pigmentation data we used may qualify as this second type of study. The Cape Verde population is an

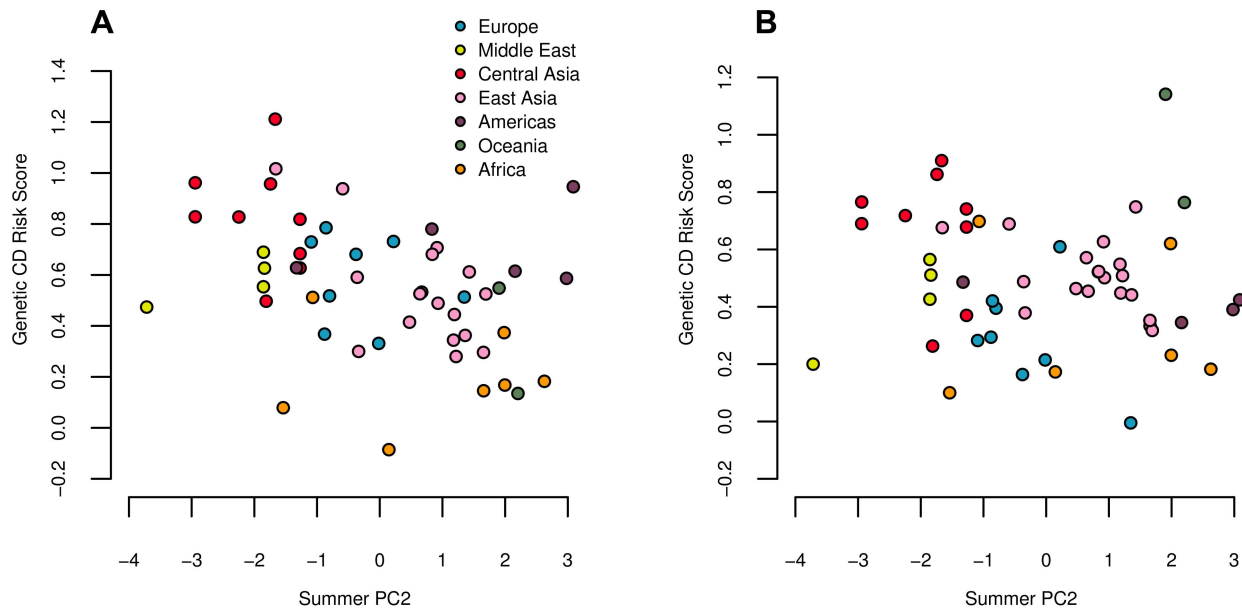


Figure 7. Estimated genetic risk score for Crohn's disease (A) and ulcerative colitis (B) risk plotted against summer PC2. Both correlations are significant against the genome wide background after controlling for population structure (Table 2). Since a large proportion of SNPs underlying these traits are shared, we note that these results are not independent. doi:10.1371/journal.pgen.1004412.g007

admixed population of African/European descent, and has substantial inter-individual variation in admixture proportion. Due to its admixed nature, the population segregates alleles which would not be at intermediate frequency in either parental population, making it an ideal mapping population.

Despite the considerable population structure, the fact that recombination continues to mix genotypes in this population means that much of the LD due to the African/European population structure has been broken up (and the remaining LD is well predicted by an individual's genome-wide admixture coefficient). Population structure seems to have been well controlled for in this study, and a number of the loci have been replicated in independent admixed populations. While we think it unlikely that the four loci we use are false associations, they could in principle suffer from the structured ascertainment issues described above, so it is unclear that the null distributions we use are strictly appropriate. That said, provided that Beza and colleagues have appropriately controlled for population structure, under neutrality there would be no reason to expect that the correlation among the loci should be strongly positive with respect to the sign of their effect on the phenotype, and thus the pattern observed is at least consistent with a history of selection, especially in light of the multiple alternative lines of evidence for adaptation on the basis of skin pigmentation [68–70,99–101].

Further work is needed to determine how best to modify the tests proposed herein to deal with GWAS performed in structured populations.

Complications of Interpretation

Our understanding of the genetic basis of variation in complex traits remains very incomplete, and as such the results of these analyses must be interpreted with caution. That said, because our methods are based simply on the rejection of a robust, neutral null model, an incomplete knowledge of the genetic basis of a given trait should only lead to a loss of statistical power, and not to a high false positive rate.

For all traits analyzed here except for skin pigmentation, the within population variance for genetic value is considerably larger than the variance between populations. This suggests that much of what we find is relatively subtle adaptation even on the level of the phenotype, and emphasizes the fact that for most genetic and phenotypic variation in humans, the majority of the variance is within populations rather than between populations (see Figures S14–S19). In many cases, the influence of the environment likely plays a stronger role in the differences between populations for true phenotypes than the subtle differences we find here (as demonstrated by the rapid change in T2D incidence with changing diet, e.g. [102]). That said, an understanding of how adaptation has shaped the genetic basis of a wide variety of phenotypes is clearly of interest, even if environmental differences dominate as the cause of present day population differences, as it informs our understanding of the biology and evolutionary history of these traits.

The larger conceptual issues relate to the interpretation of our positive findings, which we detail below. A number of these issues are inherent to the conceptual interpretation of evidence for local adaptation [103].

Effect size heterogeneity and misestimation. In all of our analyses, we have simply extrapolated GWAS effect sizes measured in one population and one environment to the entire panel of HGDP populations. It is therefore prudent to consider the validity of this assumption, as well as the implications for our analyses when it is violated. Aside from simple measurement error, there are two possible reasons that estimated effect sizes from GWAS may not reflect the true effect sizes.

The first is that most GWAS hits likely identify tag SNPs that are in strong LD with causal sites that are physically nearby on the chromosome, rather than actual causal sites themselves [94,95]. This acts to reduce the estimated effect size in the GWAS sample. More importantly for the interpretation of our signals, patterns of LD between tag SNPs and causal sites will change over evolutionary time, and so a tag SNP's allele frequencies will be

an imperfect measure of the differentiation of the causal SNP over the sampled populations. This should lead to a reduction in our power to detect the effect of selection in much the same way that power is reduced when selection acts on a trait that is genetically correlated with the trait of interest (Figure 1B). This effect will be especially pronounced when the populations under study have a shorter scale of LD than the populations in which the effect have been mapped (e.g. when applying effect sizes estimated in Europe to populations of African descent). In the case that selection has not affected the trait of interest, the effect sizes have no association whatsoever with the distribution of allele frequencies across populations unless such an association is induced by the ascertainment process, as described above. Therefore, changes in the patterns of LD between identified tag SNPs and causal sites will not lead to an excess of false positives if the loci under study have not been subject to spatially varying selection pressures.

The second is that the actual value of the additive effect at a causal site may change across environments and genetic backgrounds due to genotype-by-genotype (i.e. functional epistasis) and genotype-by-environment interactions. Although the response at a given locus due to selection depends only the additive effect of the allele in that generation, the additive effect itself is a function of the environment and the frequencies of all interacting loci. As all of these can change considerably during the course of evolution, the effects estimated in one population may not apply in other populations, either in the present day, or over history of the populations [1,104]. We first wish to stress that, as above, because our tests rely on rejection of a null model of drift, differences in additive effects among populations or over time will not lead to an excess of false positives, provided that the trait is truly neutral. Such interactions can, however, considerably complicate the interpretation of positive results. For example, different sets of alleles could be locally selected to maintain a constant phenotype across populations due to gene-by-environment interactions. Such a scenario could lead to a signal of local adaptation on a genetic level but no change in the phenotype across populations, a phenomenon known as countergradient variation [105].

It will be very difficult to know how reasonable it is to extrapolate effect sizes among populations without repeating measurements in different populations and different environments. Perhaps surprisingly, the existing evidence suggests that for a variety of highly polygenic traits, effects sizes and directions may be relatively consistent across human populations [87–95]. There is no particular reason to believe that this will hold as a general rule across traits or across species, and thus addressing this issue will require a great deal more functional genetic work and population genetic method development, a topic which we discuss briefly below in Future Directions.

Missing variants. As the majority of GWAS studies are performed in a single population they will often miss variants contributing to phenotypic variation. This can occur due to GxG or GxE interactions as outlined above, but also simply because those variants are absent (or at low frequency) due to drift or selection among the populations. Such cases will not create a false signal of selection if only drift is involved, however, they do complicate the interpretation of positive signals. A particularly dramatic example of this is offered by our analysis of skin pigmentation associated loci, whose frequencies are clearly shaped by adaptation. The alleles found by a GWAS in the Cape Verde population completely fail to predict the skin pigmentation of East Asians and Native Americans. This reflects the fact that a number of the alleles responsible for light skin pigmentation in those populations are not variable in Cape Verde due to the partially convergent adaptive evolution of light skin pigmentation [71]. As a

result, when we take the Eurasian HGDP populations we see a significant correlation between genetic skin pigmentation score and longitude ($r^2=0.15, p=0.015$), despite the fact that no such phenotypic correlation exists. While the wrong interpretation is easy to avoid here because we have a good understanding of the true phenotypic distribution, for the majority of GWAS studies such complications will be subtler and so care will have to be taken in the interpretation of positive results.

Loss of constraint and mutational pressure. One further complication in the interpretation of our results is in how loss of constraint may play a role in driving apparent signals of local adaptation. Traits evolving under uniform stabilizing selection across all populations should be less variable than predicted by our covariance model of drift, due to negative covariances among loci, and so should be underrepresented in the extreme tails of our environmental correlation statistics and the upper tail of Q_X . As such, loss of constraint (i.e. weaker stabilizing selection in some populations than others), should not on its own create a signal of local adaptation. While the loci underpinning the phenotype can be subject to more drift in those populations, there is no systemic bias in the direction of this drift. Loss of constraint, therefore, will not tend to create significant environmental correlations or systematic covariance between alleles of like effect.

An issue may arise, however, when loss of constraint is paired with biased mutational input (i.e. new mutations are more likely to push the phenotype in one direction than another [106]) or asymmetric loss of constraint (selection is relaxed on one tail of the phenotypic distribution). Under these two scenarios, alleles that (say) increase the phenotype would tend to drift up in frequency in the populations with loss of constraint, creating systematic trends and positive covariance among like effect alleles at different loci, and resulting in a positive signal under our framework. While one would be mistaken to assume that the signal was necessarily that of recent positive directional selection, these scenarios do still imply that selection pressures on the genetic basis of the phenotype vary across space. Positive tests under our methods are thus fairly robust in being signals of differential selection among populations, but are themselves agnostic about the specific processes involved. Further work is needed to establish whether these scenarios can be distinguished from recent directional selection based on only allele frequencies and effect sizes, and as always, claims of recent adaptation should be supported by multiple lines of evidence beyond those provided by population genomics alone.

Future directions. In this article we have focused on methods development and so have not fully explored the range of populations and phenotypes to which our methods could be applied. Of particular interest is the possibility of applying these methods to GWAS performed in other species where the ecological determinants of local adaptation are better understood [19,20].

One substantial difficulty with our approach, particularly in its application to other organisms, is that genome-wide association studies of highly polygenic phenotypes require very large sample sizes to map even a fraction of the total genetic variance. One promising way to partially sidestep this issue is by applying methods recently developed in animal and plant breeding. In these genomic prediction/selection approaches, one does not attempt to map individual markers, but instead concentrates on predicting an individual's genetic value for a given phenotype using all markers simultaneously [107–109]. This is accomplished by fitting simple linear models to genome-wide genotyping data, in principle allowing common SNPs to tag the majority of causal sites throughout the genome without attempting to explicitly identify them [110]. These methods have been applied to a range of

species, including humans [111–117], demonstrating that these predictions can potentially explain a relatively high fraction of the additive genetic variance within a population (and hence much of the total genetic variance). As these predictions are linear functions of genotypes, and hence allele frequencies, we might be able to predict the genetic values of sets of closely related populations for phenotypes of interest and apply very similar methods to those developed here. Such an approach may allow for substantial gains in power, as it would greatly increase the fraction of the genetic variance used in the analyses. However, if the only goal is to establish evidence for local adaptation in a given phenotype, then because measurements of true phenotypes inherently include all of the underlying loci, the optimal approach is to perform a common garden experiment and employ statistical methods such as those developed by Ovaskainen and colleagues [40,41,118], assuming such experiments can be done.

As discussed in various places above, it is unlikely that all of the loci underpinning the genetic basis of a trait will have been subject to the same selection pressures, due to their differing roles in the trait and their pleiotropic effects. One potential avenue of future investigation is whether, given a large set of loci involved in a trait, we can identify sets of loci in particular pathways or with a particular set of functional attributes that drive the signal of selection on the additive genetic basis of a trait.

Another promising extension of our approach is to deal explicitly with multiple correlated phenotypes. With the increasing number of GWAS efforts both empirical and methodological work are beginning to focus on understanding the shared genetic basis of various phenotypes [26,119]. This raises the possibility that we may be able to disentangle the genetic basis of which phenotypes are more direct targets of selection, and which are responding to correlated selection on these direct targets (for progress along these lines using Q_{ST} , see [40,120–122]). Such tools may also offer a way of incorporating G×E interactions, as multiple GWAS for the same trait in different environments can be treated as correlated traits [123].

As association data for a greater variety of populations, species, and traits becomes available, we view the methods described out here as a productive way forward in developing a quantitative framework to explore the genetic and phenotypic basis of local adaptation.

Materials and Methods

Mean Centering and Covariance Matrix Estimation

Written in matrix notation, the procedure of mean centering the estimated genetic values and dropping one population from the analysis can be expressed as

$$\tilde{\mathbf{Z}}' = \mathbf{T}\tilde{\mathbf{Z}} \quad (16)$$

where \mathbf{T} is an $M-1$ by M matrix with $\frac{M-1}{M}$ on the main diagonal, and $-\frac{1}{M}$ elsewhere.

In order to calculate the corresponding expected neutral covariance structure about this mean, we use the following procedure. Let \mathbf{G} be an M by K matrix, where each column is a vector of allele frequencies across the M populations at a particular SNP, randomly sampled from the genome according to the matching procedure described below. Let ϵ_k and ϵ_i be the mean allele frequency in columns k and i of \mathbf{G} respectively, and let \mathbf{S} be a matrix such that $s_{ij} = \frac{1}{\epsilon_i(1-\epsilon_i)}$. With these data, we can

estimate \mathbf{F} as

$$\mathbf{F} = \frac{1}{K-1} \mathbf{TGS}\mathbf{G}^T\mathbf{T}^T. \quad (17)$$

This transformation performs the operation of centering the matrix at the mean value, and rooting the analysis with one population by dropping it from the covariance matrix (the same one we dropped from the vector of estimated genetic values), resulting in a covariance matrix describing the relationship of the remaining $M-1$ populations. This procedure thus escapes the singularity introduced by centering the matrix at the observed mean of the sample.

As we do not get to observe the population allele frequencies, the entries of \mathbf{G} are the sample frequencies at the randomly chosen loci, and thus the covariance matrix \mathbf{F} also includes the effect of finite sample size. Because the noise introduced by the sampling of individuals is uncorrelated across populations (in contrast to that introduced by drift and shared history), the primary effect is to inflate the diagonal entries of the matrix by a factor of $\frac{1}{n_m}$, where n_m is the number of chromosomes sampled in population m (see the supplementary material of [46] for discussion). This means that our population structure adjusted statistics also approximately control for differences in sample size.

Standardized environmental variable. Given a vector of environmental variable measurements for each population, we apply both the \mathbf{T} and Cholesky transformation as for the estimated genetic values

$$\tilde{\mathbf{Y}}' = \mathbf{C}^{-1}\mathbf{T}\tilde{\mathbf{Y}}. \quad (18)$$

This provides us with a set of $M-1$ adjusted observations for the environmental variable which can be compared to the transformed genetic values for inference. This step is necessary as we have rotated the frame of reference of the estimated genetic values, and so we must do the same for the environmental variables to keep them both in a consistent reference frame.

Identifying Outliers with Conditional MVN Distributions

As described in the Results, we can use our multivariate normal model of relatedness to obtain the expected distribution of genetic values for an arbitrary set of populations, conditional on the observed values in some other arbitrary set.

We first partition our populations into two groups, those for which we want to obtain the expected distribution of genetic values (group 1), and those on which we condition in order to obtain this distribution (group 2). We then re-estimate the covariance matrix such that it is centered on the mean of group 2. This step is necessary because the amount of divergence between the populations in group 1 and the mean of group 2 will always be greater than the amount of divergence from the global mean, even under the neutral model, and our covariance matrix needs to reflect this fact in order to make accurate predictions. We can obtain this re-parameterized \mathbf{F} matrix as follows. If M is the total number of populations in the sample, then let q be the number of populations in group one, and let $M-q$ be the number of populations in group 2. We then define a new \mathbf{T}_R matrix such that the q columns corresponding the populations in group one have 1 on the diagonal, and 0 elsewhere, while the $M-q$ columns corresponding to group two have $\frac{M-q-1}{M-q}$ on the diagonal, and

− $\frac{1}{M-q}$ elsewhere. We can then re-estimate a covariance matrix that is centered at the mean of the $M-q$ populations in group 2. Recalling our matrices \mathbf{G} and \mathbf{S} from (17), this matrix is calculated as

$$\mathbf{F}_R = \frac{1}{K-1} \mathbf{T}_R \mathbf{G} \mathbf{S} \mathbf{G}^T \mathbf{T}_R^T \quad (19)$$

where we write \mathbf{F}_R to indicate that it is a covariance matrix that has been re-centered on the mean of group two.

Once we have calculated this re-centered covariance matrix, we can use well known results from multivariate normal theory to obtain the expected joint distribution of the genetic values for group one, conditional on the values observed in group two.

We partition our vector of genetic values and the re-centered covariance matrix such that

$$\vec{X} = \begin{bmatrix} \vec{X}_1 \\ \vec{X}_2 \end{bmatrix} \quad (20)$$

and

$$\mathbf{F}_R = \begin{bmatrix} \mathbf{F}_{11} & \mathbf{F}_{12} \\ \mathbf{F}_{21} & \mathbf{F}_{22} \end{bmatrix} \quad (21)$$

where \vec{X}_1 and \vec{X}_2 are vectors of genetic values in group 1 and 2 respectively, and \mathbf{F}_{11} , \mathbf{F}_{22} and $\mathbf{F}_{12} = \mathbf{F}_{21}^T$ are the marginal covariance matrices of populations within group 1, within group 2, and across the two groups, respectively. Letting $\mu_1 = \mu_2 = \frac{1}{M-q} \sum_{m=M-q}^M X_m$ (i.e. the sum of the elements of \vec{X}_2), we wish to obtain the distribution

$$\vec{X}_1 | \vec{X}_2, \mu_1, \mu_2 \sim MVN(\vec{\xi}, \mathbf{\Omega}), \quad (22)$$

where $\vec{\xi}$ and $\mathbf{\Omega}$ give the expected means and covariance structure of the populations in group 1, conditional on the values observed in group 2. These can be calculated as

$$\vec{\xi} = \mathbb{E}[\vec{X}_1 | \vec{X}_2, \mu_1, \mu_2] = \mu_1 \vec{1} + \mathbf{F}_{12} \mathbf{F}_{22}^{-1} (\vec{X}_2 - \mu_2 \vec{1}) \quad (23)$$

and

$$\mathbf{\Omega} = \text{Cov}[X_1 | X_2, \mu_1, \mu_2] = \mathbf{F}_{11} - \mathbf{F}_{12} \mathbf{F}_{22}^{-1} \mathbf{F}_{21}. \quad (24)$$

where the one vectors in line (23) are of length q and $M-q$ respectively.

This distribution is itself multivariate normal, and as such this framework is extremely flexible, as it allows us to obtain the expected joint distribution for arbitrary sets of populations (e.g. geographic regions or continents), or for each individual population. Further,

$$\mathbb{E} \left[\frac{1}{q} \sum_{m=1}^q X_m | \vec{\xi}, \mathbf{\Omega} \right] = \frac{1}{q} \sum_{m=1}^q \xi_m \quad (25)$$

and

$$\text{Var} \left[\frac{1}{q} \sum_{m=1}^q X_m | \vec{\xi}, \mathbf{\Omega} \right] = \frac{V_A}{q^2} \sum_{m=1}^q \sum_{n=1}^q \omega_{mn}. \quad (26)$$

where ω_{nm} denotes the elements of $\mathbf{\Omega}$. In words, the conditional expectation of the mean estimated genetic value across group 1 is equal to the mean of the conditional expectations, and its variance is equal to the mean value of the elements of the conditional covariance matrix. As such we can easily calculate a Z score (and corresponding p value) for group one as a whole as

$$Z = \frac{\frac{1}{q} \sum_{m=1}^q X_m - \frac{1}{q} \sum_{m=1}^q \xi_m}{\frac{1}{q} \sqrt{V_A \sum_{m=1}^q \sum_{n=1}^q \omega_{m,n}}}. \quad (27)$$

This Z score is a normal random variable with mean zero, variance one under the null hypothesis, and thus measures the divergence of the genetic values between the two populations relative to the null expectation under drift. Note that the observation of a significant Z score in a given population or region cannot necessarily be taken as evidence that selection has acted in that population or region, as selection in some of the populations on which we condition (especially the closely related ones) could be responsible for such a signal. As such, caution is warranted when interpreting the output of these sort of analyses, and is best done in the context of more explicit information about the demographic history, geography, and ecology of the populations.

The Linear Model at the Individual Locus Level

As with our excess variance test, explored in the main text, it is natural to ask how our environmental correlation tests can be written in terms of allele frequencies at individual loci.

As noted in (8), we can obtain for each underlying locus a set of transformed allele frequencies, which have passed through the same transformation as the estimated genetic values. We assume that each locus ℓ has a regression coefficient

$$\beta_\ell = \gamma \alpha_\ell \quad (28)$$

where γ is shared across all loci so that

$$p'_{m\ell} \sim \gamma \alpha_\ell Y_m + e_{m\ell} \quad (29)$$

where the $e_{m\ell}$ are independent and identically distributed residuals. We can find the maximum likelihood estimate $\hat{\gamma}$ by treating $\alpha_\ell Y'_m$ as the linear predictor, and taking the regression of the combined vector \vec{p}' , across all populations and loci, on the combined vector $\vec{\alpha Y'}$. As such

$$\hat{\gamma} = \frac{\text{Cov}(p', \alpha Y')}{\text{Var}(\alpha Y')} \quad (30)$$

we can decompose this into a sum across loci such that

$$\hat{\gamma} = \frac{\frac{1}{L} \sum_{\ell} \text{Cov}(p'_{\ell}, \alpha_{\ell} Y')}{\frac{1}{L} \sum_{\ell} \text{Var}(\alpha_{\ell} Y')} = \frac{1}{\sum_{\ell} \alpha_{\ell}^2} \frac{\sum_{\ell} \alpha_{\ell} \text{Cov}(p'_{\ell}, Y')}{\text{Var}(Y')} \quad (31)$$

As noted in (8), our transformed genetic values can be written as

$$X_m = 2 \sum_{\ell} \alpha_{\ell} p'_{m\ell} \quad (32)$$

and so the estimated slope ($\hat{\beta}$) of our regression ($\bar{X} = \beta \bar{Y}' + \bar{e}$) is

$$\hat{\beta} = \frac{\text{Cov}(X, Y')}{\text{Var}(Y')} = \frac{2 \sum_{\ell} \alpha_{\ell} \text{Cov}(p'_{\ell}, Y')}{\text{Var}(Y')} \quad (33)$$

Comparing these equations, the mean regression coefficient at the individual loci (31) and the regression coefficient of the estimated genetic values (33) are proportional to each other via a constant that is given by one over two times the sum of the effect sizes squared (i.e. $\gamma = \frac{1}{2 \sum_{\ell} \alpha_{\ell}^2} \beta$). Our test based on estimating the

regression of genetic values on the environmental variable is thus mathematically equivalent to an approach in which we assume that the regression coefficients of individual loci on the environmental variable are proportional to one another via a constant that is a function of the effect sizes. Such a relationship can also be demonstrated for the correlation coefficient (r^2) calculated at the genetic value level and at the individual locus level (this is not necessarily true for the rank correlation ρ), however the algebra is more complicated, and thus we do not show it here.

This is in contrast to the r^2 enrichment statistic we compute for the power simulations, in which we assume that the correlations of individual loci with the environmental variable are independent of one another, and then perform a test for whether more loci individually show strong correlations with the environmental variable than we would expect by chance.

HGDP Data and Imputation

We used imputed allele frequency data in the HGDP, where the imputation was performed as part of the phasing procedure of [59], as per the recommendations of [124]. We briefly recap their procedure here:

Phasing and imputation were done using fastPHASE [125], with the settings that allow variation in the switch rate between subpopulations. The populations were grouped into subpopulations corresponding to the clusters identified in [62]. Haplotypes from the HapMap YRI and CEU populations were included as known, as they were phased in trios and are highly accurate. HapMap JPT and CHB genotypes were also included to help with the phasing.

Choosing Null SNPs

Various components of our procedure involve sampling random sets of SNPs from across the genome. While we control for biases in our test statistics introduced by population structure through

our \mathbf{F} matrix, we are also concerned that subtle ascertainment effects of the GWAS process could lead to biased test statistics, even under neutral conditions. We control for this possibility by sampling null SNPs so as to match the joint distribution of certain properties of the ascertained GWAS SNPs. Specifically, we were concerned that the minor allele frequency (MAF) in the ascertainment population, the imputation status of the allele in the HGDP datasets, and the background selection environment experienced at a given locus, as measured by B value [61], might influence the distribution of allele frequencies across populations in ways that we could not predict.

We partitioned SNPs into a three way contingency table, with 25 bins for MAF (i.e. a bin size of 0.02), 2 bins for imputation (either imputed or not), and 10 bins for B value (B values range from 0 to 1, and thus our bin size was 0.1). For each set of null genetic values, we sampled one null SNP from the same cell in the contingency table as each of the GWAS SNPs, and assigned this null SNP the effect size associated with the GWAS SNP it was sampled to match. While we do not assign effect sizes to sampled SNPs used to estimate the covariance matrix \mathbf{F} (instead simply scaling \mathbf{F} by a weighted sum of squared effect sizes, which is mathematically equivalent under our assumption that all SNPs have the same covariance matrix), we follow the same sampling procedure to ensure that \mathbf{F} describes the expected covariance structure of the GWAS SNPs.

For the skin pigmentation GWAS [67] we do not have a good proxy present in the HGDP population, as the Cape Verdeans are an admixed population. Cape Verdeans are admixed with ~59.53% African ancestry, and 41.47% European ancestry in the sample obtained by [67] (Beleza, pers. comm., April 8, 2013). As such, we estimated genome wide allele frequencies in Cape Verde by taking a weighted mean of the frequencies in the French and Yoruban populations of the HGDP, such that $p_{CV} = 0.5953p_Y + 0.4147p_F$. We then used these estimated frequencies to assign SNPs to frequency bins.

[67] also used an admixture mapping strategy to map the genetic basis of skin pigmentation. However, if they had only mapped these loci in an admixture mapping setting we would have to condition our null model on having strong enough allele frequency differentiation between Africans and Europeans at the functional loci for admixture mapping to have power [126]. The fact that [67] mapped these loci in a GWAS framework allows us to simply reproduce the strategy, and we ignore the results of the admixture mapping study (although we note that the loci and effect sizes estimated were similar). This highlights the need for a reasonably well defined ascertainment population for our approach, a point which we comment further on in the Discussion.

Supporting Information

Figure S1 Power of tests described in the main text to detect a signal of selection on the mapped genetic basis of skin pigmentation [67] as an increasing function of the strength of selection (A), and a decreasing function of the genetic correlation between skin pigmentation and the selected trait with the effect of selection held constant at $\delta = 0.13$ (B).

(TIFF)

Figure S2 Power of tests described in the main text to detect a signal of selection on the mapped genetic basis of BMI [74] as an increasing function of the strength of selection (A), and a decreasing function of the genetic correlation between BMI and the selected trait with the effect of selection held constant at $\delta = 0.07$ (B).

(TIFF)

Figure S3 Power of tests described in the main text to detect a signal of selection on the mapped genetic basis of T2D [75] as an increasing function of the strength of selection (A), and a decreasing function of the genetic correlation between height and the selected trait with the effect of selection held constant at $\delta = 0.08$ (B).

(TIFF)

Figure S4 Power of tests described in the main text to detect a signal of selection on the mapped genetic basis of CD [26] as an increasing function of the strength of selection (A), and a decreasing function of the genetic correlation between CD and the selected trait with the effect of selection held constant at $\delta = 0.05$ (B).

(TIFF)

Figure S5 Power of tests described in the main text to detect a signal of selection on the mapped genetic basis of UC [26] as an increasing function of the strength of selection (A), and a decreasing function of the genetic correlation between UC and the selected trait with the effect of selection held constant at $\delta = 0.05$ (B).

(TIFF)

Figure S6 The two components of Q_X for the skin pigmentation dataset, as described by the left and right terms in (14). The null distribution of each component is shown as a histogram. The expected value is shown as a black bar, and the observed value as a red arrow.

(TIFF)

Figure S7 The two components of Q_X for the BMI dataset, as described by the left and right terms in (14). The null distribution of each component is shown as a histogram. The expected value is shown as a black bar, and the observed value as a red arrow.

(TIFF)

Figure S8 The two components of Q_X for the T2D dataset, as described by the left and right terms in (14). The null distribution of each component is shown as a histogram. The expected value is shown as a black bar, and the observed value as a red arrow.

(TIFF)

Figure S9 The two components of Q_X for the CD dataset, as described by the left and right terms in (14). The null distribution of each component is shown as a histogram. The expected value is shown as a black bar, and the observed value as a red arrow.

(TIFF)

Figure S10 The two components of Q_X for the UC dataset, as described by the left and right terms in (14). The null distribution of each component is shown as a histogram. The expected value is shown as a black bar, and the observed value as a red arrow.

(TIFF)

Figure S11 The genetic values for height in each HGDP population plotted against the measured sex averaged height taken from [127]. Only the subset of populations with an appropriately close match in the named population in [127]'s Appendix I are shown, values used are given in Supplementary table S1.

(TIFF)

Figure S12 The genetic skin pigmentation score for a each HGDP population plotted against the HGDP populations values on the skin pigmentation index map of Biasutti 1959. Data obtained from Supplementary table of [69]. Note that Biasutti

map is interpolated, and so values are known to be imperfect. Values used are given in Supplementary table S2.

(TIFF)

Figure S13 The genetic skin pigmentation score for a each HGDP population plotted against the HGDP populations values from the [68] mean skin reflectance (685nm) data (their Table 6). Only the subset of populations with an appropriately close match were used as in the Supplementary table of [69]. Values and populations used are given in Table S2.

(TIFF)

Figure S14 The distribution of genetic height score across all 52 HGDP populations. Grey bars represent the 95% confidence interval for the genetic height score of an individual randomly chosen from that population under Hardy-Weinberg assumptions.

(TIFF)

Figure S15 The distribution of genetic skin pigmentation score across all 52 HGDP populations. Grey bars represent the 95% confidence interval for the genetic skin pigmentation score of an individual randomly chosen from that population under Hardy-Weinberg assumptions.

(TIFF)

Figure S16 The distribution of genetic BMI score across all 52 HGDP populations. Grey bars represent the 95% confidence interval for the genetic BMI score of an individual randomly chosen from that population under Hardy-Weinberg assumptions.

(TIFF)

Figure S17 The distribution of genetic T2D risk score across all 52 HGDP populations. Grey bars represent the 95% confidence interval for the genetic T2D risk score of an individual randomly chosen from that population under Hardy-Weinberg assumptions.

(TIFF)

Figure S18 The distribution of genetic CD risk score across all 52 HGDP populations. Grey bars represent the 95% confidence interval for the genetic CD risk score of an individual randomly chosen from that population under Hardy-Weinberg assumptions.

(TIFF)

Figure S19 The distribution of genetic UC risk score across all 52 HGDP populations. Grey bars represent the 95% confidence interval for the genetic UC risk score of an individual randomly chosen from that population under Hardy-Weinberg assumptions.

(TIFF)

Table S1 Genetic height scores as compared to true heights for populations with a suitably close match in the dataset of [127]. See Figure S11 for a plot of genetic height score against sex averaged height.

(PDF)

Table S2 Genetic skin pigmentation score as compared to values from Biasutti [69,128] and [68]. We also calculate a genetic skin pigmentation score including previously reported associations at KITLG and OCA2 for comparison. See also Figures S12 and S13.

(PDF)

Table S3 Conditional analysis at the regional level for the height dataset.

(PDF)

Table S4 Conditional analysis at the individual population level for the height dataset.

(PDF)

Table S5 Conditional analysis at the regional level for the skin pigmentation dataset. (PDF)

Table S6 Conditional analysis at the individual population level for the skin pigmentation dataset. (PDF)

Table S7 Conditional analysis at the regional level for the BMI dataset. (PDF)

Table S8 Conditional analysis at the individual population level for the BMI dataset. (PDF)

Table S9 Conditional analysis at the regional level for the T2D dataset. (PDF)

Table S10 Conditional analysis at the individual population level for the T2D dataset. (PDF)

Table S11 Conditional analysis at the regional level for the CD dataset. (PDF)

Table S12 Conditional analysis at the individual population level for the CD dataset. (PDF)

Table S13 Conditional analysis at the regional level for the UC dataset. (PDF)

Table S14 Conditional analysis at the individual population level for the UC dataset. (PDF)

Table S15 Corresponding β statistics for all analyses presented in Table 2. (PDF)

Table S16 Corresponding ρ statistics for all analyses presented in Table 2. (PDF)

Acknowledgments

We would like to thank Gideon Bradburd, Yaniv Brandvain, Luke Jostins, Chuck Langley, Joe Pickrell, Jonathan Pritchard, Peter Ralph, Jeff Ross-Ibarra, Alisa Sedghifar, Michael Turelli and Michael Whitlock for helpful discussion and/or comments on earlier versions of the manuscript. We thank Josh Schraiber and Otso Ovaskainen for useful discussions via Haldane's Sieve.

Author Contributions

Conceived and designed the experiments: JJB GC. Performed the experiments: JJB. Analyzed the data: JJB. Contributed reagents/materials/analysis tools: GC. Wrote the paper: JJB GC.

References

- Fisher RA (1918) XV.—The Correlation between Relatives on the Supposition of Mendelian Inheritance. *Transactions of the Royal Society of Edinburgh* 52: 399–433.
- Provine WB (2001) *The Origins of Theoretical Population Genetics*. With a New Afterword. University Of Chicago Press.
- Turelli M, Barton NH (1990) Dynamics of polygenic characters under selection. *Theoretical Population Biology* 38: 1–57.
- Slate J (2005) Quantitative trait locus mapping in natural populations: progress, caveats and future directions. *Molecular Ecology* 14: 363–379.
- Kingsolver JG, Hoekstra HE, Hoekstra JM, Berrigan D, Vignieri SN, et al. (2001) The Strength of Phenotypic Selection in Natural Populations. *The American Naturalist* 157: 245–261.
- Hudson RR, Kreitman M, Aguadé M (1987) A test of neutral molecular evolution based on nucleotide data. *Genetics* 116: 153–159.
- McDonald JH, Kreitman M (1991) Adaptive protein evolution at the Adh locus in *Drosophila*. *Nature* 351: 652–654.
- Begun DJ, Aquadro CF (1992) Levels of naturally occurring DNA polymorphism correlate with recombination rates in *D. melanogaster*. *Nature* 356: 519–520.
- Nielsen R, Williamson S, Kim Y, Hubisz MJ, Clark AG, et al. (2005) Genomic scans for selective sweeps using SNP data. *Genome Research* 15: 1566–1575.
- Latta RG (1998) Differentiation of allelic frequencies at quantitative trait loci affecting locally adaptive traits. *American Naturalist* 151: 283–292.
- Latta RG (2003) Gene flow, adaptive population divergence and comparative population structure across loci. *New Phytologist* 161: 51–58.
- Le Corre V, Kremer A (2003) Genetic variability at neutral markers, quantitative trait loci and trait in a subdivided population under selection. *Genetics* 164: 1205–1219.
- Le Corre V, Kremer A (2012) The genetic differentiation at quantitative trait loci under local adaptation. *Molecular Ecology* 21: 1548–1566.
- Kremer A, Le Corre V (2011) Decoupling of differentiation between traits and their underlying genes in response to divergent selection. *Heredity* 108: 375–385.
- Pritchard JK, Pickrell JK, Coop G (2010) The genetics of human adaptation: hard sweeps, soft sweeps, and polygenic adaptation. *Current Biology* 20: R208–15.
- Kemper KE, Saxton SJ, Bolormaa S, Hayes BJ, Goddard ME (2014) Selection for complex traits leaves little or no classic signatures of selection. *BMC Genomics* 15: 1–14.
- Risch N, Merikangas K (1996) The future of genetic studies of complex human diseases. *Science* 273: 1516–1517.
- Visscher PM, Brown MA, McCarthy MI, Yang J (2012) Five Years of GWAS Discovery. *The American Journal of Human Genetics* 90: 7–24.
- Atwell S, Huang YS, Vilhj a lmsson BJ, Willems G, Horton M, et al. (2010) Genome-wide association study of 107 phenotypes in *Arabidopsis thaliana* inbred lines. *Nature* 465: 627–631.
- Fournier-Level A, Korte A, Cooper MD, Nordborg M, Schmitt J, et al. (2011) A Map of Local Adaptation in *Arabidopsis thaliana*. *Science* 334: 86–89.
- Mackay TFC, Richards S, Stone EA, Barbadilla A, Ayroles JF, et al. (2012) The *Drosophila melanogaster* Genetic Reference Panel. *Nature* 482: 173–178.
- Manolio TA, Collins FS, Cox NJ, Goldstein DB, Hindorf LA, et al. (2009) Finding the missing heritability of complex diseases. *Nature* 461: 747–753.
- Bloom JS, Ehrenreich IM, Loo WT, Lite TLV, Kruglyak L (2013) Finding the sources of missing heritability in a yeast cross. *Nature* 494: 234–237.
- Myles S, Davison D, Barrett J, Stoneking M, Timpson N (2008) Worldwide population differentiation at disease-associated SNPs. *BMC medical genomics* 1: 22.
- Casto AM, Feldman MW (2011) Genome-wide association study SNPs in the human genome diversity project populations: does selection affect unlinked SNPs with shared trait associations? *PLoS Genetics* 7: e1001266.
- Jostins L, Ripke S, Weersma RK, Duerr RH, McGovern DP, et al. (2012) Host-microbe interactions have shaped the genetic architecture of inflammatory bowel disease. *Nature* 491: 119–124.
- Zhang G, Muglia IJ, Chakraborty R, Akcy JM (2013) Signatures of natural selection on genetic variants affecting complex human traits. *Applied & Translational Genomics* 2: 77–93.
- Turchin MC, Chiang CW, Palmer CD, Sankararaman S, Reich D, et al. (2012) Evidence of widespread selection on standing variation in Europe at height-associated SNPs. *Nature Genetics* 44: 1015–1019.
- Fraser HB (2013) Gene expression drives local adaptation in humans. *Genome Research* 23: 1089–1096.
- Corona E, Chen R, Sikora M, Morgan AA, Patel CJ, et al. (2013) Analysis of the Genetic Basis of Disease in the Context of Worldwide Human Relationships and Migration. *PLoS Genetics* 9: e1003447.
- Fisher R (1930) *The Genetical Theory of Natural Selection*. Clarendon Press.
- Falconer D (1960) *Introduction to Quantitative Genetics*. Ronald Press.
- Prout T, Barker F (1993) F Statistics in *Drosophila buzzatii*: Selection, Population Size and Inbreed. *Genetics* 375: 369–375.
- Spitze K (1993) Population structure in *Daphnia obtusa*: quantitative genetic and allozymic variation. *Genetics* 135: 367–374.
- Lewontin RC, Krakauer J (1973) Distribution of gene frequency as a test of the theory of the selective neutrality of polymorphisms. *Genetics* 74: 175–195.
- Nei M, Maruyama T (1975) Lewontin-Krakauer test for neutral genes. *Genetics* 80: 395.
- Robertson A (1975) GENE FREQUENCY DISTRIBUTIONS AS A TEST OF SELECTIVE NEUTRALITY. *Genetics*.

38. Bonhomme M, Chevalet C, Servin B, Boitard S, Abdallah JM, et al. (2010) Detecting Selection in Population Trees: The Lewontin and Krakauer Test Extended. *Genetics*.
39. Günther T, Coop G (2013) Robust Identification of Local Adaptation from Allele Frequencies. *Genetics*.
40. Ovaskainen O, Karhunen M, Zheng C, Arias JMC, Merilä J (2011) A New Method to Uncover Signatures of Divergent and Stabilizing Selection in Quantitative Traits. *Genetics* 189: 621–632.
41. Karhunen M, Ovaskainen O, Herczeg G, MERILÄ J (2014) bringing habitat information into statistical tests of local adaptation in quantitative traits: a case study of nine-spined sticklebacks. *Evolution* 68: 559–568
42. Nicholson G, Smith AV, Jonsson F, Gustafsson O, Stefansson K, et al. (2002) Assessing population differentiation and isolation from single-nucleotide polymorphism data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64: 695–715.
43. WEIR BS, HILL WG (2002) Estimating F-statistics. *Annual review of genetics* 36: 721–750.
44. Cavalli-Storza LL, Barrai I, Edwards AWF (1964) Analysis of Human Evolution Under Random Genetic Drift. *Cold Spring Harbor Symposia on Quantitative Biology* 29: 9–20.
45. Felsenstein J (1982) How can we infer geography and history from gene frequencies? *Journal of Theoretical Biology* 96: 9–20.
46. Pickrell JK, Pritchard JK (2012) Inference of Population Splits and Mixtures from Genome-Wide Allele Frequency Data. *PLoS Genetics* 8: e1002967.
47. Coop G, Witonsky D, Di Rienzo A, Pritchard JK (2010) Using environmental correlations to identify loci underlying local adaptation. *Genetics* 185: 1411–1423.
48. Fariello MI, Boitard S, Naya H, SanCristobal M, Servin B (2013) Detecting Signatures of Selection Through Haplotype Differentiation Among Hierarchically Structured Populations. *Genetics* 193: 929–941.
49. Guillot G (2012) Detection of correlation between genotypes and environmental variables. A fast computational approach for genomewide studies. *arXiv preprint arXiv:12060889*.
50. Bradburd GS, Ralph PL, Coop GM (2013) Disentangling the effects of geographic and ecological isolation on genetic differentiation. *arXiv preprint arXiv:13023274*.
51. Rao C, Toutenburg H (1999) *Linear Models: Least Squares and Alternatives*, Springer. 2nd edition, pp. 104–106.
52. Wright S (1951) The Genetical Structure of Populations. *Annals of Eugenics* 15: 323–354.
53. Lande R (1992) Neutral theory of quantitative genetic variance in an island model with local extinction and colonization. *Evolution* 46: 381–389.
54. Whitlock MC (1999) Neutral additive genetic variance in a metapopulation. *Genetical Research* 74: 215–221.
55. Rogers AR, Harpending HC (1983) Population structure and quantitative characters. *Genetics* 105: 985–1002.
56. Whitlock MC (2008) Evolutionary inference from QST. *Molecular Ecology* 17: 1885–1896.
57. Whitlock MC, Guillaume F (2009) Testing for Spatially Divergent Selection: Comparing QST to FST. *Genetics* 183: 1055–1063.
58. Li JZ, Absher DM, Tang H, Southwick AM, Casto AM, et al. (2008) Worldwide human relationships inferred from genome-wide patterns of variation. *Science (New York, NY)* 319: 1100–1104.
59. Pickrell JK, Coop G, Novembre J, Kudaravalli S, Li JZ, et al. (2009) Signals of recent positive selection in a worldwide sample of human populations. *Genome Research* 19: 826–837.
60. Charlesworth B, Nordborg M, Charlesworth D (1997) The effects of local selection, balanced polymorphism and background selection on equilibrium patterns of genetic diversity in subdivided populations. *Genetics Research* 70: 155–174.
61. McVicker G, Gordon D, Davis C, Green P (2009) Widespread Genomic Signatures of Natural Selection in Hominid Evolution. *PLoS Genetics* 5: e1000471.
62. Rosenberg NA, Pritchard JK, Weber JL, Cann HM, Kidd KK, et al. (2002) Genetic structure of human populations. *Science (New York, NY)* 298: 2381–2385.
63. Hancock AM, Witonsky DB, Gordon AS, Eshel G, Pritchard JK, et al. (2008) Adaptations to Climate in Candidate Genes for Common Metabolic Disorders. *PLoS Genetics* 4: e32.
64. Lango Allen H, Estrada K, Lettre G, Berndt SI, Weedon MN, et al. (2010) Hundreds of variants clustered in genomic loci and biological pathways affect human height. *Nature* 467: 832–838.
65. Zaitlen N, Kraft P, Patterson N, Pasaniuc B, Bhatia G, et al. (2013) Using Extended Genealogy to Estimate Components of Heritability for 23 Quantitative and Dichotomous Traits. *PLoS Genetics* 9: e1003520.
66. Gustafsson A, Lindfors P (2009) Latitudinal patterns in human stature and sexual stature dimorphism. *Annals of Human Biology* 36: 74–87.
67. Beleza S, Johnson Na, Candille SI, Absher DM, Coram MA, et al. (2013) Genetic architecture of skin and eye color in an african-European admixed population. *PLoS Genetics* 9: e1003572.
68. Jablonski NG, Chaplin G (2000) The evolution of human skin coloration. *Journal of Human Evolution* 39: 57–106.
69. Lao O, de Gruijter JM, van Duijn K, Navarro A, Kayser M (2007) Signatures of Positive Selection in Genes Associated with Human Skin Pigmentation as Revealed from Analyses of Single Nucleotide Polymorphisms. *Annals of Human Genetics* 71: 354–369.
70. Jablonski NG, Chaplin G (2010) Colloquium Paper: Human skin pigmentation as an adaptation to UV radiation. *Proceedings of the National Academy of Sciences of the United States of America* 107: 8962–8968.
71. Norton HL, Kittles RA, Parra E, McKeigue P, Mao X, et al. (2006) Genetic Evidence for the Convergent Evolution of Light Skin in Europeans and East Asians. *Molecular Biology and Evolution* 24: 710–722.
72. Miller CT, Beleza S, Pollen AA, Schluter D, Kittles RA, et al. (2007) cis-Regulatory Changes in Kit Ligand Expression and Parallel Evolution of Pigmentation in Sticklebacks and Humans. *Cell* 131: 1179–1189.
73. Edwards M, Bigham A, Tan J, Li S, Gozdzik A, et al. (2010) Association of the OCA2 Polymorphism His615Arg with Melanin Content in East Asian Populations: Further Evidence of Convergent Evolution of Skin Pigmentation. *PLoS Genetics* 6: e1000867.
74. Speliotes EK, Willer CJ, Berndt SI, Monda KL, Thorleifsson G, et al. (2010) Association analyses of 249,796 individuals reveal 18 new loci associated with body mass index. *Nature Genetics* 42: 937–948.
75. Morris AP, Voight BF, Teslovich TM, Ferreira T, Segrè AV, et al. (2012) Large-scale association analysis provides insights into the genetic architecture and pathophysiology of type 2 diabetes. *Nature Publishing Group* 44: 981–990.
76. Helgason A, Pálsson S, Thorleifsson G, Grant SFA, Emilsson V, et al. (2007) Refining the impact of TCF7L2 gene variants on type 2 diabetes and adaptive evolution. *Nature Genetics* 39: 218–225.
77. Hancock AM, Witonsky DB, Ehler E, Alkorta-Aranburu G, Beall C, et al. (2010) Human adaptations to diet, subsistence, and ecoregion are due to subtle shifts in allele frequency. *Proceedings of the National Academy of Sciences* 107: 8924–8930.
78. Klimentidis YC, Abrams M, Wang J, Fernandez JR, Allison DB (2011) Natural selection at genomic regions associated with obesity and type-2 diabetes: East Asians and sub-Saharan Africans exhibit high levels of differentiation at type-2 diabetes regions. *Human Genetics* 129: 407–418.
79. Neel JV (1962) Diabetes Mellitus: A “Thrifty” Genotype Rendered Detrimental by “Progress”? *American journal of human genetics* 14: 353.
80. Freedman ML, Reich D, Penney KL, McDonald GJ, Mignault AA, et al. (2004) Assessing the impact of population stratification on genetic association studies. *Nature Genetics* 36: 388–393.
81. Campbell CD, Ogburn EL, Lunetta KL, Lyon HN, Freedman ML, et al. (2005) Demonstrating stratification in a European American population. *Nature Genetics* 37: 868–872.
82. Price AL, Patterson NJ, Plenge RM, Weinblatt ME, Shadick NA, et al. (2006) Principal components analysis corrects for stratification in genome-wide association studies. *Nature Genetics* 38: 904–909.
83. Kang HM, Zaitlen NA, Wade CM, Kirby A, Heckerman D, et al. (2008) Efficient Control of Population Structure in Model Organism Association Mapping. *Genetics* 178: 1709–1723.
84. Price AL, Zaitlen NA, Reich D, Patterson N (2010) New approaches to population stratification in genome-wide association studies. *Nature Publishing Group* 11: 459–463.
85. Diao L, Chen KC (2012) Local Ancestry Corrects for Population Structure in Saccharomyces cerevisiae Genome-Wide Association Studies. *Genetics* 192: 1503–1511.
86. Liu L, Zhang D, Liu H, Arendt C (2013) Robust methods for population stratification in genome wide association studies. *BMC bioinformatics* 14: 132.
87. Cho YS, Go MJ, Kim YJ, Heo JY, Oh JH, et al. (2009) A large-scale genome-wide association study of Asian populations uncovers genetic factors influencing eight quantitative traits. *Nature Genetics* 41: 527–534.
88. Cho YS, Chen CH, Hu C, Long J, Ong RTH, et al. (2011) Meta-analysis of genome-wide association studies identifies eight new loci for type 2 diabetes in east Asians. *Nature Publishing Group* 44: 67–72.
89. Voight BF, Scott LJ, Steinthorsdottir V, Morris AP, Dina C, et al. (2010) Twelve type 2 diabetes susceptibility loci identified through large-scale association analysis. *Nature Publishing Group* 42: 579–589.
90. Koener JS, Saleheen D, Sim X, Schmi J, Zhang W, et al. (2011) Genome-wide association study in individuals of South Asian ancestry identifies six new type 2 diabetes susceptibility loci. *Nature Publishing Group* 43: 984–989.
91. N'Diaye A, Chen GK, Palmer CD, Ge B, Tayo B (2011) PLOS Genetics: Identification, Replication, and Fine-Mapping of Loci Associated with Adult Height in Individuals of African Ancestry. *PLoS Genetics*.
92. Carty CL, Johnson Na, Hutter CM, Reiner AP, Peters U, et al. (2012) Genome-wide association study of body height in African Americans: the Women’s Health Initiative SNP Body Height Association Resource (SHARc). *Human Molecular Genetics* 21: 711–720.
93. Monda KL, Chen GK, Taylor KC, Palmer C, Edwards TL, et al. (2013) A meta-analysis identifies new loci associated with body mass index in individuals of African ancestry. *Nature Genetics* 45: 690–696.
94. Carlson CS, Matisse TC, North KE, Haiman CA, Fesinmeyer MD, et al. (2013) Generalization and Dilution of Association Results from European GWAS in Populations of Non-European Ancestry: The PAGE Study. *PLoS Biology* 11: e1001661.
95. Marigorta UM, Navarro A (2013) High Trans-ethnic Replicability of GWAS Results Implies Common Causal Variants. *PLoS Genetics* 9: e1003566.
96. Kang HM, Sul JH, Service SK, Zaitlen NA, Kong SY, et al. (2010) technical reports. *Nature Genetics* 42: 348–354.

97. Zhou X, Stephens M (2012) technical reports. *Nature Genetics* 44: 821–824.
98. Liu X, Ong RTH, Pillai EN, Elzein AM, Small KS, et al. (2013) Detecting and Characterizing Genomic Signatures of Positive Selection in Global Populations. *American journal of human genetics*: 1–16.
99. Sabeti PC, Varilly P, Fry B, Lohmueller J, Hostetter E, et al. (2007) Genome-wide detection and characterization of positive selection in human populations. *Nature*. *Nature*.
100. Williamson SH, Hubisz MJ, Clark AG, Payseur BA, Bustamante CD, et al. (2007) Localizing Recent Adaptive Evolution in the Human Genome. *PLoS Genetics* 3: e90.
101. Sturm Ra (2009) Molecular genetics of human pigmentation diversity. *Human Molecular Genetics* 18: R9–17.
102. Franco M, Bilal U, Ordunez P, Benet M, Morejon A, et al. (2013) Population-wide weight loss and regain in relation to diabetes burden and cardiovascular mortality in Cuba 1980–2010: repeated cross sectional surveys and ecological comparison of secular trends. *BMJ* 346: f1515–f1515.
103. Kawecki TJ, Ebert D (2004) Conceptual issues in local adaptation. *Ecology Letters* 7: 1225–1241.
104. Wade MJ (2002) A gene's eye view of epistasis, selection and speciation. *Journal of Evolutionary Biology* 15: 337–346.
105. Conover DO, Schultz ET (1995) Phenotypic similarity and the evolutionary significance of countergradient variation. *Trends in Ecology & Evolution* 10: 248–252.
106. Zhang XS, HILL WG (2008) The Anomalous Effects of Biased Mutation Revisited: Mean-Optimum Deviation and Apparent Directional Selection Under Stabilizing Selection. *Genetics* 179: 1135–1141.
107. Meuwissen TH, Hayes BJ, Goddard ME (2001) Prediction of total genetic value using genome-wide dense marker maps. *Genetics* 157: 1819–1829.
108. Hayes BJ, Bowman PJ, Chamberlain AJ, Goddard ME (2009) Invited review. *Journal of Dairy Science* 92: 433–443.
109. Meuwissen T, Hayes B, Goddard M (2013) Accelerating Improvement of Livestock with Genomic Selection. *Annual Review of Animal Biosciences* 1: 221–237.
110. Zhou X, Carbonetto P, Stephens M (2013) Polygenic modeling with bayesian sparse linear mixed models. *PLoS Genetics* 9: e1003264.
111. Yang J, Benyamin B, McEvoy BP, Gordon S, Henders AK, et al. (2010) Common SNPs explain a large proportion of the heritability for human height. *Nature Genetics* 42: 565–569.
112. Davies G, Tenesa A, Payton A, Yang J, Harris SE, et al. (2011) Genome-wide association studies establish that human intelligence is highly heritable and polygenic. *Molecular Psychiatry* 16: 996–1005.
113. Yang J, Manolio TA, Pasquale LR, Boerwinkle E, Caporaso N, et al. (2011) ng.823. *Nature Publishing Group* 43: 519–525.
114. Lee SH, DeCandia TR, Ripke S, Yang J, Sullivan PF, et al. (2012) Estimating the proportion of variation in susceptibility to schizophrenia captured by common SNPs. *Nature Genetics* 44: 247–250.
115. de los Campos G, Gianola D, Allison DB (2010) *PersPectives*. *Nature Publishing Group* 11: 880–886.
116. de los Campos G, Klimentidis YC, Vazquez AI, Allison DB (2012) Prediction of Expected Years of Life Using Whole-Genome Markers. *PLoS ONE* 7: e40964.
117. de los Campos G, Vazquez AI, Fernando R, Klimentidis YC, Sorensen D (2013) Prediction of Complex Human Traits Using the Genomic Best Linear Unbiased Predictor. *PLoS Genetics* 9: e1003608.
118. Karhunen M, Ovaskainen O (2012) Estimating Population-Level Coancestry Coefficients by an Admixture F Model. *Genetics* 192: 609–617.
119. Global Lipids Genetics Consortium, Willer CJ, Schmidt EM, Sengupta S, Peloso GM, et al. (2013) Discovery and refinement of loci associated with lipid levels. *Nature Genetics* 45: 1274–1283.
120. Kremer A, Zanetto A, Ducouso A (1997) Multilocus and multitrait measures of differentiation for gene markers and phenotypic traits. *Genetics* 145: 1229–1241.
121. Blows MW (2007) A tale of two matrices: multivariate approaches in evolutionary biology. *Journal of Evolutionary Biology* 20: 1–8.
122. Chenoweth SF, Blows MW (2008) Q(St) meets the G matrix: the dimensionality of adaptive divergence in multiple correlated quantitative traits. *Evolution; international journal of organic evolution* 62: 1437–1449.
123. Falconer DS (1952) The problem of environment and selection. *American Naturalist*: 293–298.
124. Conrad DF, Jakobsson M, Coop G, Wen X, Wall JD, et al. (2006) A worldwide survey of haplotype variation and linkage disequilibrium in the human genome. *Nature Genetics* 38: 1251–1260.
125. Scheet P, Stephens M (2006) A fast and flexible statistical model for large-scale population genotype data: applications to inferring missing genotypes and haplotypic phase. *American journal of human genetics* 78: 629–644.
126. Reich D, Patterson N (2005) Will admixture mapping work to find disease genes? *Philosophical transactions of the Royal Society of London Series B, Biological sciences* 360: 1605–1607.
127. Gustafsson A, Lindfors P (2004) Human size evolution: no evolutionary allometric relationship between male and female stature. *Journal of Human Evolution* 47: 253–266.
128. Parra EJ, Kittles RA, Shriver MD (2004) Implications of correlations between skin color and genetic ancestry for biomedical research. *Nature Genetics* 36: S54–S60.