

2015-12-30, Article, Discoveries

Title: The role of deleterious substitutions in crop genomes

Thomas J. Y. Kono¹, Fengli Fu¹, Mohsen Mohammadi^{1,4}, Paul J. Hoffman¹, Chaochih Liu¹, Robert M. Stupar¹, Kevin P. Smith¹, Peter Tiffin², Justin C. Fay³, and Peter L. Morrell¹

1: Department of Agronomy and Plant Genetics, University of Minnesota, 411 Borlaug Hall, 1991 Upper Buford Circle, Saint Paul, MN 55108

2: Department of Plant Biology, University of Minnesota, 140 Gortner Laboratory, 1479 Gortner Avenue, Saint Paul, MN 55108

3: Department of Genetics, Washington University, 5526 Center for Genome Sciences, 4444 Forest Park Parkway, Saint Louis, MO 63108

4: Department of Agronomy, Purdue University, 915 West State Street, West Lafayette, IN 47907

Research was performed in the Department of Agronomy and Plant Genetics at the University of Minnesota.

Corresponding Author:

Peter L. Morrell

411 Borlaug Hall

1991 Upper Buford Circle

Saint Paul, MN 55108

USA

Fax: 612-625-1268

E-mail: pmorrell@umn.edu

Abstract

Populations experience a continual input of new mutations with fitness effects ranging from lethal to adaptive. While the distribution of fitness effects (DFE) of new mutations is not directly observable, many mutations likely have either no effect on organismal fitness or are deleterious. Historically, it has been hypothesized that a population may carry many mildly deleterious variants as segregating variation, which reduces the mean absolute fitness of the population. Recent advances in sequencing technology and sequence conservation-based metrics for predicting the functional effect of a variant permit examination of the persistence of deleterious variants in populations. The issue of segregating deleterious variation is particularly important for crop improvement, because the demographic history of domestication and breeding allows deleterious variants to persist and reach moderate frequency, potentially reducing crop productivity. In this study, we use exome resequencing of thirteen cultivated barley lines and genome resequencing of seven cultivated soybean lines to investigate the prevalence and genomic distribution of deleterious SNPs in the protein-coding regions of the genomes of two crops. We find that putatively deleterious SNPs are best identified with multiple prediction approaches, and that SNPs that cause protein truncation make up a minority of all putatively deleterious SNPs. We also report the implementation of a SNP annotation tool (BAD_Mutations) that makes use of a likelihood ratio test based on alignment of all currently publicly available Angiosperm genomes.

Introduction

Mutation produces a constant influx of new variants into populations. Each mutation has a fitness effect that varies from lethal to neutral to advantageous. While the distribution of fitness effects of new mutations is not directly observable (Eyre-Walker and Keightley 2007), most mutations with fitness impacts are deleterious (Keightley and Lynch 2003). Deleterious mutations are typically identified as changes at phylogenetically-conserved sites (Doniger et al. 2008), or loss of protein function (Yampolsky et al. 2005). Strongly deleterious variants (particularly those with dominant effects) are quickly purged from populations by purifying selection. Likewise, strongly advantageous variants increase in frequency, and ultimately fix due to positive selection (Robertson 1960; Smith and Haigh 1974). Weakly deleterious variants have the potential to persist in populations and cumulatively contribute significantly to reductions in fitness (Fay et al. 2001; Eyre-Walker et al. 2006; Doniger et al. 2008).

Considering a single variant in a population, three parameters affect its segregation: the effective population size (N_e), the selective coefficient against homozygous individuals (s), and the dominance coefficient (h). The effects of N_e and s are relatively simple; variants are primarily subject to genetic drift rather than selection if $N_e s < 1$ (Kimura et al. 1963). The effect of h is not as straightforward, as it depends on the frequency of outcrossing. In populations with a high degree of inbreeding, many individuals will be homozygous, which reduces the importance of h in determining the efficacy of selection against the variant. In populations that are outcrossing, an individual deleterious variant will occur primarily in the heterozygous state, and h will determine how “visible” the variant is to selection, with higher values of h increasing the strength of selection (Charlesworth and Charlesworth 1999). A completely recessive deleterious variant may remain effectively neutral as long as the frequency of the variant is low enough that substantial numbers of homozygous individuals are not produced. Conversely, a completely dominant deleterious variant will be quickly purged from the population (Lande and Schemske 1985). On average, deleterious variants segregating in a population are predicted to be partially recessive (Simmons and Crow 1977), allowing them to remain “hidden” from the action of purifying selection,

and reach moderate frequencies. Indeed data from a gene knockout library in yeast (Shoemaker et al. 1996) indicate that protein loss-of-function variants have an average dominance coefficient of 0.2 (Agrawal and Whitlock 2012).

Effective rates of recombination also have important impacts on the number and distribution of deleterious mutations in the genome. Low recombination regions are prone to the irreversible accumulation of deleterious variants. This phenomenon is known as the “ratchet effect” (Muller 1964). In finite populations with low recombination, the continual input of deleterious mutations and stochastic variation in reproduction causes the loss of individuals with the fewest deleterious variants. Lack of recombination precludes the selective elimination of chromosomal segments carrying deleterious variants, and thus they can increase in an inexorable fashion (Muller 1964). Nordborg (2000) demonstrates that under high levels of inbreeding, effective recombination can be decreased by almost 20-fold relative to an outbreeding population. While inbreeding populations are especially susceptible to ratchet effects on a genome-wide scale, even outbreeding species have genomic regions with limited effective recombination (Arnheim et al. 2003; McMullen et al. 2009). Both simulation studies (Felsenstein 1974) and empirical investigations in *Drosophila melanogaster* (Campos et al. 2012, 2014) indicate that deleterious variants accumulate in regions of limited recombination.

Efforts to identify individual deleterious variants and quantify them in individuals have led to a new branch of genomics research. In humans, examination of the contribution of rare deleterious variants to heritable disease has contributed to the emergence of personalized genomics as a field of study (reviewed in Abecasis et al. 2010; Cooper et al. 2010; Marth et al. 2011). Current estimates suggest that an average human may carry ~300 loss-of-function variants (Abecasis et al. 2010; Agrawal and Whitlock 2012). Individual humans carry approximately three lethal equivalents (mutations that would be lethal in the homozygous state) (Gao et al. 2015; Henn et al. 2015), and up to tens of thousands of weakly deleterious variants in coding and functional noncoding regions of the genome (Arbiza et al. 2013). These variants are enriched for mutations that are causative for diseases (Kryukov et al. 2007; Marth et al. 2011). As such they are expected

to have appreciable negative selection coefficients ($N_e s$) and be kept at low frequencies due to the action of purifying selection.

Humans are not unique in harboring substantial numbers of deleterious variants. It is estimated that almost 40% of nonsynonymous variants in *Saccharomyces cerevisiae* have deleterious effects (Doniger et al. 2008) and 20% of nonsynonymous variants in rice (Lu et al. 2006), *Arabidopsis thaliana* (Günther and Schmid 2010), and maize (Mezmouk and Ross-Ibarra 2014) are deleterious. In dogs, Cruz et al. (2008) identified an excess of nonsynonymous SNPs segregating in domesticated dogs with respect to grey wolves. A similar pattern has been found in horses (Schubert et al. 2014), suggesting that an increased prevalence of deleterious variants may be a “cost of domestication.”

Genetic bottlenecks associated with domestication (Eyre-Walker et al. 1998) may allow deleterious variants to drift to higher frequency (Robertson 1960). The selective sweeps associated with domestication and improvement (Wright et al. 2005) would decrease nucleotide diversity in affected genomic regions (Smith and Haigh 1974; Kaplan et al. 1989), and subsequently reduce the effective recombination rate. The selective and demographic processes of domestication and improvement lead to three basic hypotheses about the distribution of deleterious variants in crop plants: *i*) the relative proportion of deleterious variants will be higher in domesticates than wild relatives; *ii*) deleterious variants will be enriched near loci of agronomic importance subjected to strong selection during domestication and improvement; *iii*) the relative proportion of deleterious variants will be lower in elite cultivars than landraces due to strong selection for yield (Gaut et al. 2015).

Approaches to identify deleterious mutations take one of two forms. Quantitative genetic methods have been proposed that make use of phenotypic measurements to investigate the aggregate impact of potentially deleterious alleles (Kelly 1999). These approaches require phenotypic measurements of pedigreed individuals to estimate the net effect of potentially deleterious alleles on trait variation. While quantitative genetic approaches allow researchers to estimate the contribution of deleterious alleles to additive genetic variance to a particular trait, they do not

yield information about any individual genetic variant. Bioinformatic approaches, on the other hand, make use of measures of sequence conservation to identify variants with the greatest probability of being deleterious. When combined with genome-scale resequencing, they permit the identification of large numbers of putatively deleterious variants. Commonly applied approaches include SIFT (Sorting Intolerant From Tolerated) (Ng 2003), PolyPhen2 (Polymorphism Phenotyping) (Adzhubei et al. 2010), and a likelihood ratio test (LRT) (Chun and Fay 2009). These sequence conservation approaches operate in the absence of phenotypic data, but allow assessment of individual sequence variants. Recent advances in resequencing and sequence conservation methods have led to the suggestion that removal of deleterious variants from breeding populations presents a novel path for crop improvement (Morrell et al. 2011).

In this study, we investigate the distribution of deleterious variants in thirteen elite barley (*Hordeum vulgare* ssp. *vulgare*) and seven elite soybean (*Glycine max*) cultivars using exome and whole genome resequencing. We seek to answer four questions about the presence of deleterious variants: *i*) How many deleterious variants do individual cultivars harbor, and what proportion of these are nonsense (early stop codons) versus nonsynonymous (missense) variants? *ii*) What proportion of nonsynonymous variation is inferred to be deleterious? *iii*) How many known phenotype-altering SNPs are inferred to be deleterious? *iv*) How does the relative frequency of deleterious variants vary with recombination rate? We identify an average of ~1,000 deleterious variants per accession in our barley sample and ~700 deleterious variants per accession in our soybean sample. Approximately 40% of the deleterious variants are private to one individual in both species, suggesting the potential for selection for individuals with a reduced number of deleterious variants. Approximately 3-6% of nonsynonymous variants are inferred to be deleterious by all three approaches, and known causative SNPs annotate as deleterious at a much higher proportion than the genomic average. In soybean, where appropriate recombination rates are available, the proportion of deleterious variants is negatively correlated with recombination rate.

Results

Variant Calling and Identification of Deleterious SNPs

Resequencing and read mapping followed by read de-duplication resulted in an average coverage of ~39X exome coverage for our barley samples and ~38X genome coverage in soybean. Average heterozygosity was 2.5% in our barley sample, and 0% in our soybean sample, reflecting the inbreeding of our accessions. Heterozygous variant calls in soybean were all in reads with low mapping score, possibly due to the highly duplicated nature of the soybean genome (Schmutz et al. 2010). After realignment and variant recalibration, we identified 652,797 SNPs in thirteen cultivated and two wild barley lines. The majority of these SNPs were noncoding, with 522,863 occurring outside of CDS annotations. Of the coding SNPs, 70,069 were synonymous, and 59,865 were nonsynonymous. The list of differences from reference carried by each barley sample is summarized in Table 1, and a per-approach summary of deleterious variants is given in Table 2. SIFT identified 13,626 SNPs as deleterious, PolyPhen2 identified 13,534 SNPs to be deleterious, and the LRT called 17,865 deleterious. The intersection of all three methods gives a much smaller set of deleterious variants, with a total of 4,872 nonsynonymous SNPs identified as deleterious. While individual methods identified ~18% of nonsynonymous variants as deleterious, the intersect of methods identifies 5.7%. A derived site frequency spectrum (SFS) of synonymous, nonsynonymous, and putatively deleterious SNPs in our barley sample is shown in Figure 1A.

In soybean, we called 586,102 SNPs in gene regions. Of these, 542,558 occur in the flanking regions of a gene model. We identify 73,577 SNPs with a synonymous consequence, and 99,685 with a nonsynonymous consequence (Table 3). SNPs in the various classes sum to greater than the total number of SNPs as a single SNP in multiple transcripts can have multiple functional annotations. For instance, a SNP may be intronic in one transcript, but be in an exon of a different one. SIFT identified 7,694 of the nonsynonymous SNPs as deleterious, PolyPhen2 identified 14,933 as deleterious, and the LRT identified 11,223 as deleterious. Similarly to the barley sample, the proportion of putatively deleterious variants was similar across prediction approaches, with the exception of SIFT, which failed to find alignments for many genes. The overlap of prediction ap-

proaches identified 3,041 (2.6%) of nonsynonymous variants to be deleterious (Table 4). Derived allele frequency distributions are shown in Figure 1B. Variants inferred to be deleterious are generally at lower derived allele frequency than other classes of variation, implying that these variants are truly deleterious.

Nonsense variants made up a relatively small proportion of putatively deleterious variants. In our barley sample, we identify a total of 711 nonsense variants, 14.5% of our putatively deleterious variants. In soybean, we identify 1,081 nonsense variants, which make up 15.7% of putatively deleterious variants. Nonsense variants have a higher heterozygosity than tolerated, silent, or deleterious missense variants (Figure S2). While the absolute differences in heterozygosity were small due to the inbred nature of our samples, the pattern suggests that nonsense variants are more strongly deleterious than just missense variants.

Deleterious Mutations and Causative Variants

Bioinformatic approaches to identifying deleterious variants rely on sequence constraint to estimate protein functional impact. An example of a deleterious variant showing a derived base substitution that alters a phylogenetically conserved codon is shown in Figure 2. The variants identified in these approaches should be enriched for variants that cause large phenotypic changes. To explore how frequently known causative SNPs annotate as deleterious, we compiled a list of 23 nonsynonymous variants inferred to contribute to known phenotypic variation in barley and 11 in soybean and tested the effect of these variants in our prediction pipeline. Because they mostly alter agronomically important phenotypes, they are likely to be truly deleterious in natural populations. Of 23 putative causative mutations in barley, 6 (25%) of them were inferred to be deleterious (Table S5). Of the 11 soybean putatively causative mutations, 5 (45%) of them were inferred to be deleterious. This contrasts with the genome-wide average of ~3-6%, showing that variants that annotate as deleterious are more likely to impact phenotypes.

Deleterious Mutations and Genetic Map Distance

The purging of deleterious variants from populations is greatly affected by the effective recombination rate, which is related to the ratio of genetic distance to physical distance. To examine the relationship between the number of deleterious variants and recombination rate, we used a high-density genetic map from a soybean recombinant inbred line family (Lee et al. 2015). The soybean map was based on a subset of the SoySNP50K genotyping platform (Song et al. 2013). There was a weak but significant correlation between recombination rate and the proportion of nonsynonymous SNPs inferred to be deleterious ($r^2 = 0.007$, $p < 0.001$, Figures 3, S3). We did not examine this relationship in barley because the barley reference genome assembly (Mayer et al. 2012) contains limited physical distance information.

Discussion

Questions regarding the prevalence of deleterious variants date back over half a century (Fisher 1930; Muller 1950). In finite populations, the segregation of deleterious variants can have a substantial impact on population mean fitness (Kimura et al. 1963). While it has been argued that the concept of a reduction of fitness relative to a hypothetical optimal genotype is irrelevant (Wallace 1970), mutation accumulation studies have shown that the accumulation of deleterious mutations has a significant effect on absolute fitness (Schultz et al. 1999; Shaw et al. 2002).

Our results demonstrate that a large number of putatively deleterious variants persist in individual cultivars in both barley and soybeans. The approaches used in this study predict the probability that a given amino acid or nucleotide substitution disrupts protein function. Mutations that alter phenotypes may be especially likely to annotate as deleterious, and we show that a high proportion of inferred causative mutations annotate as deleterious. It should be noted that variants identified as deleterious may affect a phenotype that is adaptive in only part of the species range or has a transient selective advantage – i.e., locally or temporally adaptive phenotypes. If the portion of the range in which the phenotype is adaptive is small or the selective advantage is transient, such variants will be kept at low frequencies and also be identified as deleterious. Just as few variants are expected to be globally advantageous, a portion of deleterious variants are likely to

not be globally disadvantageous. Such variants could be either locally or temporally advantageous, with a fitness advantage under some circumstances contributing to their maintenance in populations (Tiffin and Ross-Ibarra 2014).

At the molecular level, variants occurring in minor transcripts of genes may exhibit conditional neutrality (Tiffin and Ross-Ibarra 2014), and $N_e s$ will be too low for purifying selection to act. Gan et al. (2011) identified many isoforms of genes among a diverse panel of *Arabidopsis thaliana* accessions, as well as compensatory mutations for a majority of frameshift mutations. Genetic variants that annotated as nonsynonymous or nonsense using the *A. thaliana* reference are frequently spliced out of the transcript such that the gene still produces a full-length and functional product. In a similar vein, deleterious variants are often accompanied by multiple compensatory mutations that alleviate their fitness effects (Poon and Otto 2000; Poon and Chao 2005). The occurrence of the bulk of putatively deleterious variants in the rarest frequency classes (Fig 1), and a higher level of observed heterozygosity for putatively deleterious variants (Figure S5) are both consistent with action of purifying selection on variants with negative impacts on fitness. Putatively disease-causing variants in human populations have also been observed to occur at low frequencies and to occur over a more geographically restricted range (Marth et al. 2011).

Comparison of Identification Methods

Each of the methods used here to identify deleterious variants makes use of sequence constraint across a phylogenetic relationship. They differ in terms of the models used to assess the functional effect of a variant. SIFT uses a heuristic, which determines if a nonsynonymous variant alters a conserved site based on an alignment build from PSI-BLAST results (Ng 2003). Polyphen2 is similar but, additionally identifies potential disruptions in secondary or tertiary structure of the encoded protein (when this information is available) (Adzhubei et al. 2010). Both of these approaches estimate codon conservation from a multiple sequence alignment, but do not use phylogenetic relationships in their predictions. PolyPhen2 identified the largest number of variants as deleterious. The reason for this may be that the data used to train the PolyPhen2 model is from human disease-causing and neutral variants. Nonhuman systems may differ fundamental-

ly as to which amino acid substitutions tend to have strong functional impact, which would reduce prediction accuracy in other species (Adzhubei et al. 2010). The LRT explicitly calculates the local synonymous substitution rate, and uses that to test whether an individual codon is under selective constraint or evolving neutrally (Chun and Fay 2009). It is a hypothesis-driven approach, and compares the likelihood of two evolutionary scenarios. Variants in selectively constrained codons are considered to be deleterious.

The SNPs predicted to be deleterious differ somewhat between prediction approaches. Even though SIFT and PolyPhen2 identify similar proportions of nonsynonymous SNPs as deleterious, they overlap at ~50% of sites (Table 2). SNPs identified through at least two approaches, seem more likely to be deleterious, based on lower average derived allele frequencies (Figure S4). Comparisons of the distribution of Grantham scores (Grantham 1974) show high similarity in the severity of amino acid replacements that are predicted to be deleterious by each approach (Figure S5). The effects of reference bias are apparent in SIFT and PolyPhen2. In barley and soybean, the reference genotypes are 'Morex' and 'Williams 82' respectively. Even when polarizing by ancestral and derived alleles, these genotypes show considerably fewer inferred deleterious variants (Table 2; Table 4).

Deleterious Variants in Crop Breeding

Identification and elimination of deleterious variants has been proposed as a potential means of improving plant fitness and crop yield (Morrell et al. 2011). Current plant breeding strategies using genome-wide prediction rely on estimating genome-wide marker effects on quantitative traits of interest (Meuwissen et al. 2001). Genome-wide prediction has been shown to be effective in both animals (Schaeffer 2006) and plants (Heffner et al. 2011; Jacobson et al. 2014), but these approaches rely on estimating marker contributions to a quantitative trait (i.e., a measured phenotypic effect). The genetic architecture of these traits suggests that our ability to quantify the effects of individual loci will reach practical limits before we can identify loci contributing to their variance (Rockman 2012). Many traits of agronomic interest, particularly yield in grain crops, are quantitative and have a complex genetic basis. As such, they are under the influence of

environmental effects and many loci (Falconer and Mackay 1996). QTL mapping approaches to identifying favorable variants for agronomic traits will reach practical limits, even for variants of large effect (King et al. 2012). Current genome wide prediction and selection methodologies rely on estimating the combined effects of markers across the genome (Meuwissen et al. 2001), but this approach is limited by recombination rate and the ability to measure phenotypes of interest. The identification and purging of deleterious variants should provide a complementary approach to current breeding methodologies (Morrell et al. 2011).

Rise of Deleterious Variants Into Populations

The number of segregating deleterious variants in a species is very different from the number of de novo deleterious mutations in each generation, commonly identified as U . In humans, U is estimated at ~2 new deleterious variants per genome per generation (Agrawal and Whitlock 2012) and estimates from Arabidopsis suggest that U is approximately 0.1 (Schultz et al. 1999). U is the product of the per-base pair mutation rate, the genome size, and the fraction of the genome that is deleterious when mutated (Charlesworth 2012). Even though new mutations are constantly arising, the standing load of deleterious variation greatly exceeds the rate at which they arise (Charlesworth et al. 2004; Charlesworth 2012). However, our results show that ~40% of our inferred deleterious variants are private to individual cultivars, suggesting that they can be purged from breeding programs.

In the current study, we restricted our analyses to protein coding regions, but additional recent evidence suggests that deleterious variants can accumulate in conserved noncoding sequences, such as transcription factor binding sites (Arbiza et al. 2013). As such, analysis of the protein-coding regions of genomes presents a lower-bound on the estimates of the number of deleterious variants segregating in populations. Efforts to identify deleterious variants in noncoding sequence are limited by scant knowledge of functional constraints on noncoding genomic regions, and difficulty in aligning noncoding regions from all but the most closely related taxa (Doniger et al. 2008). Annotation of noncoding sequence will uncover additional deleterious variants, but a majority of putatively deleterious variants will be in coding regions. The several thousand puta-

tively deleterious variants we identify per individual cultivar should provide ample targets for selection of recombinant progeny in a breeding program.

Materials and Methods

Plant Material and DNA Sequencing

The exome resequencing data reported here includes thirteen cultivated barleys, and two wild barley accessions. Barley exome capture was based on a 60 Mb liquid-phase Nimblegen capture design (Mascher et al. 2013). For the soybean sample, we resequenced whole genomes of seven elite soybean cultivars and used previously-generated whole genome sequence of *Glycine soja* (Kim et al. 2010). Each sample was prepared and sequenced with manufacturer protocols (Illumina, San Diego, CA) to at least 25x coverage of the target with 76bp, 100bp or 151bp paired-end reads. A summary of samples and sequencing statistics is given in Table S1.

Read Mapping and SNP Calling

DNA sequence handling followed the “Genome Analysis Tool Kit (GATK) Best Practices” workflow from the Broad Institute (broadinstitute.org/gatk/guide/topic?name=best-practices). Our workflow for read mapping and SNP calling is depicted in Figure S1. First, reads were checked for proper length, Phred score distribution, and *k*-mer contamination with FastQC (bioinformatics.babraham.ac.uk/projects/fastqc/). Primer and adapter sequence contamination was then trimmed from barley reads using Scythe (github.com/vsbuffalo/scythe), using a prior on contamination rate of 0.05. Low-quality bases were then removed with Sickle (github.com/najoshi/sickle), with a minimum average window Phred quality of 25, and window size of 10% of the read length. Soybean reads were trimmed using the fastqc-mcf tool in the ea-utils package (code.google.com/p/ea-utils/). Post-alignment processing and SNP calling were performed with the GATK v. 3.1 (McKenna et al. 2010; DePristo et al. 2011).

Barley reads were aligned to the Morex draft genome sequence (Mayer et al. 2012) using

BWA-MEM (Li and Durbin 2009). We tuned the alignment reporting parameter and the gapping parameters to allow ~2% mismatch between the reads and reference sequence, which is roughly equivalent to the highest estimated nucleotide diversity observed at a locus in barley coding sequence (Morrell et al. 2003, 2006, 2014). The resulting SAM file was trimmed of unmapped reads with Samtools (Li et al. 2009), sorted, and trimmed of duplicate reads with Picard tools (picard.sourceforge.net/). We then realigned around indels, using a set of 100 previously known indels from Sanger resequencing of 25 loci (Caldwell et al. 2006; Morrell and Clegg 2007; Morrell et al. 2014). Sequence coverage was estimated with 'bedtools genomecov,' using the regions included in the Nimblegen barley exome capture design (https://sftp.rch.cm/diagnostics/sequencing/nimblegen_annotations/ez_barley_exome/barley_exome.zip). Individual sample alignments were then merged into a multisample alignment for variant calling. A preliminary set of variants was called with the GATK HaplotypeCaller with a heterozygosity (average pairwise diversity) value of 0.008, based on average coding sequence diversity reported for cultivated barley (Morrell et al. 2014). This preliminary set of variants was filtered to sites with a genotype score of 40 or greater, heterozygous calls in at most two individuals, and read depth of at least five reads. We then used the filtered variants, SNPs identified in the Sanger resequencing data set, and 9,605 SNPs from genotyping assays: 5,010 from the James Hutton Institute (Comadran et al. 2012), and 4,595 from Illumina GoldenGate assays (Close et al. 2009) as input for the GATK VariantRecalibrator to obtain a final set of variant calls.

Processing of soybean samples is as described above, but with the following modifications. Soy reads were aligned to the Williams 82 reference genome sequence (Schmutz et al. 2010). Mismatch and reporting parameters for the cultivated samples were adjusted to allow for ~1% mismatch between reads and reference, which is approximately the highest typical genic sequence diversity in soybean cultivars (Hyten et al. 2006). The alignments were trimmed and sorted as described above. Preliminary variants were called as in the barley sample, but with a heterozygosity value of 0.001, which is the nucleotide diversity reported by Hyten et al. (2006). Final variant calls were obtained in the same way as described for the barley sample, using SNPs on the SoySNP50K chip (Song et al. 2013) as known variants.

Read mapping scripts, variant calling scripts, and variant filtering scripts for both barley and soybean are available on GitHub at (github.com/MorrellLAB/Deleterious_Mutations).

SNP Classification

Barley SNPs were identified as coding or noncoding using the Generic Feature Format v3 (GFF) file provided with the reference genome (Mayer et al. 2012). A custom Python script was then used to identify coding barley SNPs as synonymous or nonsynonymous. Soybean SNPs were assigned using primary transcripts using the Variant Effect Predictor (VEP) from Ensembl (ensembl.org/info/docs/tools/vep/index.html). Nonsynonymous SNPs were then assessed using SIFT (Ng 2003), PolyPhen2 (Adzhubei et al. 2010) using the ‘HumDiv’ model, and a likelihood ratio test comparing codon evolution under selective constraint to neutral evolution (Chun and Fay 2009). For the likelihood ratio test, we used the phylogenetic relationships between 37 Angiosperm species based on genic sequence from complete plant genome sequences available through Phytozome (phytozome.jgi.doe.gov/) and Ensembl Plants (plants.ensembl.org/). The LRT is implemented as a Python package we call ‘BAD_Mutations’ (BLAST Aligned-Deleterious Mutations; github.com/MorrellLAB/BAD_Mutations). Coding sequences from each genome were downloaded and converted into BLAST databases. The coding sequence from the query species was used to identify the best match from each species using TBLASTX. The best match from each species was then aligned using PASTA (Mirab et al. 2014), a phylogeny-aware alignment tool. The resulting alignment was then used as input to the likelihood ratio test for the affected codon. The LRT was performed on codons with a minimum of 10 species represented in the alignment at the queried codon. Reference sequences were masked from the alignment to reduce the effect of reference bias (Simons et al. 2014). A SNP was identified as deleterious if the p-value for the test was less than 0.05, with a Bonferroni correction applied based on the number of tested codons, and if either the alternate or reference allele was not seen in any of the other species. A full list of species names and genome assembly and annotation versions used is available in Table S4.

Inference of Ancestral State

Prediction of deleterious mutations is complicated by reference bias (Chun and Fay 2009; Simons et al. 2014), which manifests in two ways. First, individuals that are closely related to the reference line used for the reference genome will appear to have fewer genetic variants, and thus fewer inferred nonsynonymous and deleterious variants. Second, when the reference strain carries a derived allele at a polymorphic site, that site is generally not predicted to be deleterious (Simons et al. 2014). To address the issue of reference bias, we polarized all coding variants by ancestral and derived state, rather than reference and non-reference state. Ancestral states were inferred for SNPs in gene regions by inferring the majority state in the most closely related clade from the consensus phylogenetic tree for the species included in the LRT. For barley, the ancestral states were inferred from gene alignments of *Aegilops tauschii*, *Brachypodium distachyon*, and *Tritium urartu*. For soybean, ancestral states were inferred using *Medicago truncatula* and *Phaseolus vulgaris*. This approach precludes universal inference of ancestral state for noncoding variants. However, examination of alignments of intergenic sequence in Triticeae species and in *Glycine* species showed that alignments outside of protein coding sequence is not reliable for ancestral state inference (data not shown).

Acknowledgements

The authors thank Brandon Gaut, Michael Kantar, and Ana Poets for helpful comments on an earlier version of the manuscript. This work was supported by a USDA NIFA National Needs Fellowship (Appropriation No. 5430-21000-006-00D) and a MnDrive 2014 Food Security Fellowship, in support of TJYK. Support was also provided by the Minnesota Agricultural Experiment Station Variety Development fund and U.S. NSF Plant Genome Program (DBI-1339393). This research was carried out with hardware and software support provided by the Minnesota Supercomputing Institute (MSI) at the University of Minnesota.

Figures

Figure 1

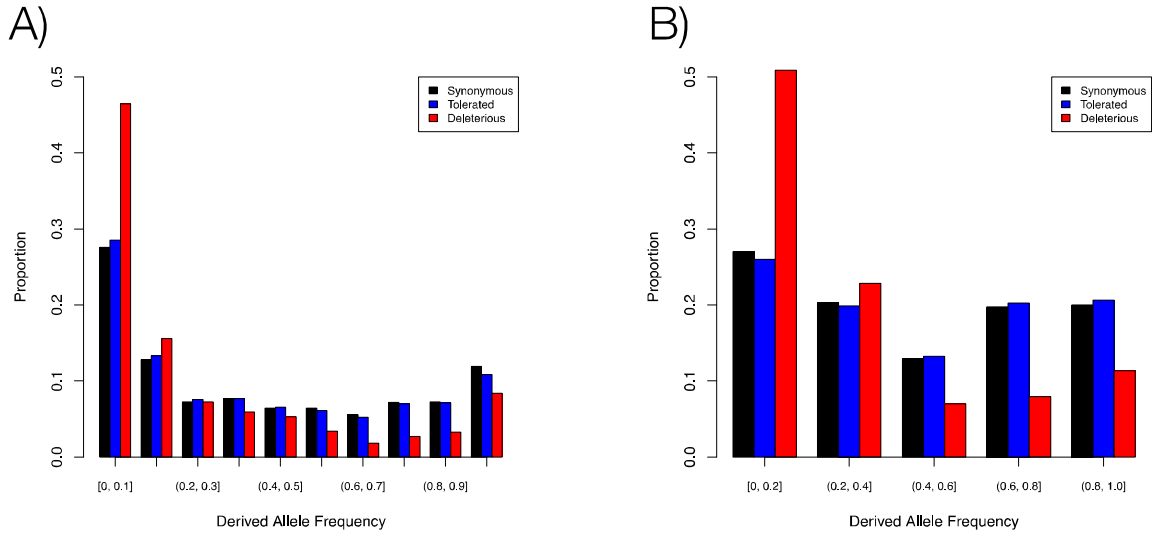


Figure 2

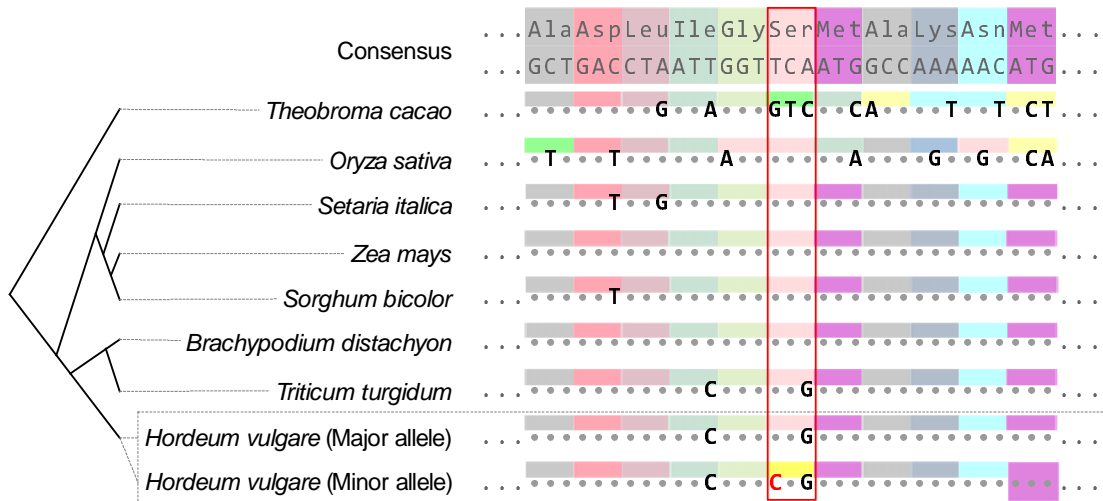



Figure 3



Tables

Table 1. Counts of variants in various classes for each barley sample.

Sample	Diff. From Morex	Noncoding	Syn.	Nonsyn.	Nonsense
Barke	162,954	130,832	17,800	14,322	86
Bonus	135,540	108,762	15,184	11,594	65
Borwina	139,222	111,613	15,699	11,910	62
Bowman	130,335	105,575	13,645	11,115	77

Foma	156,846	126,608	16,746	13,492	84
Gull	128,671	103,143	14,555	10,973	55
Harrington	153,203	124,264	15,995	12,944	89
Haruna Nijo	155,245	125,143	16,655	13,447	83
Igri	161,224	130,100	17,275	13,849	98
Kindred	86,932	69,996	9,458	7,478	34
Morex	5,239	4,602	253	384	4
Steptoe	148,785	120,245	15,839	12,701	91
Vogelsanger Gold	146,303	117,557	16,168	12,578	70
FT11	243,049	195,009	26,886	21,154	147
OUH602	197,405	158,377	21,701	17,327	115

Diff. From Morex = Differences from reference; Syn. = Synonymous; Nonsyn. = Nonsynonymous

Table 2. Counts of variants that are predicted to be deleterious by three prediction methods in barley samples. Numbers in parentheses are proportions of nonsynonymous variants in each sample.

Sample	SIFT	PPH	LRT	Intersect
Barke	3,609 (0.199)	3,507 (0.194)	3,339 (0.185)	1,070 (0.059)
Bonus	3,066 (0.189)	2,911 (0.180)	2,941 (0.182)	882 (0.054)
Borwina	3,083 (0.190)	2,932 (0.181)	2,977 (0.184)	894 (0.055)
Bowman	3,335 (0.186)	3,219 (0.180)	3,248 (0.182)	998 (0.056)
Foma	3,879 (0.200)	3,666 (0.189)	3,598 (0.186)	1,151 (0.059)
Gull	3,076 (0.187)	2,946 (0.179)	2,998 (0.183)	885 (0.054)
Harrington	3,599 (0.196)	3,536 (0.192)	3,335 (0.182)	1,063 (0.058)
Haruna Nijo	3,576 (0.194)	3,507 (0.191)	3,354 (0.182)	1,078 (0.059)
Igri	3,694 (0.199)	3,585 (0.193)	3,350 (0.181)	1,077 (0.058)
Kindred	3,131 (0.175)	2,995 (0.168)	3,204 (0.180)	931 (0.052)
Morex	2,492 (0.146)	2,298 (0.135)	3,034 (0.178)	700 (0.041)
Steptoe	3,479 (0.195)	3,393 (0.190)	3,266 (0.183)	1,061 (0.06)
Vogelsanger Gold	3,172 (0.193)	3,100 (0.188)	2,995 (0.182)	910 (0.055)
FT11	3,948 (0.225)	3,937 (0.224)	3,290 (0.187)	1,220 (0.069)
OUH602	3,864 (0.209)	3,886 (0.210)	3,391 (0.183)	1,174 (0.064)

Joint	13,626 (0.228)	13,534 (0.226)	11,574 (0.193)	4,275 (0.071)
-------	----------------	----------------	----------------	---------------

Table 3. Counts of variants in various classes for each soybean sample.

Sample	Diff. From Wm82	Noncoding	Syn.	Nonsyn.	Nonsense
Archer	71,400	38,056	12,397	16,948	457
IA3023	60,388	33,905	11,335	14,703	426
M92-220	85,861	45,183	14,644	19,843	591
Minsoy	151,622	77,466	23,614	31,430	927
Noir1	128,689	65,508	20,292	27,309	828
Williams	3,340	7,308	3,688	3,854	87
Williams 82	9,497	6,267	3,248	3,359	70
<i>G. soja</i>	163,365	83,940	24,120	32,114	930
Archer	71,400	38,056	12,397	16,948	457

Diff. From Wm82 = Differences from reference; Syn. = Synonymous; Nonsyn. = Nonsynonymous

Table 4. Counts of variants that are predicted to be deleterious by three prediction methods in soybean samples. Numbers in parentheses are proportions of nonsynonymous variants in each sample.

Sample	SIFT	PPH	LRT	Intersect
Archer	1,987 (0.062)	3,847 (0.120)	3,166 (0.099)	773 (0.024)
IA3023	1,994 (0.062)	3,837 (0.119)	3,178 (0.099)	792 (0.025)
M92-220	2,142 (0.066)	4,269 (0.132)	3,397 (0.105)	860 (0.027)
Minsoy	2,686 (0.081)	5,257 (0.158)	3,977 (0.119)	1,135 (0.034)
Noir 1	2,417 (0.073)	4,951 (0.150)	3,865 (0.117)	1,035 (0.031)
Williams	1,408 (0.048)	2,673 (0.090)	2,425 (0.082)	485 (0.016)
Williams 82	1,394 (0.047)	2,631 (0.089)	2,393 (0.081)	478 (0.016)
<i>G. soja</i>	1,751 (0.074)	3,583 (0.150)	2,675 (0.112)	716 (0.030)
Joint	7,694 (0.076)	14,933 (0.147)	11,223 (0.110)	3,041 (0.030)

Figure Legends

Figure 1: Unfolded site frequency spectrum for coding regions showing deleterious, tolerated, and synonymous SNPs for barley and soybean. Ancestral state was inferred by majority state in the LRT gene alignments. A) is based on thirteen domesticated barley accessions and two wild accessions while B) is based on seven cultivated soybean accessions and one wild accession.

Figure 2: A sample alignment used to infer a Serine to Proline mutation as deleterious in *Ppd-H1*. The alignment is built from sequences used by SIFT, and the affected codon is highlighted in red.

Figure 3: Comparison between the recombination rate and proportion of nonsynonymous SNPs inferred to be deleterious in soybean on chromosome 4. The cM/Mb values are calculated from a genetic map using the SoySNP6K (Lee *et al.* 2015). Red points are the proportion of nonsynonymous variants that are inferred to be deleterious, and black points are cM/Mb values between adjacent markers.

Figure S1: A schematic for the read mapping and SNP calling workflow. Boxes with bold borders denote the start and end points of workflow. Rounded boxes with light grey borders are the tools that are used at each step in the pipeline. Boxes with dashed borders denote external datasets that are used in various steps of the pipeline.

Figure S2: The distributions of per-SNP heterozygosity for tolerated nonsynonymous and silent SNPs, deleterious missense SNPs, and nonsense SNPs. Nonsense SNPs tend to be heterozygous more often than deleterious or tolerated SNPs.

Figure S3: Correlation between recombination rate (cM/Mb) and proportion of nonsynonymous SNPs inferred to be deleterious genome-wide in our soybean sample.

Figure S4: Unfolded site frequency spectrum for SNPs in A) barley and B) soybean predicted to

be deleterious by one, two, or three prediction approaches. SNPs predicted by only one approach are not as strongly skewed toward rare variants, suggesting that the intersection of multiple prediction approaches gives the most reliable prediction of deleterious variants.

Figure S5: Distribution of Grantham score for nonsynonymous variants predicted to be deleterious by each prediction approach. Each approach and the intersection of each approach gives a very similar distribution of Grantham scores. Vertical lines show the mean of the distribution.

References

- Abecasis GR, Altshuler D, Auton A, Brooks LD, Durbin RM, Gibbs RA, Hurles ME, McVean GA. 2010. A map of human genome variation from population-scale sequencing. *Nature*. 467:1061-1073.
- Adzhubei IA, Schmidt S, Peshkin L, Ramensky VE, Gerasimova A, Bork P, Kondrashov AS, Sunyaev SR. 2010. A method and server for predicting damaging missense mutations. *Nat Methods*. 7:248-249.
- Agrawal AF, Whitlock MC. 2012. Mutation load: the fitness of individuals in populations where deleterious alleles are abundant. *Annu Rev Ecol Evol Syst*. 43:115-135.
- Arbiza L, Gronau I, Aksoy BA, Hubisz MJ, Gulko B, Keinan A, Siepel A. 2013. Genome-wide inference of natural selection on human transcription factor binding sites. *Nat Genet*. 45:723-729.
- Arnheim N, Calabrese P, Nordborg M. 2003. Hot and cold spots of recombination in the human genome: the reason we should find them and how this can be achieved. *Am J Hum Genet*. 73:5-16.
- Caldwell KS, Russell J, Langridge P, Powell W. 2006. Extreme population-dependent linkage disequilibrium detected in an inbreeding plant species, *Hordeum vulgare*. *Genetics*. 172:557-567.
- Campos JL, Charlesworth B, Haddrill PR. 2012. Molecular evolution in nonrecombining regions of the *Drosophila melanogaster* genome. *Genome Biol Evol*. 4:278-288.

- Campos JL, Halligan DL, Haddrill PR, Charlesworth B. 2014. The relation between recombination rate and patterns of molecular evolution and variation in *Drosophila melanogaster*. *Mol Biol Evol.* 31:1010-1028.
- Charlesworth B. 2012. The effects of deleterious mutations on evolution at linked sites. *Genetics.* 190:5-22.
- Charlesworth B, Borthwick H, Bartolomé C, Pignatelli P. 2004. Estimates of the genomic mutation rate for detrimental alleles in *Drosophila melanogaster*. *Genetics.* 167:815-826.
- Charlesworth B, Charlesworth D. 1999. The genetic basis of inbreeding depression. *Genet Res.* 74:329-340.
- Chun S, Fay JC. 2009. Identification of deleterious mutations within three human genomes. *Genome Res.* 19:1553-1561.
- Close TJ, Bhat PR, Lonardi S et al. 2009. Development and implementation of high-throughput SNP genotyping in barley. *BMC Genomics.* 10:582.
- Comadran J, Kilian B, Russell J et al. 2012. Natural variation in a homolog of *Antirrhinum* CENTRORADIALIS contributed to spring growth habit and environmental adaptation in cultivated barley. *Nat Genet.* 44:1388-1392.
- Cooper DN, Chen JM, Ball EV, Howells K, Mort M, Phillips AD, Chuzhanova N, Krawczak M, Kehrer-Sawatzki H, Stenson PD. 2010. Genes, mutations, and human inherited disease at the dawn of the age of personalized genomics. *Hum Mutat.* 31:631-655.
- Cruz F, Vilà C, Webster MT. 2008. The legacy of domestication: accumulation of deleterious mutations in the dog genome. *Mol Biol Evol.* 25:2331-2336.
- DePristo MA, Banks E, Poplin R et al. 2011. A framework for variation discovery and genotyping using next-generation DNA sequencing data. *Nat Genet.* 43:491-498.
- Doniger SW, Kim HS, Swain D, Corcuera D, Williams M, Yang SP, Fay JC. 2008. A catalog of neutral and deleterious polymorphism in yeast. *PLoS Genet.* 4:e1000183.
- Eyre-Walker A, Keightley PD. 2007. The distribution of fitness effects of new mutations. *Nat Rev Genet.* 8:610-618.

- Eyre-Walker A, Woolfit M, Phelps T. 2006. The distribution of fitness effects of new deleterious amino acid mutations in humans. *Genetics*. 173:891-900.
- Eyre-Walker A, Gaut RL, Hilton H, Feldman DL, Gaut BS. 1998. Investigation of the bottleneck leading to the domestication of maize. *Proc Natl Acad Sci U S A*. 95:4441-4446.
- Falconer DS, Mackay TFC. 1996. Introduction to quantitative genetics. London: Longman.
- Fay JC, Wyckoff GJ, Wu C-I. 2001. Positive and negative selection on the human genome. *Genetics*. 158:1227-1234.
- Felsenstein J. 1974. The evolutionary advantage of recombination. *Genetics*. 78:737-756.
- Fisher RA. 1930. The distribution of gene ratios for rare mutations. *Proceedings of the Royal Society of Edinburgh*. 50:205-220.
- Gan X, Stegle O, Behr J et al. 2011. Multiple reference genomes and transcriptomes for *Arabidopsis thaliana*. *Nature*. 477:419-423.
- Gao Z, Waggoner D, Stephens M, Ober C, Przeworski M. 2015. An estimate of the average number of recessive lethal mutations carried by humans. *Genetics*. 199:1243-1254.
- Gaut BS, Diez CM, PL M. In Review. Genomics and the contrasting dynamics of annual and perennial domestication. *Trends in Genetics*.
- Grantham R. 1974. Amino Acid Difference Formula to Help Explain Protein Evolution. *Science*. 185:862-864.
- Günther T, Schmid KJ. 2010. Deleterious amino acid polymorphisms in *Arabidopsis thaliana* and rice. *Theor Appl Genet*. 121:157-168.
- Heffner EL, Jannink J-L, Iwata H, Souza E, Sorrells ME. 2011. Genomic selection accuracy for grain quality traits in biparental wheat populations. *Crop Science*. 51:2597.
- Henn BM, Botigué LR, Bustamante CD, Clark AG, Gravel S. 2015. Estimating the mutation load in human genomes. *Nat Rev Genet*. 16:333-343.
- Hyten DL, Song Q, Zhu Y, Choi IY, Nelson RL, Costa JM, Specht JE, Shoemaker RC, Cregan PB. 2006. Impacts of genetic bottlenecks on soybean genome diversity. *Proc Natl Acad Sci U S A*. 103:16666-16671.

- Jacobson A, Lian L, Zhong S, Bernardo R. 2014. General Combining Ability Model for Genomewide Selection in a Biparental Cross. *Crop Science*. 54:895.
- Kaplan NL, Hudson RR, Langley CH. 1989. The "hitchhiking effect" revisited. *Genetics*. 123:887-899.
- Keightley PD, Lynch M. 2003. Toward a realistic model of mutations affecting fitness. *Evolution*. 57:683-685.
- Kelly JK. 1999. An experimental method for evaluating the contribution of deleterious mutations to quantitative trait variation. *Genet Res*. 73:263-273.
- Kim MY, Lee S, Van K et al. 2010. Whole-genome sequencing and intensive analysis of the undomesticated soybean (*Glycine soja* Sieb. and Zucc.) genome. *Proc Natl Acad Sci U S A*. 107:22032-22037.
- Kimura M, Maruyama T, Crow JF. 1963. The mutation load in small populations. *Genetics*. 48:1303.
- King EG, Merkes CM, McNeil CL, Hoofer SR, Sen S, Broman KW, Long AD, Macdonald SJ. 2012. Genetic dissection of a model complex trait using the *Drosophila* Synthetic Population Resource. *Genome Res*. 22:1558-1566.
- Kryukov GV, Pennacchio LA, Sunyaev SR. 2007. Most rare missense alleles are deleterious in humans: implications for complex disease and association studies. *Am J Hum Genet*. 80:727-739.
- Lande R, Schemske DW. 1985. The evolution of self-fertilization and inbreeding depression in plants. I. genetic models. *Evolution*. 39
- Lee S, Freewalt KR, McHale LK, Song Q, Jun T-H, Michel AP, Dorrance AE, Mian MAR. 2015. A high-resolution genetic linkage map of soybean based on 357 recombinant inbred lines genotyped with BARCSoySNP6K. *Mol Breeding*. 35
- Li H, Durbin R. 2009. Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics*. 25:1754-1760.

- Li H, Handsaker B, Wysoker A, Fennell T, Ruan J, Homer N, Marth G, Abecasis G, Durbin R, 1000 GPDPS. 2009. The Sequence Alignment/Map format and SAMtools. *Bioinformatics*. 25:2078-2079.
- Lu J, Tang T, Tang H, Huang J, Shi S, Wu CI. 2006. The accumulation of deleterious mutations in rice genomes: a hypothesis on the cost of domestication. *Trends Genet*. 22:126-131.
- Marth GT, Yu F, Indap AR et al. 2011. The functional spectrum of low-frequency coding variation. *Genome Biol*. 12:R84.
- Mascher M, Richmond TA, Gerhardt DJ et al. 2013. Barley whole exome capture: a tool for genomic research in the genus *Hordeum* and beyond. *Plant J*. 76:494-505.
- Mayer KF, Waugh R, Brown JW et al. 2012. A physical, genetic and functional sequence assembly of the barley genome. *Nature*. 491:711-716.
- McKenna A, Hanna M, Banks E et al. 2010. The Genome Analysis Toolkit: a MapReduce framework for analyzing next-generation DNA sequencing data. *Genome Res*. 20:1297-1303.
- McMullen MD, Kresovich S, Villeda HS et al. 2009. Genetic properties of the maize nested association mapping population. *Science*. 325:737-740.
- Meuwissen THE, Hayes BJ, Goddard ME. 2001. Prediction of total genetic value using genome-wide dense marker maps. *Genetics*. 157:1819-1829.
- Mezmouk S, Ross-Ibarra J. 2014. The pattern and distribution of deleterious mutations in maize. *G3*. 4:163-171.
- Mirab S, Nguyen N, Warnow T. 2014. PASTA: ultra-large multiple sequence alignment. In: Hutchison D, Kanade T, Kittler J et al., editors. Cham: Springer International Publishing. p. 177-191.
- Morrell PL, Buckler ES, Ross-Ibarra J. 2011. Crop genomics: advances and applications. *Nat Rev Genet*. 13:85-96.
- Morrell PL, Clegg MT. 2007. Genetic evidence for a second domestication of barley (*Hordeum vulgare*) east of the Fertile Crescent. *Proc Natl Acad Sci U S A*. 104:3289-3294.

- Morrell PL, Gonzales AM, Meyer KK, Clegg MT. 2014. Resequencing data indicate a modest effect of domestication on diversity in barley: a cultigen with multiple origins. *J Hered.* 105:253-264.
- Morrell PL, Lundy KE, Clegg MT. 2003. Distinct geographic patterns of genetic diversity are maintained in wild barley (*Hordeum vulgare* ssp. *spontaneum*) despite migration. *Proc Natl Acad Sci U S A.* 100:10812-10817.
- Morrell PL, Toleno DM, Lundy KE, Clegg MT. 2006. Estimating the contribution of mutation, recombination and gene conversion in the generation of haplotypic diversity. *Genetics.* 173:1705-1723.
- Muller HJ. 1964. The relation of recombination to mutational advance. *Mutation Research/Fundamental and Molecular Mechanisms of Mutagenesis.* 1:2-9.
- Muller HJ. 1950. Our load of mutations. *American journal of human genetics.* 2:111.
- Ng PC. 2003. SIFT: predicting amino acid changes that affect protein function. *Nucleic Acids Research.* 31:3812-3814.
- Nordborg M. 2000. Linkage disequilibrium, gene trees and selfing: an ancestral recombination graph with partial self-fertilization. *Genetics.* 154:923-929.
- Poon A, Chao L. 2005. The rate of compensatory mutation in the DNA bacteriophage phiX174. *Genetics.* 170:989-999.
- Poon A, Otto SP. 2000. Compensating for our load of mutations: freezing the meltdown of small populations. *Evolution.* 54:1467-1479.
- Robertson A. 1960. A Theory of Limits in Artificial Selection. *Proceedings of the Royal Society B: Biological Sciences.* 153:234-249.
- Rockman MV. 2012. The QTN program and the alleles that matter for evolution: all that's gold does not glitter. *Evolution.* 66:1-17.
- Schaeffer LR. 2006. Strategy for applying genome-wide selection in dairy cattle. *J Anim Breed Genet.* 123:218-223.
- Schmutz J, Cannon SB, Schlueter J et al. 2010. Genome sequence of the palaeopolyploid soybean. *Nature.* 463:178-183.

- Schubert M, Jónsson H, Chang D et al. 2014. Prehistoric genomes reveal the genetic foundation and cost of horse domestication. *Proc Natl Acad Sci U S A*. 111:E5661-9.
- Schultz ST, Lynch M, Willis JH. 1999. Spontaneous deleterious mutation in *Arabidopsis thaliana*. *Proceedings of the National Academy of Sciences*. 96:11393-11398.
- Shaw FH, Geyer CJ, Shaw RG. 2002. A comprehensive model of mutations affecting fitness and inferences for *Arabidopsis thaliana*. *Evolution*. 56:453-463.
- Shoemaker DD, Lashkari DA, Morris D, Mittmann M, Davis RW. 1996. Quantitative phenotypic analysis of yeast deletion mutants using a highly parallel molecular bar-coding strategy. *Nat Genet*. 14:450-456.
- Simmons MJ, Crow JF. 1977. Mutations affecting fitness in *Drosophila* populations. *Annu Rev Genet*. 11:49-78.
- Simons YB, Turchin MC, Pritchard JK, Sella G. 2014. The deleterious mutation load is insensitive to recent population history. *Nat Genet*. 46:220-224.
- Smith JM, Haigh J. 1974. The hitch-hiking effect of a favourable gene. *Genet Res*. 23:23.
- Song Q, Hyten DL, Jia G, Quigley CV, Fickus EW, Nelson RL, Cregan PB. 2013. Development and evaluation of SoySNP50K, a high-density genotyping array for soybean. *PLoS One*. 8:e54985.
- Tiffin P, Ross-Ibarra J. 2014. Advances and limits of using population genetics to understand local adaptation. *Trends Ecol Evol*. 29:673-680.
- Wallace B. 1970. Genetic load, its biological and conceptual aspects. Prentice-Hall.
- Wright SI, Bi IV, Schroeder SG, Yamasaki M, Doebley JF, McMullen MD, Gaut BS. 2005. The effects of artificial selection on the maize genome. *Science*. 308:1310-1314.
- Yampolsky LY, Kondrashov FA, Kondrashov AS. 2005. Distribution of the strength of selection against amino acid replacements in human proteins. *Hum Mol Genet*. 14:3191-3201.

Author Contributions

TJYK, RMS, PT, and PLM designed the research. KPS and PLM provided input on which

barley lines to sample, and RMS provided sequence data for soybean lines. Barley read mapping, variant calling and assessment with SIFT and PolyPhen2 were performed by TJYK and CL. Soybean data analysis was performed by FF with assistance from TJYK. Code for the likelihood ratio test was developed by TJYK, PJH, and JCF. Breeding history and causative mutations list were provided by MM. TJYK and PLM wrote the manuscript.

Table Captions

Table 1: Counts of SNPs in various classes in thirteen barley samples. Numbers reported are comparisons against the reference genome, which makes it possible to include noncoding variants, where ancestral state cannot be estimated unambiguously.

Table 2: Per-method and per-sample counts of deleterious variants for barley. Numbers reported are comparisons against ancestral state. The proportion of nonsynonymous variants that is inferred to be deleterious by each prediction method in each accession is shown in parentheses.

Table 3: Counts of SNPs in various classes in seven soybean samples. Numbers reported are comparisons against the reference genome, which makes it possible to include noncoding variants, where ancestral state cannot be estimated unambiguously.

Table 4: Per-method and per-sample of counts of deleterious variants in soybean. Numbers reported are comparisons against ancestral state. The proportion of nonsynonymous variants that is inferred to be deleterious by each prediction method in each accession is shown in parentheses.

Table S1: Accessions used in this study. The final coverage reported is the average depth over the targeted region.

Table S2: Origin information for the barley accessions used in this study.


Table S3: Origin information for the soybean accessions used in this study.

Table S4: List of all species and genome assembly versions, annotation versions, and data sources


for sequences used in the likelihood ratio test.

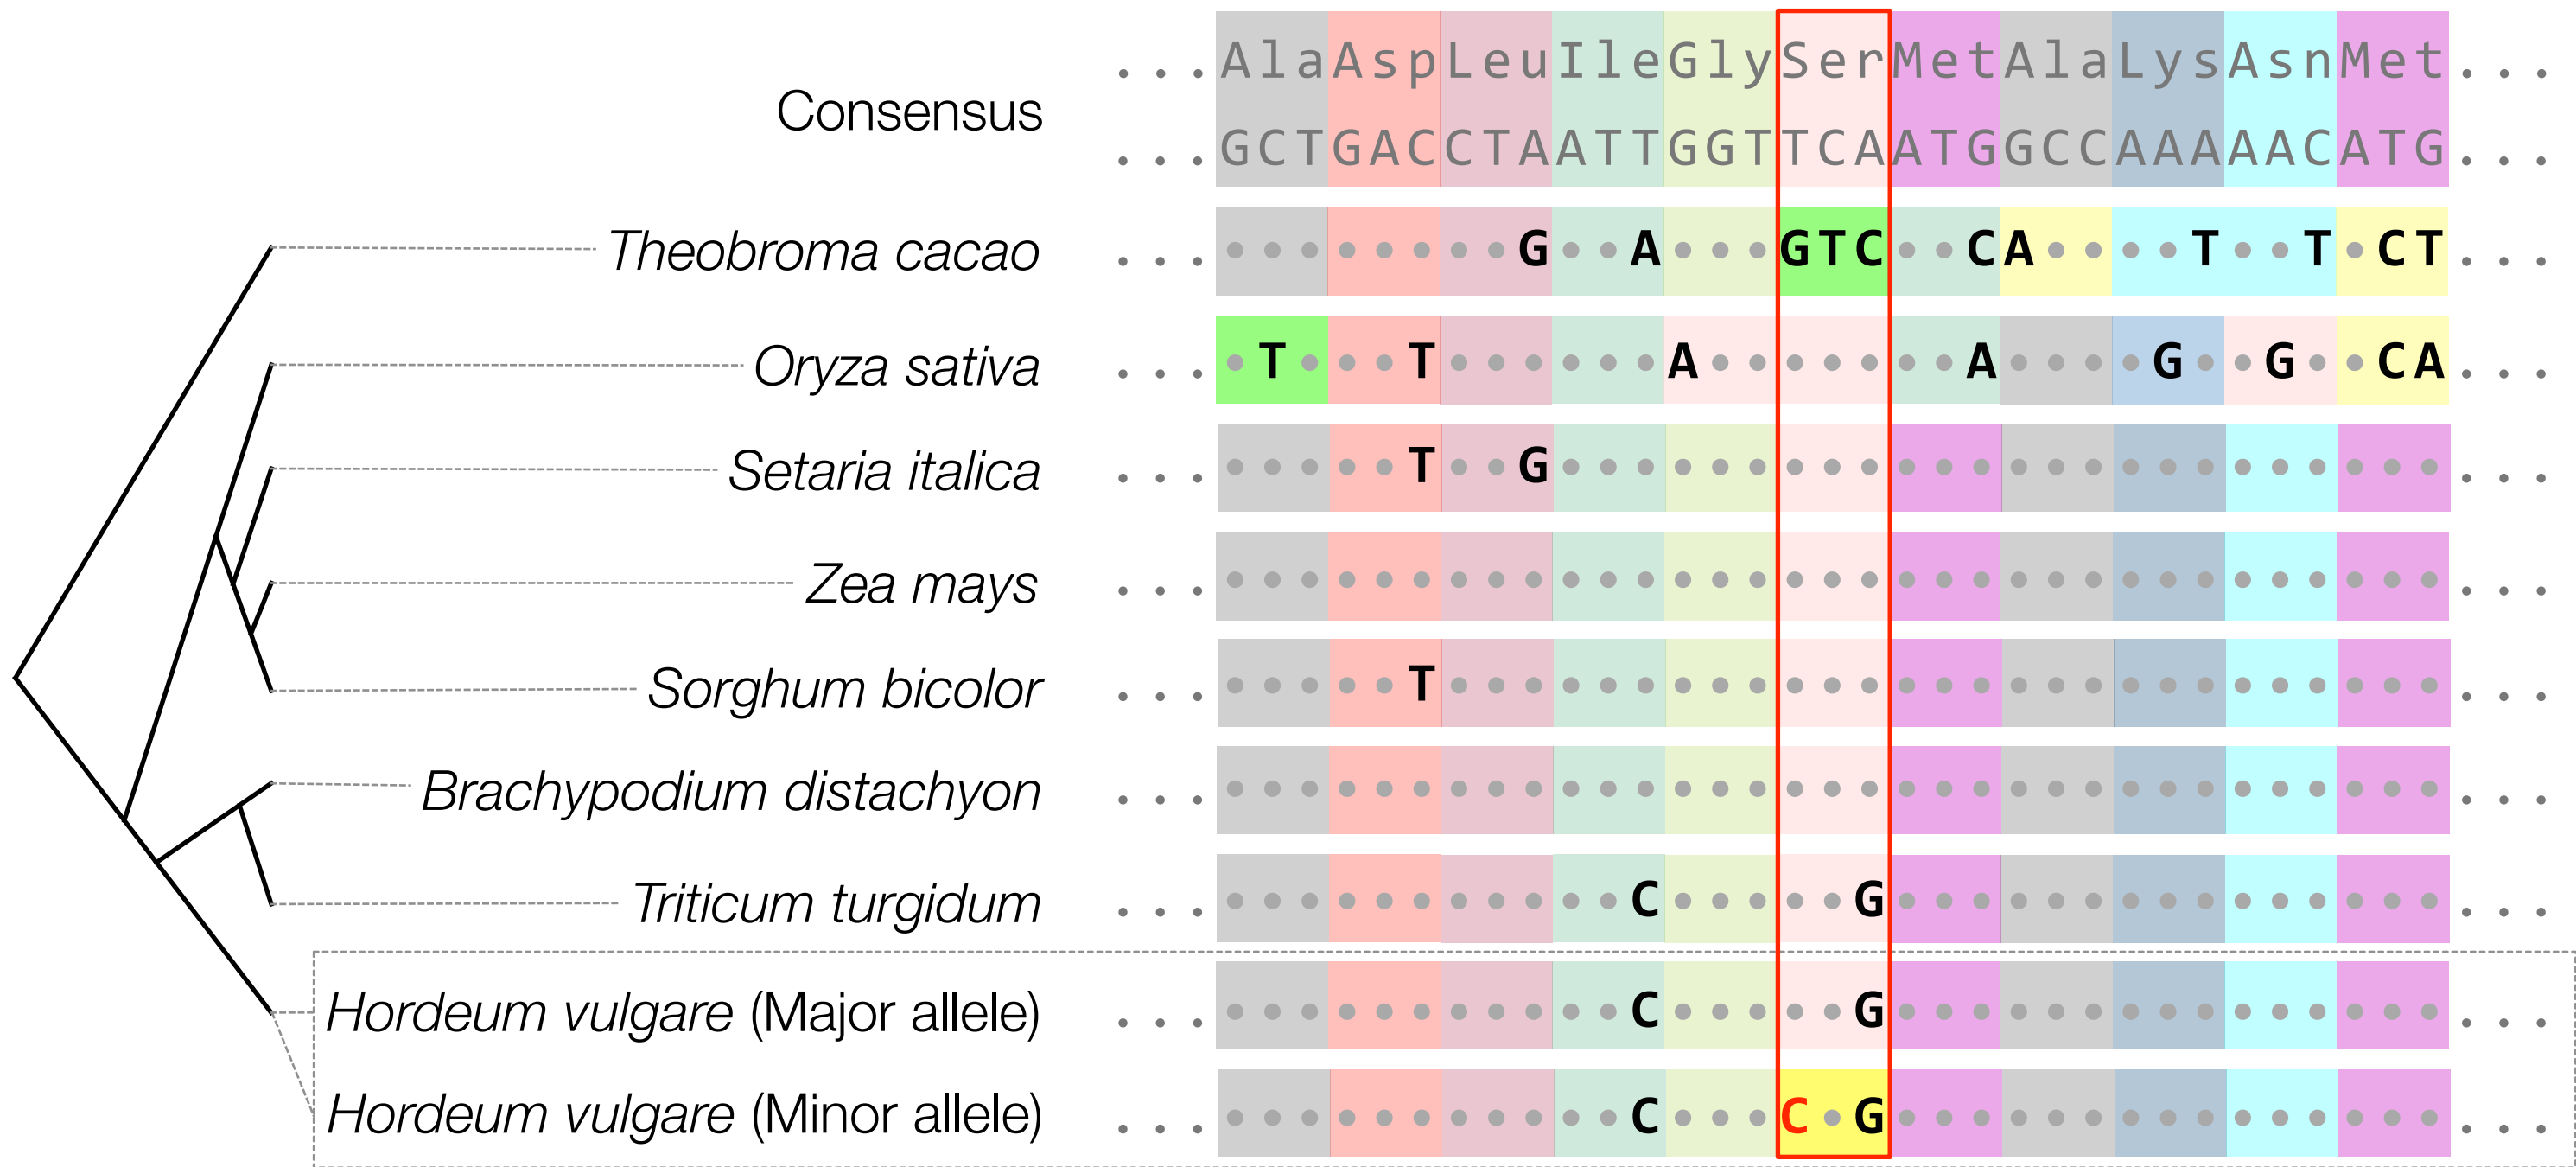
Table S5: List of cloned genes with SNPs causing phenotypic differences, and predictions for each SNP. Causative SNPs annotate as deleterious with a higher frequency than the genomic average.

A)




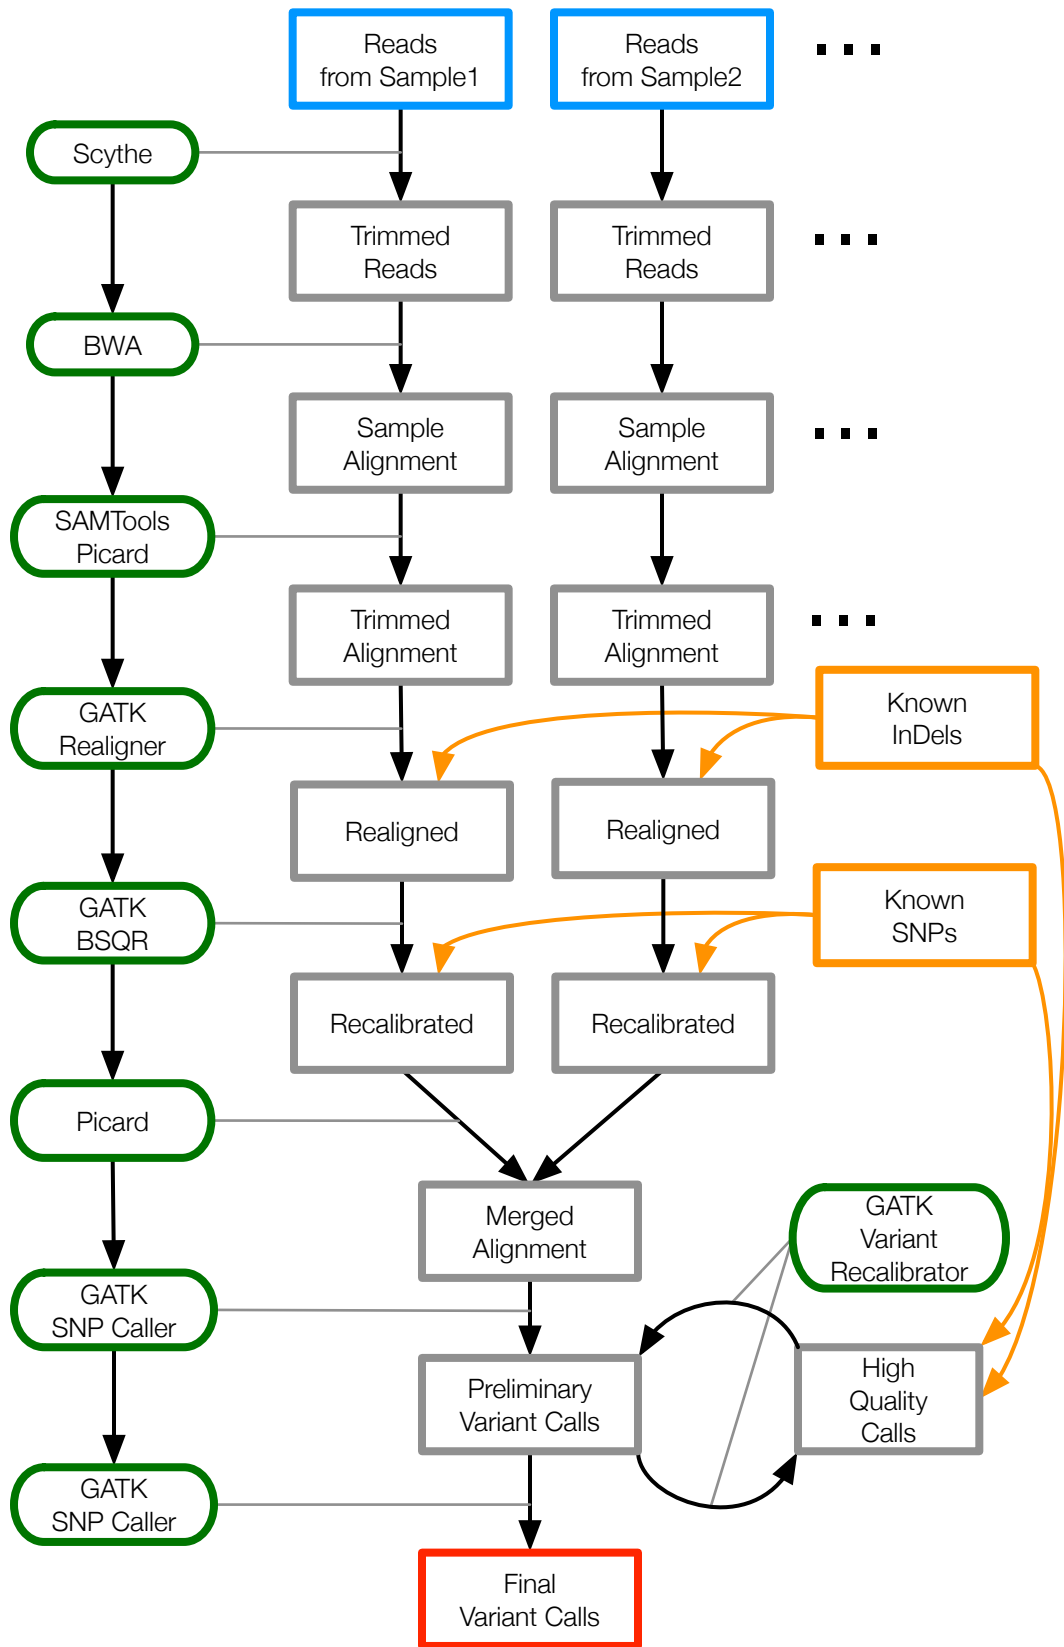
B)






Proportion of Nonsynonymous SNPs Called Deleterious and Recombination Rate on Gm03






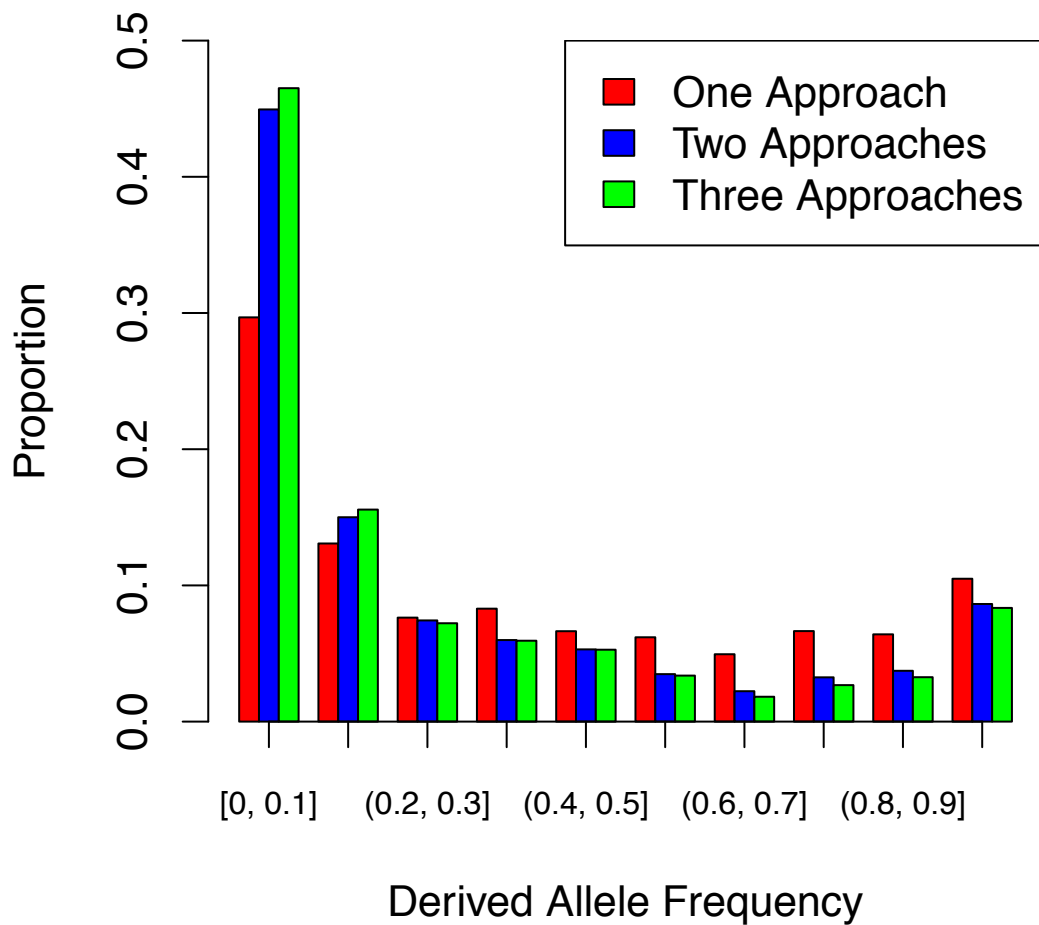
Distribution of Heterozygosity for Deleterious and Tolerated SNPs




Proportion of Nonsyn. SNPs Called Deleterious Over Recombination Rate



A)



B)



Distribution of Grantham Scores

