

Supplementary Information

Contents

1	Design and construction of the combinatorial CRISPR gene-drive system and library	3
1.1	Strains	3
1.2	Mutations and their selection	3
1.3	Construction of guide crRNA plasmids	4
1.4	Final barcoding procedure	5
1.5	Hierarchical mating procedure	6
2	Genotype verification	8
2.1	Whole-genome sequencing	8
2.2	Locus-specific multiplex PCR	9
2.3	Counting alleles for each locus in each well	9
2.4	Statistical inference of gene-drive failures	10
2.5	Other genotyping	11
3	Combinatorial indexing and sequencing of barcodes	12
3.1	Combinatorial pooling and sequencing	12
3.2	Barcode assignment to single wells	12
4	Bulk phenotyping	13
4.1	Growth experiments	13
4.2	Amplicon barcode sequencing	13
4.3	Fitness inference for time-dependent barcode frequencies	14
4.4	Comparison between technical, biological replicates	15
5	Quantitative analysis of epistasis	17
5.1	Estimation of parameters	17
5.2	Variance partitioning	18
6	Analysis of fitness-correlated trends	18
6.1	Fitting regression slopes to determine fitness-correlated trends	19
6.2	Decomposition of fitness-correlated trends	23
6.3	Quantifying the effect of landscape size in the analysis of fitness-correlated trends	27
7	Captions for Data Tables	32
7.1	Data Table S1	32
7.2	Data Table S2	32
7.3	Data Table S3	32

List of Figures

S1	gRNA integration plasmids	5
S2	Barcoding plasmid	6
S3	Mating scheme	8
S4	Comparison of technical and biological replicates for all ploidies and assay environments	16
S5	Slopes and their reversions in the $\Delta\phi$ formulation	20
S6	Comparison of fitness correlated trends for a simple case	21
S7	Comparison of fitness correlated trends for a complicated case	22
S8	Comparison of fitness correlated trends for a case where the orientation matters . . .	22
S9	Effect of parameters on global regression in the ϕ_{wt}/ϕ_{mut} formulation	24
S10	Absolute fit comparison in term removal analysis.	25
S11	Relative fit ratios for each locus.	26
S12	Terms, epistasis sufficient to recapitulate global slopes.	27
S13	Scatterplots of ϕ_{wt}/ϕ_{mut} for all loci in haploid form.	28
S14	Scatterplots of ϕ_{wt}/ϕ_{mut} for all loci in homozygous form.	29
S15	FCT resolved as idiosyncratic in nature with respect to landscape size in 4NQO (haploid)	31
S16	FCT resolved as idiosyncratic in nature with respect to landscape size at 37C (haploid)	31
S17	Final relative fit ratio as a function of reproducibility in biological replicates	32

List of Tables

S1	Mutations constructed in the experiment	4
----	---	---

1 Design and construction of the combinatorial CRISPR gene-drive system and library

3 1.1 Strains

4 The two parental strains used in this study, YAN548 and YAN564, differ at their mating type and
5 are derived from the BY4742 [43] (S288C: MATa, his3 Δ 1, ura3 Δ 0, leu2 Δ 0, lys2 Δ 0) with several
6 modifications required for our combinatorial CRISPR gene-drive strategy. We chose to work in this
7 background due to its history in studies of epistasis in yeast [4] and ease of transformation [44].

8 S288C is a poor sporulator [45], and we introduced the RME1 promoter allele known to in-
9 crease sporulation efficiency (ins-108A) in BY4742, creating YAN404. YAN407 was generated
10 from YAN404 by mating-type switching using a centromeric plasmid carrying the HO endonuclease
11 (pAN216a_pGAL1-HO_pSTE2-HIS3_pSTE3-LEU2). We then introduced the Cre recombinase un-
12 der the control of the galactose promoter at the YBR209W locus using Delitto Perfetto [46], yielding
13 YAN525 and YAN526. The CAN1 gene was subsequently replaced with a mating type reporter
14 construct [47] (pSTE2-SpHIS5-pSTE3-LEU2) which expresses the HIS5 gene from *Schizosaccha-*
15 *romyces pombe* (orthologous to the *S. cerevisiae* HIS3) in MATa cells, and the LEU2 gene in MAT α
16 cells.

17 Cas9 was introduced close to the HO locus under the control of an estradiol-inducible promoter
18 [48] (HO::SpCas9-B112-ER), generating the final strains YAN548 and YAN564. Preliminary work
19 has shown that 2 μ M β -estradiol is sufficient for robust Cas9 induction.

20 Starting strains containing specific mutations were constructed using dsDNA oligo-mediated re-
21 pair using Cas9-mediated double-strand break. To do so, we created a centromeric plasmid carrying
22 the URA3 gene that expressed the guide-RNA. Yeast cells were grown with β -estradiol to induce
23 Cas9, and transformed at log-phase with the guide-RNA expressing plasmid and a double-stranded
24 DNA oligonucleotide with the desired mutation. Cells were then recovered on SD-URA with β -
25 estradiol to maintain expression of Cas9 and the guide-RNA. A parallel transformation can be done
26 to assess the targetting efficiency as an efficient guide-RNA usually leads to far fewer surviving
27 colonies during the transformation due to the toxicity of unrepaired Cas9-mediated double-strand
28 break. Large colonies from the transformation were then grown in YPD overnight and spread on
29 media containing 5-FOA (1 g/L) to counterselect the plasmid expressing the guide-RNA. All strains
30 were then verified by Sanger sequencing.

31 1.2 Mutations and their selection

32 Mutations for our combinatorially-complete fitness landscape were chosen based on several factors.
33 First, we used prior information from published and unpublished experiments that suggested fitness
34 effects for our mutations in at least one environment. Second, due to the need to minimize guide-
35 RNA recognition after the desired mutation is made, we focused on amino acid changes because
36 synonymous mutations could also be incorporated. Third, mutations were chosen that would
37 target a variety of cellular processes to maximize our ability to detect global epistasis. Finally,
38 mutations were chosen that could be efficiently made and not negatively impact our CRISPR-Cas9
39 system described here (i.e., mutations should not make strains sterile, impair sporulation, or impact
40 galactose metabolism).

Mutation	Sequence Information	Reference
----------	----------------------	-----------

WHI2 L262S Chr XV	Guide RNA: ATGGATATGTTGTGCTCCTC L262S DNA: GAcATGagtTGtTCCTCCGGA L262L DNA: GAcATGcTaTGtTCCTCCGGA	[37]
PMA1 S234C Chr VII Essential gene	Guide RNA: TGCTATTACTGGTGAATCCTT S234C DNA: ACTGGTGAATgccTtGCTGTC S234S DNA: ACTGGTGAATCccTtGCTGTC	[38]
MKT1 D30G Chr XIV	Guide RNA: ATGGTTGACGTCTATATCCA D30G DNA: ACCCTGGgaATtGAtGTtAAC D30D DNA: ACCCTGGAcATtGAtGTtAAC	[35]
RHO5 G10S Chr XIV	Guide RNA: ATAATTGGTGATGGTGCAGT G10S DNA: ATatcaGAcGGaGCAGTAGGT G10G DNA: ATaGGaGAcGGaGCAGTAGGT	Our lab
AKL1 S176P Chr II	Guide RNA: TCGCGATGGATCAAGGACAC S176P DNA: CCTGTGcCtcTaATtCAcaGa S176S DNA: CCTGTGTcTcTaATtCAcaGa	[33]
BUL2 L883F Chr XIII	Guide RNA: CACAAACACGTTTCAAGATT L883F DNA: TGCCCAATtTcGAgAcTtGT L883L DNA: TGCCCAATtTgGAgAcTtGT	[34]
FAS1 G588A Chr XI Essential gene	Guide RNA: AATCGGTAGACCACCTTTAT G588A DNA: ATCGcacGtCCtCCaTTATT G588G DNA: ATCGGacGtCCtCCaTTATT	[36]
NCS2 H71L Chr XIV	Guide RNA: CTGAATCAGAATGTGATAAG H71L DNA: CTCCCCTTtgagtttgagtGA H71H DNA: CTCCCCTTtgagtCAcagtGA	[32]
SCH9 P220S Chr VIII	Guide RNA: TCTAATGGTCCTGAGTCACT P220S DNA: AAcGGatCaGAaTCACTAGGC P220P DNA: AAcGGaCCaGAaTCACTAGGC	[39]
RPI1 E102D Chr IX	Guide RNA: GTAATGAATGCTATATCCTC E102D DNA: GAGCCTGAcGAcATtGctTTC E102E DNA: GAGCCTGAaGAcATtGctTTC	Our lab

Table S1: Mutations constructed in the experiment. Lower case letters represent mutated sequences with respect to the wild-type DNA.

41 1.3 Construction of guide crRNA plasmids

42 Our combinatorial CRISPR gene-drive system allows a hierarchical construction of guide crRNA
43 arrays into a benign locus, by taking advantage of Cre-Lox recombination. Previously, we identified
44 three orthogonal and unidirectional recombination sites that are necessary for our design. Briefly,
45 our gene-drive system makes use of three types of recombining plasmids with three distinct pairs of
46 drug markers, which we refer to as type HygMX-KanMX, KanMX-NatMX, and NatMX-HygMX.
47 The three drug markers - HygMX, KanMX, and NatMX - are resistance cassettes for hygromycin
48 B, G418, and nourseothricin, respectively, and differ additionally by the use of paralogous TEF
49 promoters and synthetic terminators as in [31]. Each type is based on an HO-targeting plasmid
50 pAN3H0a (Figure S1), which contains the two drug marker cassettes for selection as well as ho-
51 mologous sequences that lead to integration of insert sequences with high efficiency. The insert

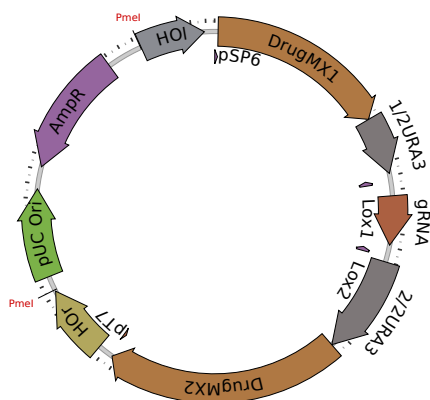


Figure S1: gRNA integration plasmids. We use three integration plasmids with different drug markers, Lox sites, and URA3 frameshift configurations as explained in Section 1.3.

52 sequences between the two drug cassettes contain one of 10 guide-RNA cassettes (each with an
 53 SNR52 promoter mutated at non-functional regions to reduce the rate of unintended homologous
 54 recombination, the guide-RNA, the structural RNA element and the SUP4 terminator [49]). In
 55 addition, each drug marker is linked to their own half of URA3 (frameshifted for each drug such
 56 that the first half of URA3 only functions properly when the correctly framed second half of URA3
 57 is used) which contains a splice donor or acceptor (from QCR10 [50]) and their own orthogonal
 58 Lox site (LoxP, Lox2272, or Lox5171, with arm mutations to allow only a single recombination
 59 event between them [31]). In the configuration found at integration, the URA3 is not functional.
 60 However, when recombined properly by Cre recombinase, a configuration which brings like drug
 61 markers on the same chromosome (HygMX-HygMX, for example) will produce a functional URA3,
 62 which we can select with media lacking uracil and counterselect with media containing 5-FOA.

63 This system allows diploids created by mating two strains with compatible marker configurations
 64 to be selected on media containing all three drugs (described later in section Section 1.5). Compat-
 65 ible configurations will always include a common drug that will yield a functional URA3 after re-
 66 combination. For example, the HygMX-KanMX configuration is compatible with KanMX-NatMX
 67 (which will form HygMX-NatMX and KanMX-KanMX after recombination) or with NatMX-
 68 HygMX (which will form NatMX-KanMX and HygMX-HygMX after recombination). The re-
 69 combined 'landing pads' are thus compatible with each other (for example, HygMX-NatMX is
 70 compatible with NatMX-KanMX, which when recombined will form HygMX-KanMX and NatMX-
 71 NatMX).

72 1.4 Final barcoding procedure

73 To allow bulk phenotyping of the strains, we introduced a 22mer DNA barcode (16 random nu-
 74 cleotides and 6 known spacer nucleotides) alongside a complete LYS2 ORF at the *LYS2* locus via
 75 homologous recombination in the AKL1-RPI1 double-mutant strains prior to the final mating step.

76 To produce a library of uniquely barcoded plasmids, we generated an entry vector with 702 bp
 77 homologous region upstream of the LYS2 deletion, the deleted 293 bp region immediately upstream
 78 of the ORF, the 4179 bp LYS2 ORF, and then a 39 bp tGuo1 synthetic terminator. Downstream
 79 of this terminator was a primer-binding site, pBC1, followed by the *ccdB* gene, which is toxic in

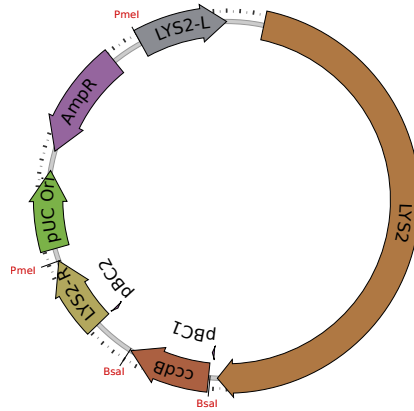


Figure S2: Barcoding plasmid before barcode insertion. We integrate a random barcode at the LYS2 locus to uniquely tag each individual in the pool.

80 *E. coli* strain DH10B. This gene was followed by 300 bp of semi-random DNA sequence (used as
 81 “filler” for obtaining PCR bands distinct from primer dimer bands), the pBC2 primer-binding site,
 82 and 589bp homologous to the region immediately downstream of the LYS2 deletion. Barcodes
 83 were cloned into this plasmid at the *ccdB* locus via Golden Gate assembly [51, 52] in 8 independent
 84 replicates, separately cloning in DH10B via electroporation and selecting on LB+Ampicillin sodium
 85 salt (100 µg/mL) agar plates (1% tryptone, 0.5% yeast extract, 0.5% sodium chloride, 1.5% agar)
 86 after an hour of recovery in SOC media (2% tryptone, 0.5% yeast extract, 8.56mM sodium chloride,
 87 2.5mM potassium chloride, 10mM magnesium chloride, 10mM magnesium sulfate, 20mM glucose).
 88 Plates, which bore at least 30,000-40,000 transformant colonies each, were each scraped and cultured
 89 in 5mL LB+Amp media prior to miniprepping to isolate plasmid.

90 To barcode AKL1-RPI1 double mutants, we first isolated 10 individual colonies of each of the
 91 4 possible double-mutant genotypes. We split these 10 colonies into two sets of 5. Each set of
 92 5 colonies was cultured, pooled, and transformed with one of the eight barcode plasmid libraries,
 93 which had previously been cut with PmeI to linearize the region for integration. Transformants
 94 were selected on SD-Lys agar plates and, to the best of our abilities, individually picked into SD-Lys
 95 media for continued purifying growth.

96 1.5 Hierarchical mating procedure

97 The basic procedure for a cycle of mating, drive, recombination, and sporulation is as follows:

98 Strains with compatible guide-RNA “landing pads” and opposite mating type were mixed to
 99 generate diploids in YPD plus ampicillin (100 µg/mL) via mating for 12-24 hours. Cells were then
 100 passaged to YPG (1% yeast extract, 2% peptone, 2% galactose) plus ampicillin liquid media con-
 101 taining hygromycin B (at 300 µg/mL), geneticin (at 200 µg/mL), and nourseothricin (at 100 µg/mL)
 102 to select for diploids, with selection sustained for at least 3 generations. Cells were then transferred
 103 to YPG containing all four drugs and at least 2 µM β -Estradiol to induce Cre-recombinase and
 104 Cas9, with selection for at least 10 generations. This generates homozygous diploids at the loci
 105 targeted by Cas9, and combines the guide-RNA from the homologous HO loci onto the same
 106 chromosome. The cells were then grown in SD-Ura with β -estradiol for at least 15 generations to
 107 select for successful Cre-Lox recombinants. They then were induced to sporulate by 16-24 h growth

108 in YPA (1% yeast extract, 2% peptone, 1% potassium acetate) followed by culture in SPO media
109 (1% potassium acetate, 0.005% zinc acetate). After 3-5 days of sporulation, haploids containing all
110 the mutated loci and recombined gRNA loci were selected with at least 15 generations of growth
111 in S/MSG-D (1.67% yeast nitrogen base lacking ammonium sulfate, 1% monosodium glutamate,
112 2% dextrose) lacking histidine or leucine (selecting for MAT_a and MAT_α respectively), containing
113 two of the three antibiotic drugs (depending on the landing pad configuration) and 1g/L 5-FOA
114 to counterselect diploids. Finally, selected populations were screened for “leakers” by growing a
115 single colony or a small number of cells (less than about 1000) in YPD, followed by a transfer into
116 YPD containing the drug to which the desired haploids should not be resistant. Only specimens
117 sensitive to this third drug were preserved as a frozen archive and then passaged into the next
118 mating-drive-recombination-sporulation step.

119 In practice, this procedure included a variety of manipulations. This range of manipulations
120 demonstrates that our method is flexible and can be adapted to work within various technical
121 constraints. For example, when handling few strains, microtiter plates are not necessary and the
122 protocol can be performed in standard culture tubes. In the case of the initial double-mutant
123 mating, for instance, mating was in most cases conducted on YPD-agar patches, which were then
124 scraped and transferred into the YPG diploid selection media. All other matings were conducted
125 in about 90 μ L YPD liquid media in wells of 96-well round-bottom microtiter plates. Similarly,
126 selection of haploids after sporulation was sometimes conducted in microtiter plates (128 μ L total
127 volume), and other times by streaking to individual colonies on SD-Leu or SD-His agar plates (with-
128 out 5-FOA counterselection). For all cycles except the experiment’s final cycle, individual colonies
129 were isolated and screened at the conclusion. Finally, depending on the scale of the cycle, diploid
130 selection, recombinant selection, presporulation, and sporulation steps were conducted in either
131 microtiter plates (shallow for selections (128 μ L media), 2-mL deep-well plates for presporulation
132 and sporulation) or test tubes (5 mL media unless otherwise stated).

133 Presporulation: Microtiter plate-based presporulation was carried out by pipetting 20 μ L sat-
134 urated SD-Ura culture into 480 μ L of YPA. Plates were shaken at 1050 rpm at 30°C for 24 hours
135 under a breathable membrane (VWR, 60941-086) before sporulation. Tube-based presporulation
136 was carried out by inoculating 5 mL YPA with 150 μ L saturated SD-Ura culture and incubating
137 on a roller drum at 30°C for 16-24 hours.

138 Sporulation: Microtiter plate-based sporulation was carried out by pelleting presporulated cells
139 at 2000 g for 2 min, washing by resuspension in 400 μ L water, pelleting once again, and resuspending
140 in 400 μ L sporulation media. These plates were sealed with a breathable membrane, secured with
141 tape to plate shakers, and shaken at 1350 rpm at room temperature for 4-5 days. Tube-based
142 sporulation was carried out by pelleting tube-presporulated cell cultures and resuspending in 2 mL
143 sporulation media, incubating at room temperature on a roller drum for 3-4 days.

144 Homozygotes from the final cycle were incubated for 5 generations in YPD+Amp prior to
145 archival freezing, but only after fully selecting for recombination of the landing pad loci with SD-
146 Ura+ β -estradiol.

147 The final 20 generations of haploid selection in the final cycle were conducted in typical haploid
148 selection media, but lacking lysine, in order to select only for those haploids which retained the
149 barcode next to the LYS2 marker (which segregated in a Mendelian fashion).

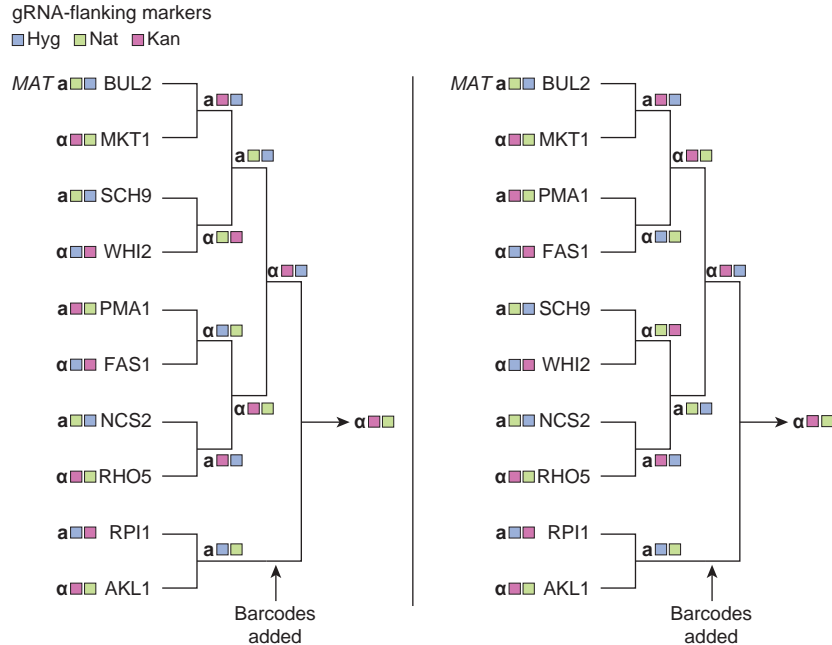


Figure S3: Parallel mating scheme. Biological replicates of the final strains were created via different mating paths.

150 2 Genotype verification

151 2.1 Whole-genome sequencing

152 To verify the lack of systematic off-target Cas9-mediated modifications, and to rule out pervasive
 153 aneuploidies, we performed whole-genome sequencing on 96 random clones (3 random wells from
 154 each of 32 plates which contained 64 different strains based on the mating procedure outlined in
 155 Section 1.5) [53] and sequenced each to approximately 100x coverage. This identified a single case
 156 of aneuploidy with elevated read counts at three chromosomes that were consistent with disomy
 157 [54]. In addition, it identified five credible non-synonymous mutations occurring on more than 1
 158 strain (strongly indicating that the mutations were introduced in the hierarchical mating scheme
 159 described in Section 1.5). Two of these (in ERG6 and QRI7) were present in just two strains each,
 160 and the other three (in SPT7, HSL7, and FRS1) were present in 5, 6, and 33 strains, respectively.
 161 In addition, some extra mutations were identified in single clones, which is not inconsistent given
 162 the rate of mutations during meiosis (70% of clones had no mutations, 10% had one, 5% had two,
 163 and the rest had poor sequencing coverage leading to what we believe are bad variant calling).
 164 These results suggest that Cas9 does not introduce a gross excess of off-target mutations in the
 165 genome, and that although unintended mutations do occur in our system (due to Cas9 or meiosis)
 166 these mutations are unlikely to dominate the estimation of parameters for modeling the fitness
 167 landscape. Notably, as explained in Main Text, biological replicates (independent crosses) were
 168 typically in agreement with each other.

169 To understand whether the mutations in SPT7, HSL7, and FRS1 may have systematic effects on
 170 our genotypes, we looked at whether they were present exclusively on any single- or double-mutant
 171 backgrounds. We found that the FRS1 mutation was present across most backgrounds, but the

172 mutations in SPT7 and HSL7 were only present on specific AKL1-RPI1 backgrounds. Notably,
173 fewer than half of the instances of these backgrounds in our WGS data bore these mutations.

174 **2.2 Locus-specific multiplex PCR**

175 To genotype the entire haploid library at all 10 primary loci and 3 putative segregating off-target
176 loci (FRS1, SPT7, and HSL7), we pursued a multiplexed strategy. We began by lysing all 2048
177 wells (not all of which contained cells) with 20 μ L yeast lysis buffer (5mg/mL Zymolyase 20T,
178 100mM Sodium Phosphate pH 7.4, 10 mM DTT) and 5 μ L of cells straight from the freezer stock.
179 The enzymatic reaction was placed at 37°C for at least 45 min and then at 95°C for 2 min. The
180 released DNA could then be stored in the freezer overnight.

181 Immediately prior to the first-round PCR, we boiled these products for a minute to mix the
182 lysates. We then added and mixed in 25 μ L of water to the lysis products to dilute and facilitate
183 liquid handling. Then, we added 2 μ L of this lysis product to the PCR master mix for the first
184 round PCR, mixing after addition. This master mix was for a 25 μ L Phire reaction and contained
185 1.3 μ L of pooled 100 μ M primers. These primers represented all 13 loci. The 13 primers that added
186 N7 adapters to the amplicons were common across all wells. The 13 primers that added S5 adapters
187 contained 6 bp inline indices. These indices existed in 8 versions, each unique to a different set of
188 4 plates in the library (54°C annealing, 45s extension). These primers may be found in Data Table
189 S1.

190 The following day, PCR round 1 products were combined into 4 pools, taking 4 μ L from each
191 well. We cleaned up these pools with a 1x bead purification step (AMPure beads by Beckman
192 Coulter) (starting volume = 42 μ L , eluted in 35 μ L). We used KAPA polymerase for a second
193 round of 25 μ L PCRs to anneal unique pairs of S5/N7 indices to the amplicons across 4 reaction
194 plates, using 2 μ L of purified round 1 product (63°C annealing, 45s extension). Several unsuccessful
195 reactions were repeated as necessary with diluted template.

196 Round 2 reaction products were then pooled and cleaned via gel extraction, followed by a final
197 bead purification step to remove any remaining small fragments.

198 The library was sequenced on a NextSeq mid output lane resulting in an average coverage of
199 about 2700x per locus per well in the genotype library. Loci varied in their overall coverage: the
200 average coverage per BUL2 locus was just about 80x, whereas the average coverage for WHI2 was
201 about 7300x. Other than BUL2, all other loci had an average coverage of at least 400x.

202 Some loci for specific wells were missing from our dataset, or otherwise had very low coverage.
203 To patch these holes in our genotyping data, we amplified with locus-appropriate primers in a first-
204 round reaction to anneal S5/N7 adapters. This reaction used Phire polymerase (54°C annealing,
205 45s extension) and 2 μ L of diluted lysate as template. These reaction products were cleaned up
206 with 1x Ampure beads and eluted in 30 μ L water. We took 2 μ L of this reaction product into the
207 second round KAPA Hifi PCR reaction, which annealed pairs of S5 and N7 indices unique to each
208 reaction (63°C annealing, 45s extension). Each reaction product was then cleaned up separately
209 using 0.8x Ampure beads on 6 μ L of reaction product diluted in 10 μ L water. The final product
210 was eluted in 25 μ L and pooled for sequencing on a MiSeq Nano lane.

211 **2.3 Counting alleles for each locus in each well**

212 Once we received the Illumina reads, we counted the number of reads of each allele at each locus in
213 each well. To do this, we followed a procedure similar to [25], examining each read in each 8-well
214 sequencing library (corresponding to individual fastq files) in turn. First, we checked that the first 6
215 bp of read 1 corresponded to a 6-bp inline index, allowing for 1 bp of mismatch. Then, we evaluated

216 read quality by ensuring that the quality score of the 22bp downstream from the inline index was at
217 least 25. If a read met these conditions, we identified the locus associated with the read by checking
218 for the presence of a characteristic 8-bp sequence either upstream or downstream of the defined
219 allele, allowing identical matches only. For reads matching an identifiable locus, we extracted the
220 20- to 23-bp allele, sequentially using a list of decreasingly stringent regular expressions (using the
221 python regex module [55]):

```
222 '(left 8bp)(length of allele)(right 8bp)',  
223 '(left 8bp)(length of allele-2,length of allele+2)(right 8bp)',  
224 '(left 8bp){e≤1}(length of allele)(right 8bp){e≤1}',  
225 '(left 8bp){e≤1}(length of allele-2,length of allele+2)(right 8bp){e≤1}',
```

226 For lists of the exact alleles and 8-bp sequences searched, see Data Table S1.

227 Overall, fewer than 0.5% of reads were excluded on the basis of these criteria, with no more
228 than 1.2% for a single library.

229 All alleles that occurred at least 10 times in at least one well AND were present at at least 1%
230 frequency for the corresponding locus in at least one well were given a unique identifier and assigned
231 as either a WT, Mut, or Other allele. “WT” alleles included properly repaired pseudo-WT alleles
232 plus other versions with some or even none of the desired synonymous changes. This includes loci
233 in which it appears no gene drive occurred (i.e., sequences identical to the unmutated parental BY
234 sequence). “Mut” alleles included any with the desired missense change, regardless of the presence
235 or absence of other synonymous alleles. “Other” alleles included those whose amino acid sequence
236 matched neither the WT nor Mut sequence, including errant missense changes and frameshifts.
237 Any remaining alleles were grouped together and designated “na.”

238 2.4 Statistical inference of gene-drive failures

239 One difficulty of verifying locus correctness by PCR in the final haploid library is that the strains
240 are not clonal (they are derived from the Cas9 gene-drive hierarchical mating procedure, see Sec-
241 tion 1.5). Thus, we needed to remove wells that had evidence of a mixture of genotypes, or strong
242 evidence of the incorrect genotype. We noticed that our multiplex PCR verification protocol in
243 Section 2.2 produced evidence of genotype mixtures at a higher rate than anticipated. However, we
244 observed that these supposedly incorrect wells were found more frequently when post-first-round
245 PCR pools were “mixed” at a given locus (i.e., were expected to have both WT and Mut alleles
246 present). This indicated to us that primers from the first-round PCRs were leaking through, thus
247 incorrectly indexing the reads, and/or PCR chimeras were forming.

248 We developed a statistical model to accurately estimate the true mixture proportion within
249 each well. For each post-first-round pool, we calculated the pool-wide frequencies of all alleles in
250 that pool (based on their unique identifiers). Then, we modeled a pool-wide probability p that a
251 given read is a “true” read and not a chimeric read. We assume that the expected frequency of
252 a false allele in a given well will be $(1 - p) \cdot$ the poolwide frequency of that allele, whereas a true
253 allele in a given well will have an expected frequency of $p + (1 - p) \cdot$ the poolwide frequency of that
254 allele. For a range of values of p , constraining p to be at least 50%, we calculated the probability
255 of the data under a multinomial model and obtained the maximum likelihood estimate of p . As
256 necessary, we constrained the likelihood surface to satisfy the constraint that all alleles must be
257 present at a frequency between 0 and 1, inclusive.

258 After obtaining these adjusted allele frequencies, we set out to distinguish which wells were
259 acceptably versus unacceptably “pure.” Since rates of apparent chimera formation varied signifi-
260 cantly across loci, we developed a separate purity threshold for each locus. We did this by sorting
261 wells by the percent of non-dominant alleles at a given locus, excluding “na” alleles. We plotted

262 these proportions against the ascending rank on the x axis, forming a “hockey stick”-like curve
263 that shoots upward at the high end of the distribution. We found the “elbow” of this curve, i.e.,
264 the proportion of non-dominant allele at which the curve is furthest from a “hypotenuse” line con-
265 necting the end of the handle to the tip of the blade of the proverbial hockey stick. We obtained
266 the following thresholds from this approach:

267 Approximate thresholds
268 BUL2 3.04%, gives 93.3% pure
269 FAS1 1.20%, gives 95.1% pure
270 MKT1 1.13%, gives 96.6% pure
271 NCS2 1.94%, gives 96.3% pure
272 PMA1 2.29%, gives 91.8% pure
273 RHO5 1.58%, gives 94.8% pure
274 SCH9 0.64%, gives 94.5% pure
275 WHI2 0.45%, gives 92.9% pure
276 AKL1 3.27%, gives 95.3% pure
277 RPI1 2.36%, gives 95.1% pure

278 For the sake of comparison, we note that overall drive failure rates inferred from sequencing
279 the quadruple and octuple mutants – which was not done in a pooled, chimera-genic way – were
280 close to 2%. In addition, many gene drive events that in fact failed may not be counted here, since
281 a failed drive event that yields the unmutated WT allele when the WT allele is desired will be
282 retained.

283 All told, 1282 wells matched their expected genotype at all loci (67.8%). Since we had biological
284 replicates of each genotype, 875 out of 1024 possible genotypes (85.4%) were represented among
285 these wells. See Data Table S2 for a complete list of wells, barcodes, and genotypes that passed
286 these filters.

287 2.5 Other genotyping

288 In addition to genotyping the final products of the experiment, we genotyped one or more mutated
289 clones per genotype after each cycle of mating, drive, recombination, and selection before proceeding
290 with the next cycle.

291 Genotyping of the double mutants (after cycle 1) was conducted via Sanger sequencing and
292 visual examination of traces for the expected alleles.

293 Genotyping of the quadruple mutants (after cycle 2) and octuple mutants (after cycle 3) was
294 conducted via next generation sequencing.

295 Quadruple mutants were lysed in 50 μ L yeast lysis buffer (5 mg/mL Zymolyase 20T (Nacalai
296 Tesque), 1 M sorbitol, 100 mM sodium phosphate pH 7.4, and 20 mM DTT), boiled at 95°C for
297 2 minutes and 2 μ L lysed cells were taken into a 25 μ L Phire polymerase PCR reaction with
298 1.25 μ L each of the 4 pairs of appropriate primers for 4 loci, respectively (54°C annealing, 30s
299 extension). After this first round of PCR, we purified the product with 0.8x beads and did a sec-
300 ond round KAPA Hifi polymerase PCR (25 μ L) to append unique S5, N7 indices to each colony
301 isolate. The final product was purified with 0.8x beads once again and sequenced libraries on MiSeq
302 Nano 2x150bp.

303 Octuple mutants were lysed with 5 μ L of saturated culture in 50 μ L yeast lysis buffer as previ-
304 ously described. The boiled lysis product was diluted two-fold, 2 μ L of the lysis was used in into
305 24 μ L Phire polymerase PCR reaction containing 1 μ L of each of 16 10 μ M primers, each of which
306 add the S5, N7 adapter sequences (54°C annealing, 30s extension). Round 1 PCR products were
307 purified via bead cleanup at 0.8x beads ratio, and eluted with 25 μ L water. Before cleanup, some

308 of wells were diluted with an additional 10 μ L to bring volume up (evaporation of PCR reactions
309 is frequent), then 12 μ L taken as starting volume for cleanup. We then performed the round 2
310 PCR reaction with unique pairs of S5 and N7 primers for each well, taking 2 μ L of cleaned up
311 DNA template into a 25 μ L KAPA reaction (63°C annealing, 45s extension time). Round 2 PCR
312 products were diluted with an additional 10 μ L water and pooled (3-4 μ L of each well). We then
313 performed a 0.7x bead cleanup and submitted the final purified pool for NextSeq Mid throughput
314 1x150bp lane.

315 **3 Combinatorial indexing and sequencing of barcodes**

316 **3.1 Combinatorial pooling and sequencing**

317 To map the barcodes to individual wells, we took a combinatorial indexing approach. Uniquely
318 barcoded AKL1-RPI1 double mutants were cultured in the central 64 wells of 32 96-well microtiter
319 plates (rows A-H, columns 3-10). With the help of a Biomek liquid handler, we took 10 μ L of each
320 well-mixed well culture into either of two new 96-well plates, in which wells had been seeded with
321 30 μ L of YPD to facilitate automated liquid dispensing. 70 μ L of pooled culture from each well of
322 these two plates was used to form 8 row-specific pools, and the process was repeated form 8 column
323 pools. Each pool contained approximately 1.1 mL of culture. Separately, for each of the 32 plates,
324 20 μ L from each of 64 wells was pooled to form 32 plate pools of about 1.3 mL each.

325 To prepare libraries for sequencing, we extracted genomic DNA from each of the 48 pools,
326 eluting in 50 μ L water. In an initial PCR step using primers 5xx>pBC1-F and 7xx>pBC2-R, we
327 amplified the barcode loci in each pool, attaching S5 and N7 adapters to each amplicon. For these
328 reactions, we used 0.5-5 μ L of genomic DNA in a 25 μ L Q5 reaction (34 cycles, 54°C annealing,
329 45s extension). After purifying amplicons via a cleanup with 0.8x ampure beads and eluting into
330 33 μ L water, we performed a second round of PCR with 1 μ L of purified DNA template and unique
331 pairs of S5 and N7 primers (KAPA 50 μ L reaction, 34 cycles, 63°C annealing, 45s extension). Final
332 PCR products were pooled, with 2 μ L of each plate pool and 8 μ L of each row and column pool
333 (total volume about 200 μ L). Half of this was taken for a 2-sided bead selection, first with 0.5x
334 beads, and next with 0.2x more beads for a 0.7x selection.

335 Libraries were sequenced on a NextSeq mid-output lane yielding an average coverage of about
336 8700 reads per barcode per pool.

337 **3.2 Barcode assignment to single wells**

338 Combinatorial indexing allows one to uniquely triangulate a barcode to a specific well. However,
339 errors due to sequencing, apparent cross-contamination due to chimeric reads or lower read cover-
340 age for some particular combinatorial pool can make some assignments ambiguous. We therefore
341 performed this assignment using a greedy procedure. First, barcodes that uniquely map to a single
342 well were identified. This yielded 2332 barcodes (out of 2348) that mapped to 2029 wells. Evidently,
343 some wells contained multiple barcodes that stem from imprecise colony picking. 16 barcodes ap-
344 peared to map ambiguously to multiple wells. Manual inspection found that 12 of these could be
345 explained by spurious reads in other pools, which meant we only had to remove four wells with
346 conflicting barcodes.

347 We additionally found about 40 wells that appeared to grow extremely slowly in SD-Ura+ β -
348 estradiol+Amp, perhaps due to picking petite colonies. All were of the same AKL1-RPI1 genotype
349 (AKL1 176S, RPI1 102D) and from the same barcode transformant pool, leading us to believe this
350 may be due to private mutations in one of the 5 replicate pooled colonies. We manually identified,

351 removed, and repicked these remaining wells from the opposite transformant pool of that same
352 AKL1-RPI1 genotype, in which we had seen no issues. In addition, we repicked about a dozen
353 barcode transformant colonies for wells with unassigned barcodes.

354 The second set of barcodes was assigned again by combinatorial indexing, this time with only
355 8 rows and 8 columns, and some spurious remaining wells that did not have a well-defined barcode
356 were also confirmed by Sanger sequencing.

357 4 Bulk phenotyping

358 4.1 Growth experiments

359 The complete frozen pool was grown in 5 mL YPD by inoculating approximately 10^7 total cells
360 to produce the starting population. We then diluted these populations by 1:2⁷ daily by passaging
361 781 μ L into 5 mL fresh media (of some particular environment) in 15 mL culture tubes on roller
362 drums. Whole population pellets, obtained from 1.5 mL of saturated culture, were stored imme-
363 diately at -20°C for later sequencing. As previously described [31], this protocol results in about
364 7 generations per day, with a daily bottleneck size of about 10^8 in most assay environments. We
365 performed two replicates of each assay and sampled for 49 generations (7 timepoints). Only 5
366 timepoints (representing 7, 14, 28, 42, and 49 generations) were sequenced.

367 The six environments chosen were: YPD + 0.4% acetic acid (YPDA), YPD + 6 mM guanidium
368 chloride (gu), YPD + 35 μ M suloctidil (suloc), YPD + 0.8 M NaCl (salt), YPD at 37°C (37C),
369 and SD + 10 ng/mL 4NQO (4NQO). (All environments besides 37C were at 30°C.) The YPDA
370 environment was chosen because preliminary experiments suggested that it had a tendency to reveal
371 phenotypic variance and it previously had been studied in our lab ([31]). Gu was chosen because of
372 its known large target size from separate work in our lab which identified a change in sign for the
373 effect of the MKT1 D30G mutation. Suloc and 4NQO were chosen because previous work in our lab
374 showed these environments to have low genotype correlation with other YPD-based environments.
375 37C and salt were chosen because several of the genes under study were previously reported to be
376 mutated in evolution under that stressor or be a QTL in that stressor (NCS2 in high temperature;
377 PMA1, RPI1, and RHO5 were all mutated in NaCl evolution experiments; see Table 1.2).

378 The degree of the stressor in suloc, YPDA, salt, and 4NQO environments was chosen empiri-
379 cally to maximize the stress while still permitting 7 generations of growth per day over the entire
380 phenotyping assay.

381 4.2 Amplicon barcode sequencing

382 Genomic DNA from cell pellets were processed as in [31]. Briefly, DNA was obtained by zymolyase-
383 mediated cell lysis (5 mg/mL Zymolyase 20T (Nacalai Tesque), 1 M sorbitol, 100 mM sodium phos-
384 phate pH 7.4, 10 mM EDTA, 0.5% 3-(N,N-Dimethylmyristylammonio)propanesulfonate (Sigma,
385 T7763), 200 μ g/mL RNase A, and 20 mM DTT) and binding on silica mini-preparative columns
386 with guanidine thiocyanate buffer (4 volumes of 100 mM MES pH 5, 4.125 M guanidine thiocyanate,
387 25% isopropanol, and 10 mM EDTA). After binding, the columns were washed with a first wash
388 buffer (10% guanidine thiocyanate, 25% isopropanol, 10 mM EDTA) and then a second wash buffer
389 (80% ethanol, 10 mM Tris pH 8), followed by elution into elution buffer (10 mM Tris pH 8.5).
390 1.5 mL of pelleted cells eluted into 100 μ L routinely provided about 1-2 μ g of total DNA.

391 PCR of the barcodes was performed using a two-stage procedure previously described to attach
392 unique molecular identifiers (UMIs) to PCR fragments (see [31] for a detailed protocol). Primers
393 used in the first-stage PCR contained a priming sequence, a 7-12-nucleotide multiplexing index,

394 8 random nucleotides as UMIs, and an overhang that matched the Tn5 transposome. These
395 two primers had the configurations P1 = TCGTCG GCAGCG TCAGAT GTGTAT AAGAGA
396 CAGNNN NNNNNY YYYYYY AAGGTA CGATTC TGACGC A, P2 = GTCTCG TGGGCT
397 CGGAGA TGTGTA TAAGAG ACAGNN NNNNNN YYYYYY YAGTTG TCTCTG CTCTCG
398 CTA. Here N corresponds to degenerate bases used as UMIs, and Y corresponds to multiplexing
399 indexes.

400 These primers anneal on either side of the barcode sequence integrated just downstream of
401 LYS2, at the pBC1 and pBC2 sites, respectively. After attachment of molecular identifiers to
402 template molecules during three PCR cycles (20 μ L Q5 Polymerase reaction, 50°C annealing, 30s
403 extension), the first-stage amplicons were cleaned using Ampure beads using an automated liquid
404 handling protocol established for a Biomek FXp, with 1.25x Ampure beads, eluting in 35 μ L. Of
405 the elution of this clean-up, 30 μ L was used directly as template for the second-stage PCR with
406 primers that contained multiplexing indexes and adapters that anneal to the Illumina flowcells (P5
407 and P7 primers). After 35 PCR cycles (50 μ L KAPA Hifi Polymerase reaction, 63°C annealing, 30s
408 extension), these final products were then purified using Ampure beads, quantified, and pooled to
409 approximately equimolar concentration. The PCR products were sequenced with a NovaSeq S1 full
410 flow cell (Illumina) by paired-end sequencing (2 x 50 bp, reading 80 bp from the P1 direction and
411 20 bp from the P2 direction).

412 We first processed our raw sequencing reads to identify and extract the indices and barcode
413 sequences as in [31]. Using the barcodes previously identified in Section 3.2, we can make “correc-
414 tions” to reads with sequencing errors by direct lookup of the lowest Levenshtein distance to the
415 dictionary of verified barcodes.

416 Finally, we can calculate the counts of each error-corrected true barcode by removing duplicate
417 reads, using the unique molecular identifiers from the first-stage PCRs. Frequencies calculated
418 from these counts are used to infer fitnesses for all segregants, as explained in Section 4.3. After
419 all filtering, our final mean sequencing coverages were over 1500 reads per barcode per timepoint
420 per replicate (averaged across all assays).

421 4.3 Fitness inference for time-dependent barcode frequencies

422 Strain fitnesses can be inferred from relative barcode frequencies over time (see Refs. [31] and
423 [56] for expanded information on joint inference of fitnesses using barcode frequencies). Briefly,
424 fitnesses are regressed as the change in relative log frequencies of strains against a selected ref-
425 erence per generation. This parameter is approximately the difference in instantaneous growth
426 rate between lineages under exponential growth. Most genotypes in our data are represented by
427 more than one barcode in the same assay (representing biological replicates), and each barcode was
428 measured in two technical replicates. In theory, we could jointly infer the biological replicates and
429 constrain their fitnesses to be equal. This would yield, for a combinatorially complete landscape,
430 exactly 2^N fitnesses which could be fit exactly with 2^N coefficients (later described in Section 5.1).
431 However, strains with the same desired genotype may not always be identical at all other loci in
432 the genome (due to new mutations or off-target effects). By only performing the joint inference
433 on technical replicates, variance left unexplained by a full model containing 2^N coefficients can
434 be regarded as biological variation at other loci and some measurement error (described in more
435 detail in Section 5.2). This joint inference is intuitively similar to a weighted average of the two
436 technical measurements, with weights proportional to the evidence within each replicate (which is
437 a combination of the number of reads and the number of timepoints measured). A standard error
438 for the inferred fitness parameter can be obtained through the inference process by the square root
439 of the inverse of the Fisher information at the maximum likelihood. This standard error can be

440 interpreted as the error that can be attributed to the (overdispersed) binomial sampling error. For
441 our analyses, we removed datapoints with standard error above 1.

442 4.4 Comparison between technical, biological replicates

443 As shown in Figure S3, biological replicates were made for all final strains in the experiment by
444 proceeding through a parallel mating scheme. However, due to gene-drive failures, some strains were
445 not found in replicate, and it may be useful to ask the following questions: 1) How trustworthy
446 are the strains without replicates? and 2) What is the average effect of unintended mutations
447 introduced within our cross? To answer these questions, we can compare the inferred fitnesses of
448 technical (comparing the same barcode across assays) and biological replicates (comparing barcodes
449 that correspond to the same expected genotype).

450 Decomposing the observed phenotypic variance due to measurement error can be done by the
451 standard reliability estimates. The Pearson’s correlation coefficient between two technical replicates
452 is an estimate of the R^2 between the true fitness value and one fitness measurement for the barcode.
453 If one takes the mean of the r technical replicates, then:

$$\frac{\sigma_{err}^2/r}{\sigma_{gen}^2 + \sigma_{err}^2/r} = \frac{1 - \langle \rho_{r_i, r_j} \rangle}{1 + (r - 1) \langle \rho_{r_i, r_j} \rangle}. \quad (1)$$

454 Decomposing the phenotypic variance due to extra variance in the genetic component can be
455 done by a similar process, by comparing the measurement values between strains bearing different
456 barcodes but expected to have the same genotype. Here, to perform this calculation, we constrain
457 ourselves to pairs of strains with the same genotype that each have a single barcode in their well
458 so that a single comparison can be made. The correlation coefficient between biological replicates
459 can be interpreted in a similar way to technical replicates, but the deviation from 1 here will reflect
460 both error due to extra variation in the genome and error due to measurement error (but without
461 tube-to-tube variation). For the purpose of our manuscript, we assume that this tube-to-tube
462 variation is negligible.

463 In plots of technical and biological replicates, density-based coloration was determined by cal-
464 culating each point’s mean distance to its five nearest neighbors. Distances were transformed using
465 the scikit MinMaxScaler() function and plotted with normalized colors based on a reversed viridis
466 colormap.

467 Technical and biological replicate comparisons for all data can be viewed in Fig. S4.

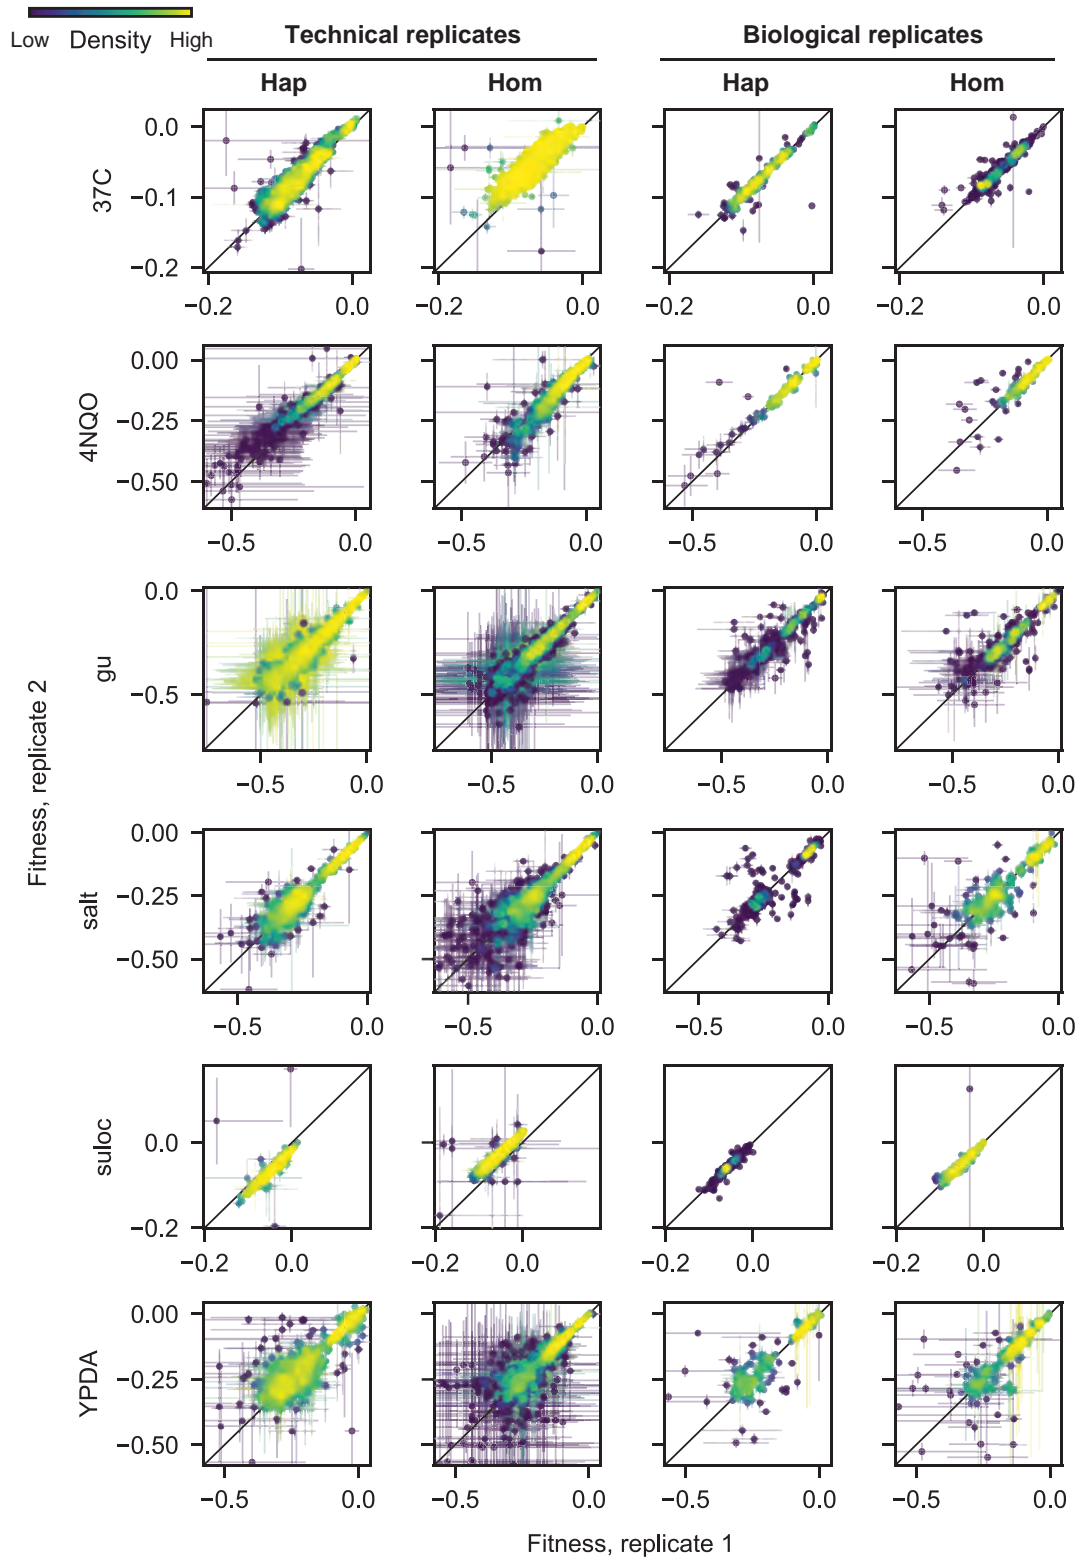


Figure S4: Comparison of technical and biological replicates for all ploidies and assay environments

468 5 Quantitative analysis of epistasis

469 5.1 Estimation of parameters

470 We model fitness ϕ as a function of the the underlying genotype which can be expressed as a sum
471 of combinations of N biallelic loci x_1, x_2, \dots, x_N that take on values $x_i = \pm 1$.

$$\phi = \bar{\phi} + \sum_i s_i x_i + \sum_{i>j} s_{ij} x_i x_j + \sum_{i>j>k} s_{ijk} x_i x_j x_k + \dots \quad (2)$$

472 This modeling framework casts additive first-order terms as the background-average effect of the
473 mutation, which is distinct from the effect of the mutation on some arbitrary wild-type genotype.
474 The terms s represent half the fitness difference between groups of individuals with and without the
475 mutation, or alternatively the expected deviation from the mean, positive or negative, for groups
476 with or without the mutation respectively. Pairwise epistatic effects are the background-average
477 perturbation that can be fit beyond the additive first-order term, and higher order epistatic terms
478 are similarly modeled. This view offers several advantages: 1) if one decides to choose a particular
479 genotype as the “wild-type”, only the signs of the terms need to change; 2) each coefficient is
480 estimated by partitioning half of the genotypes (each coefficient corresponds to a distinct slice
481 of the data), meaning each coefficient is equally powered; and 3) the coefficients are in principle
482 orthogonal from each other (there is no expected collinearity between the genotypic values of any
483 pair or combination of coefficients). This means that there is no “order” of coefficient fitting (one
484 does not have to fit the additive terms first), and fitting one coefficient does not influence another.

485 Coefficients from the equation above can always be estimated by least-squares regression when
486 all 2^N genotypes have a phenotypic measurement, though we note that we have in practice on aver-
487 age more than 1 phenotypic measurement per genotypes due to our biological replicates. However,
488 we may expect this formula to be sparser: not all mutations should have an effect, and not all pairs
489 of mutations should have a pairwise epistatic effect. We can regularize the estimation procedure to
490 yield a sparse subset using the LASSO procedure, which penalizes the least-squares regression by
491 the sum of absolute magnitudes of coefficients:

$$\min_s \left\{ \|\phi - \hat{\phi}\|_2^2 + \lambda \|s\|_1 \right\}. \quad (3)$$

492 In the absence of collinearity (as stated above, our formulation has no collinearity between
493 parameters), the LASSO operation is known to be consistent and asymptotically selects the correct
494 subset of non-zero parameters [57]. Sparsity is controlled by the λ parameter, which can be found
495 by cross-validation (in our case, 5-fold cross-validation was performed to reduce the extent of
496 overfitting). This approach removes coefficients that are approximately the same scale as the
497 noise. To provide 95% confidence intervals on the LASSO estimates, we performed 500 bootstrap
498 resampling with replacement of the data followed by model selection.

499 As discussed previously in Section 2.1, we identified extra mutations present in multiple strains
500 (FRS1, SPT7, HSL7). Because the SPT7 and HSL7 mutations likely occurred during the mating
501 process (Section 1.5), they may lead to specific signals of epistasis if they themselves have an effect.
502 We briefly assessed this possibility by plotting the distribution of fitnesses for individuals with and
503 without the mutation (constraining on the backgrounds in which the mutations were identified).
504 In visually examining these plots, we were unable to find evidence of a systematic effect for these
505 mutations. Therefore, these mutations were removed from consideration before building the model
506 by LASSO. On the other hand, FRS1 was likely present in one of the original parents of the
507 experiment and thus was found in approximately 50% of final strains. Though we did identify a

508 possible effect for this mutation in some environments, because it is not systematically distributed
 509 across the library, it is only expected to affect one of the higher order epistatic terms. (We cannot
 510 distinguish the effect of the epistatic term for the combination of strains that have FRS1 mutated
 511 and the effect of the FRS1 mutation). However, note that we have produced strains in replicate.
 512 Thus, the effect of the FRS1 mutation is unlikely to be consistently found in the same strains, and
 513 its signal will therefore be unlikely to dominate the epistatic term.

514 In general, the broad-sense heritability captured by the model is very high as both biological
 515 and technical replicates show high correlation (see Fig. S4). Thus, correlation of fitness measure-
 516 ments between environments can reveal the similarities between model coefficients. If measurement
 517 noise was too great such that it would dilute the correlation coefficients, then comparison between
 518 the predicted fitnesses may provide a better picture of environmental similarities (given that the
 519 coefficients were adequately estimated).

520 5.2 Variance partitioning

521 The phenotypic variability in the dataset can be partitioned into various components to quantify
 522 their relative importance. In our experiment, we are interested in not just the broad-sense heri-
 523 tability due to our focal loci (H^2 , or the variability due to genetic components), but also in the
 524 heritability due to specific additive and epistatic components. When the model coefficients are
 525 orthogonal, the phenotypic variance due to genetic components is trivially obtained by the sum of
 526 squares of each coefficient:

$$\sigma_{\text{gen}}^2 = \sum_i s_i^2 + \sum_{i>j} s_{ij}^2 + \sum_{i>j>k} s_{ijk}^2 + \dots \quad (4)$$

527 Partitioning the variance by subsets of coefficients – for example partitioning by first order
 528 terms or pairwise epistatic terms – is therefore straightforward.

$$\sigma_{\text{1st}}^2 = \frac{\sum_i s_i^2}{\sigma_{\text{gen}}^2} \quad (5)$$

$$\sigma_{\text{2nd}}^2 = \frac{\sum_{i>j} s_{ij}^2}{\sigma_{\text{gen}}^2} \quad (6)$$

529 However, we note that the coefficients are estimated from the data, and variance partitioning
 530 in this manner produces a bias. Removal of this bias is the major motivation behind mixed linear
 531 models that estimate narrow-sense heritability [58]. This caveat is not a major concern for our
 532 study, though, since extra sources of variation are either negligible (all the phenotypes are measured
 533 simultaneously in the same tube) or can be well estimated (measurement error can be estimated by
 534 replication). None of these extra sources of variation are expected to fundamentally alter only some
 535 of the coefficients or some subset of coefficients, and thus these relative partitions are expected to
 536 be unbiased.

537 6 Analysis of fitness-correlated trends

538 All epistatic interactions are ultimately the consequence of biophysical, physiological, or functional
 539 interactions, which depend on the specific details of the mutations involved. However, recent work
 540 has suggested that overall statistical patterns of epistasis follow regular and predictable fitness-
 541 mediated trends. In this section, we describe the framework we use to study these fitness-correlated

542 trends, and analyze the extent to which they can emerge as the consequence of specific idiosyncratic
 543 interactions, instead of from “global” effects involving non-specific fitness-mediated interactions
 544 among mutations.

545 6.1 Fitting regression slopes to determine fitness-correlated trends

546 Fitness-correlated trends (FCTs), such as diminishing returns or increasing costs, have often been
 547 analyzed by regressing the fitness effect of a mutation, $s = \Delta\phi = \phi_{mut} - \phi_{wt}$, against the fitness
 548 of the background in which it occurs, ϕ_{wt} . We refer to this as the $\Delta\phi$ formulation: we say that
 549 there is no FCT if $\Delta\phi$ is constant over a wide range of background fitness, while a negative rela-
 550 tionship between $\Delta\phi$ and ϕ_{wt} corresponds to diminishing returns/increasing costs (and a positive
 551 relationship corresponds to increasing returns/diminishing costs). However, care must be taken
 552 when performing this analysis, because when we regress $\Delta\phi$ against ϕ_{wt} , measurement errors in
 553 ϕ_{wt} will lead to a negative correlation even in the absence of true fitness-correlated trends [59].

554 A further complication with this formulation is that the regression slope we obtain depends in
 555 a complex way on the polarization we choose for the mutation (i.e., which allele is considered the
 556 wild-type and which is the mutant). To see this, consider the following simple linear model for $\Delta\phi$
 557 as a function of ϕ_{wt} :

$$\Delta\phi \equiv \phi_{mut} - \phi_{wt} = a_1 + b_1\phi_{wt}, \quad (7)$$

558 and the analogous model for the fitness effect of the reversion, $\Delta\tilde{\phi}$, as a function of ϕ_{mut} :

$$\Delta\tilde{\phi} \equiv \phi_{wt} - \phi_{mut} = a_2 + b_2\phi_{mut}. \quad (8)$$

559 Fitting data to these models using standard methods for ordinary least-squares, we find that the
 560 relationship between the regression slopes b_1 and b_2 is given by

$$b_2 = -\frac{b_1 + V}{1 + 2b_1 + V}, \quad (9)$$

561 where we have defined

$$V = \frac{\text{Var}[\Delta\phi]}{\text{Var}[\phi_{wt}]}. \quad (10)$$

562 We can use these equations to gain some intuition for the effect of V on the regression slopes
 563 and their reversions (i.e., a change in polarization). First, $V \geq 0$ by construction, and $V = 0$ only
 564 if there is no measurement error or no idiosyncratic epistasis, which in some extreme cases could be
 565 interpreted as measurement error for all measurements. As expected, it is only possible to lack an
 566 FCT in both polarizations ($b_1 = b_2 = 0$) if $V = 0$. Of note, the numerator of V can be decomposed
 567 to $\text{Var}[\phi_{mut}] + \text{Var}[\phi_{wt}] - 2\text{Cov}(\phi_{mut}, \phi_{wt})$, which shows that without a specific relationship between
 568 fitnesses of individuals with and without the mutation, $V > 0$, and an FCT will always emerge in
 569 at least one of the two polarizations.

570 Since in practice V is always positive, we can see that, as shown in Figure S5 and from Equa-
 571 tion 9, when $b_1 \geq 0$, then $b_2 < 0$, no matter the difference in scale of V and b_1 . Thus, in practice,
 572 increasing returns (or diminishing cost) epistasis or no FCT in one polarization of which allele is
 573 the “WT” always shows as diminishing returns (or increasing cost) in the reversion (when the allele
 574 is considered to be the “Mut” instead).

575 When $b_1 < 0$, or when there is diminishing returns in this polarization, then the behavior of
 576 b_2 depends on the scale of b_1 and V . First, some scenarios lead to $b_2 = 0$, or no FCT in the
 577 reversion, and these scenarios occur at the critical boundary where $V = -b_1$. Another critical
 578 boundary occurs where $V = -1 - 2b_1$, which leads to an asymptotic boundary where $b_2 \rightarrow \pm\infty$.




Figure S5: Relevant regimes for slopes and their reversions in the $\Delta\phi$ formulation. V is from Equation 10, b_1 and b_2 are least-squares regression slopes when an allele is labeled as the WT allele or the mutated allele (i.e., the reversion).

579 When $b_1 < 0$, only a small region between the critical boundaries leads to $b_2 > 0$ (the reversion
580 is increasing returns or diminishing cost epistasis). Outside the critical boundaries, $b_2 < 0$ and
581 therefore diminishing returns or increasing costs is found in both polarizations of the allele. Thus,
582 across the full space of possible parameters, diminishing returns and increasing costs – both of
583 which present as a negative regression slope – are more likely to emerge than positive regression
584 slopes in this $\Delta\phi$ formulation (though we note that biology may not explore this entire parameter
585 space uniformly), and slopes when mutations are reverted cannot always be anticipated intuitively.
586 We can also ask when $b_2 = b_1$: this will happen when $b_1 = -0.5V$. Because $V \geq 0$, this will only
587 happen when $b_1 < 0$ (and therefore $b_2 < 0$). Another fact from this equality is that if $b_2 = b_1$,
588 then the denominator of V has to be equal in the reversion. This means that $b_2 = b_1$ implies
589 $\text{Var}(\phi_{wt}) = \text{Var}(\phi_{mut})$.

590 Note, these complications are still present when using other regression techniques such as total
591 least squares that take into account measurement errors in ϕ_{wt} and ϕ_{mut} [41].

592 In contrast, we can resolve some of these complications by making two changes to the analysis:
593 (1) plotting ϕ_{mut} directly against ϕ_{wt} , and (2) regressing a linear relationship based on the total
594 least squares. Firstly, this approach avoids some problems with correlation in measurement errors.
595 In this formulation (i.e., the ϕ_{wt}/ϕ_{mut} formulation), measurement errors in both strains (or errors
596 in the dependent and independent variable) are taken into account [41] (we use the standard errors
597 estimated from Section 4.3), and we have the model functions:

$$\phi_{mut} = a_3 + b_3\phi_{wt} \tag{11}$$

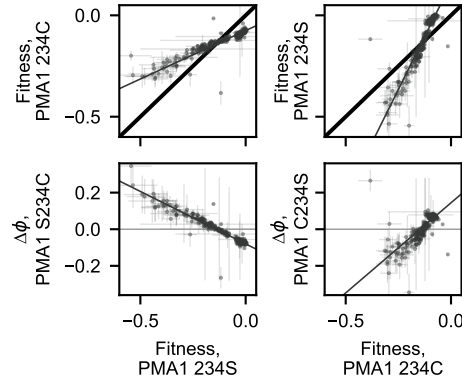


Figure S6: Comparison of fitness correlated trends for a simple case where the reversion of the focal mutation is straightforward. Haploid, 4NQO environment.

598 and the reversion:

$$\phi_{wt} = a_4 + b_4\phi_{mut} \quad (12)$$

599 Secondly, this framing and regression method (taking into account errors in both axes) also
 600 behaves far more intuitively: the slope in one direction is always the reciprocal of the other (i.e.,
 601 $b_3 = 1/b_4$).

602 To obtain some intuition of how to interpret FCTs in this ϕ_{wt}/ϕ_{mut} formulation, we can first
 603 attempt to interpret $b_3 = 1 = b_4$. This only occurs if $\text{Var}(\phi_{wt}) = \text{Var}(\phi_{mut})$, a property of the
 604 regression method. As described earlier, this is the regime where $b_1 = b_2 \leq 0$, and $b_1 = b_2 = 0$
 605 only if $\text{Var}(\phi_{mut} - \phi_{wt}) = 0$. Thus, a caveat of this ϕ_{wt}/ϕ_{mut} formulation is that a slope of 1 does
 606 not always indicate the absence of an FCT. In contrast, when $b_3 \neq 1$, then either $b_1 \neq 0$ or $b_2 \neq 0$.
 607 This can be shown by the fact that $b_3 \neq 1$ only when $\text{Var}(\phi_{wt}) \neq \text{Var}(\phi_{mut})$. This case necessarily
 608 implies $\text{Var}(\phi_{mut} - \phi_{wt}) \neq 0$, which is the necessary condition for $V > 0$.

609 We summarize these behaviors with some example figures from our data. First, we show an
 610 example of intuitive behavior, comparing both regressions and with mutational reversions (Fig-
 611 ure S6). In this simple example, regression of the fitness effect of the PMA1 234C mutation leads
 612 to a case of diminishing returns and increasing cost epistasis. When the mutation is “reverted,”
 613 or we regress the effect of the 234S mutation, we obtain the opposite FCT (diminishing costs, or
 614 increasing returns). These trends are also well-captured in the ϕ_{wt}/ϕ_{mut} formulation.

615 On the other hand, many examples are far less intuitive (Figure S7). In this example, regressing
 616 the effect of the WHI2 262L mutation leads to diminishing returns. However, regressing the effect of
 617 the reversion (262S) also leads to diminishing returns. In the ϕ_{wt}/ϕ_{mut} formulation, slopes behave
 618 as expected (the reversion is the reciprocal).

619 Examples where FCTs can only be interpreted in one of the mutational orientation are also
 620 found (Figure S8). In this example, the PMA1 234C mutation apparently shows no FCT, while its
 621 reversion displays increasing cost epistasis. On the other hand, the ϕ_{wt}/ϕ_{mut} formulation robustly
 622 shows a slope different from 1 and again behaves as the reciprocal when the mutation is reverted.
 623 Thus, because different slopes in this formulation do not readily yield an interpretation of the
 624 type of FCT (diminishing returns vs increasing returns), we refrain from using these plots for this
 625 purpose. Instead, we focus on this formulation’s ability to robustly identify FCTs when it exists.




Figure S7: Comparison of fitness correlated trends for a complicated case where the reversion of the focal mutation is not intuitive. Haploid, high-temperature environment.




Figure S8: Comparison of fitness correlated trends for a case where reversion may be interpreted as having no FCT. Haploid, acetic acid environment.

626 For our analysis, a final complication emerges from having biological replicates. In the parameter
627 estimation above (Section 5.1), this does not pose a problem (there is simply unexplained variation).
628 However, for the purpose of analyzing fitness correlated trends, if strains have two replicates for
629 the wild-type and two replicates for the mutant, then there are 4 possible comparisons and it is
630 no longer clear how to regress this effect of the mutation. To resolve this, we perform the analysis
631 on the average fitness of each genotype, which can be interpreted as the best estimate of the true
632 fitness of the genotype. The standard error of the average genotype fitness was computed as the
633 mean of the errors associated with the fitnesses that were averaged.

634 We have shown that, in general, if slopes different from 1 are obtained in the ϕ_{wt}/ϕ_{mut} formu-
635 lation, then we can interpret the data as displaying FCTs. However, what yields slopes different
636 from 1? If these formulations are readily interchangeable, then we may expect a single idiosyncratic
637 epistatic term involving the focal mutation, positive or negative, to be sufficient. However, we find
638 that this is not the case: in this formulation, we find that this epistatic interaction must also involve
639 a mutation with a non-zero additive effect.

640 To illustrate this, we begin with a simple schematic considering two loci A and B on top of
641 a background of other mutations with some fitness variance (Figure S9). We denote alternative
642 alleles at these loci as their letter case (A/a, and B/b), and the deviation from the mean fitness
643 between genotypes of alternative alleles for locus A as: $s_A = \phi_A - (\phi_A + \phi_a)/2$. When $s_A = 0$,
644 $s_B = 0$, and $s_{AB} = 0$, then plotting ϕ_A vs ϕ_a must yield a general “cloud” of points with a slope of 1
645 (Figure S9, top left panel). Partitioning the cloud of points by genotypes with the B and b alleles,
646 respectively, only yields two superimposed clouds (because the effect of having the mutation at
647 locus B, s_B , is zero). When $s_A = 0$, $s_B \neq 0$, and $s_{AB} = 0$, then the two clouds separate themselves
648 along the 1:1 line (Figure S9, top right panel). The regression slope for ϕ_A vs ϕ_a is still 1. The
649 case where $s_A = 0$, $s_B = 0$, and $s_{AB} \neq 0$ is more complicated. Setting $s_{AB} = E$, a constant, we
650 find the mean deviation in fitnesses $\phi_{AB} = E$, $\phi_{aB} = -E$, $\phi_{Ab} = -E$, and $\phi_{ab} = E$. If we focus
651 on plotting ϕ_{ab} against ϕ_{Ab} , we find that the negative deviation due to the epistatic coefficient for
652 ϕ_{Ab} moves the cloud to the left, while the positive deviation due to the epistatic coefficient for ϕ_{ab}
653 moves the cloud up. These coordinated movements yield a diagonal movement orthogonal to the
654 1:1 line. The same logic can be applied to plotting ϕ_{aB} against ϕ_{AB} , however in this case the cloud
655 moves to the right and down. Thus, the two clouds separate themselves in the direction of a slope
656 of -1 when an epistatic term is present (Figure S9, bottom left panel). The regression slope for
657 ϕ_A vs ϕ_a is still 1 even in this case and will eventually flip to be -1 as clouds separate themselves
658 farther and farther. Putting these orthogonal movements together, we find that the non-zero terms
659 for $s_A = 0$, $s_B \neq 0$, and $s_{AB} \neq 0$ lead to joint cloud movements (Figure S9, bottom right panel).
660 The regression slope for ϕ_A vs ϕ_a in this final case will never be one. Because these conditions
661 include the sufficient condition for FCTs in the $\Delta\phi$ formulation, our analyses on FCTs with this
662 ϕ_{wt}/ϕ_{mut} formulation are conservative, and we use this formulation for its advantages: 1) errors in
663 fitness measurements are taken into account for both ϕ_{wt} and ϕ_{mut} , 2) the slope for the mutation
664 reversion is the reciprocal, and 3) slopes different from 1 are always FCTs.

665 6.2 Decomposition of fitness-correlated trends

666 To understand whether idiosyncratic interactions lead to fitness-correlated trends, we proceeded
667 down two analytical avenues.

668 In the first, we examined the observed genotype fitnesses and removed epistatic terms one at
669 a time to see whether slopes converged to 1. Operationally, this involved first finding the global
670 linear regression line that fit the data best for a given locus in a given ploidy and environment.
671 We compared that regression to the best-fit line with slope of 1 by looking at the weighted sum of

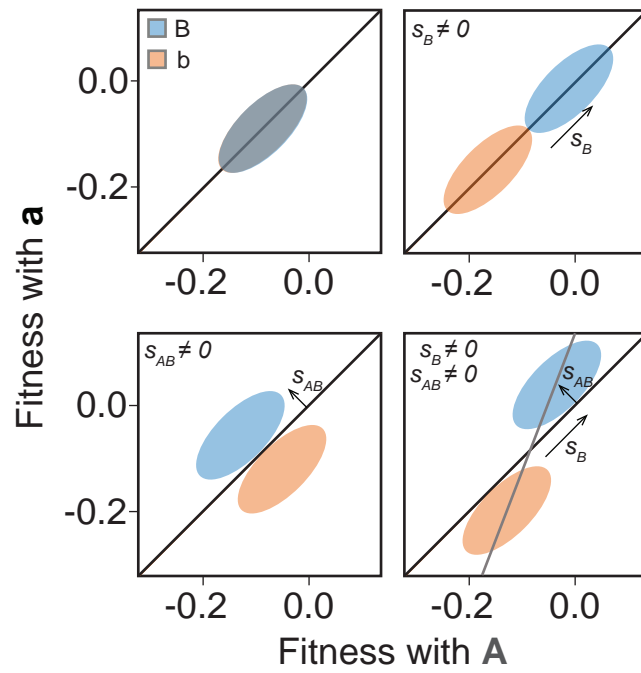


Figure S9: Effect of parameters on global regression in the ϕ_{wt}/ϕ_{mut} formulation. Clouds shown are for when $s_A = 0$. Top right shows the effect of $s_B \neq 0$, bottom left shows the effect of $s_{AB} \neq 0$ and bottom right shows the effect of both $s_b \neq 0$ and $s_{AB} \neq 0$.

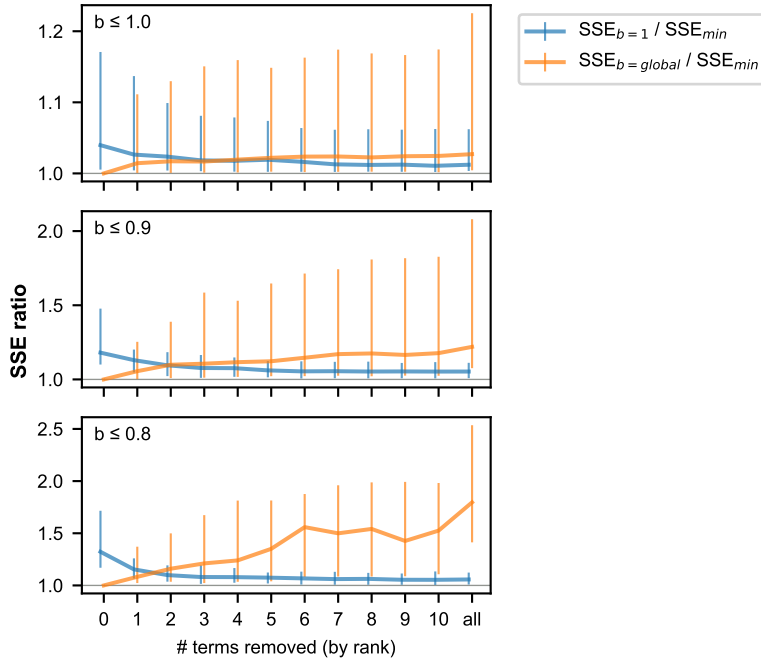


Figure S10: Three panels, one for each b threshold, showing the change in the ratios of the SSE for lines of slope $b = 1$ and $b = \text{global}$ as compared with the SSE for an unconstrained regression that minimizes SSE. Vertical bars indicate interquartile ranges.

672 squared errors (SSE), where a lower sum indicates a better fit to the data. After doing this, we
673 found the residual difference between the observed genotype fitnesses and the genotype fitnesses
674 as predicted by our full model of additive and epistatic terms. Then, we set the largest epistatic
675 term involving the focal locus to zero, regenerated the model fitness values, and added the residual
676 differences. To this dataset, we fit a line with the original global slope and a line with the slope 1,
677 again finding the SSE for each. We also fit a totally new regression line that minimized the SSE.
678 We then iterated this process, consecutively removing 10 epistatic terms and re-evaluating the fit
679 of the $b = 1$ and $b = \text{global}$ lines each time. Main text Figure 3E shows how the relative fit of these
680 two lines changes across ploidies, environments, and loci. Figure S10 shows how the SSE for $b = 1$
681 and $b = \text{global}$ compare to the minimized SSE as terms are progressively removed, revealing that
682 a slope of 1 tends to approach an idealized fit as terms are removed, while the global slope tends
683 to drift away. Figure S11 provides a more detailed look at how the ratio of SSEs for $b = 1$ and
684 $b = \text{global}$ change as terms are removed for each locus in each ploidy and environment.

685 In a converse analysis, we examined genotype fitnesses generated by our model of additive
686 and epistatic terms. For a focal locus in a given ploidy and environment, we first stripped away all
687 epistatic terms related to interactions between the focal locus and other loci, such that only additive
688 terms and interactions among background loci contributed to the modeled genotype fitnesses. This
689 produced a perfectly straight line with a slope of 1 and an intercept proportional to the background-
690 averaged additive effect of the focal mutation (as described in Section 5.1, this is twice the estimated
691 parameter s_i where i is the focal locus). We ranked the epistatic terms involving the focal mutation
692 by their effect size. Then, starting with the largest, we incorporated one term at a time into the
693 modeled genotype fitnesses. After each term was added, we replotted the fitness of genotypes with

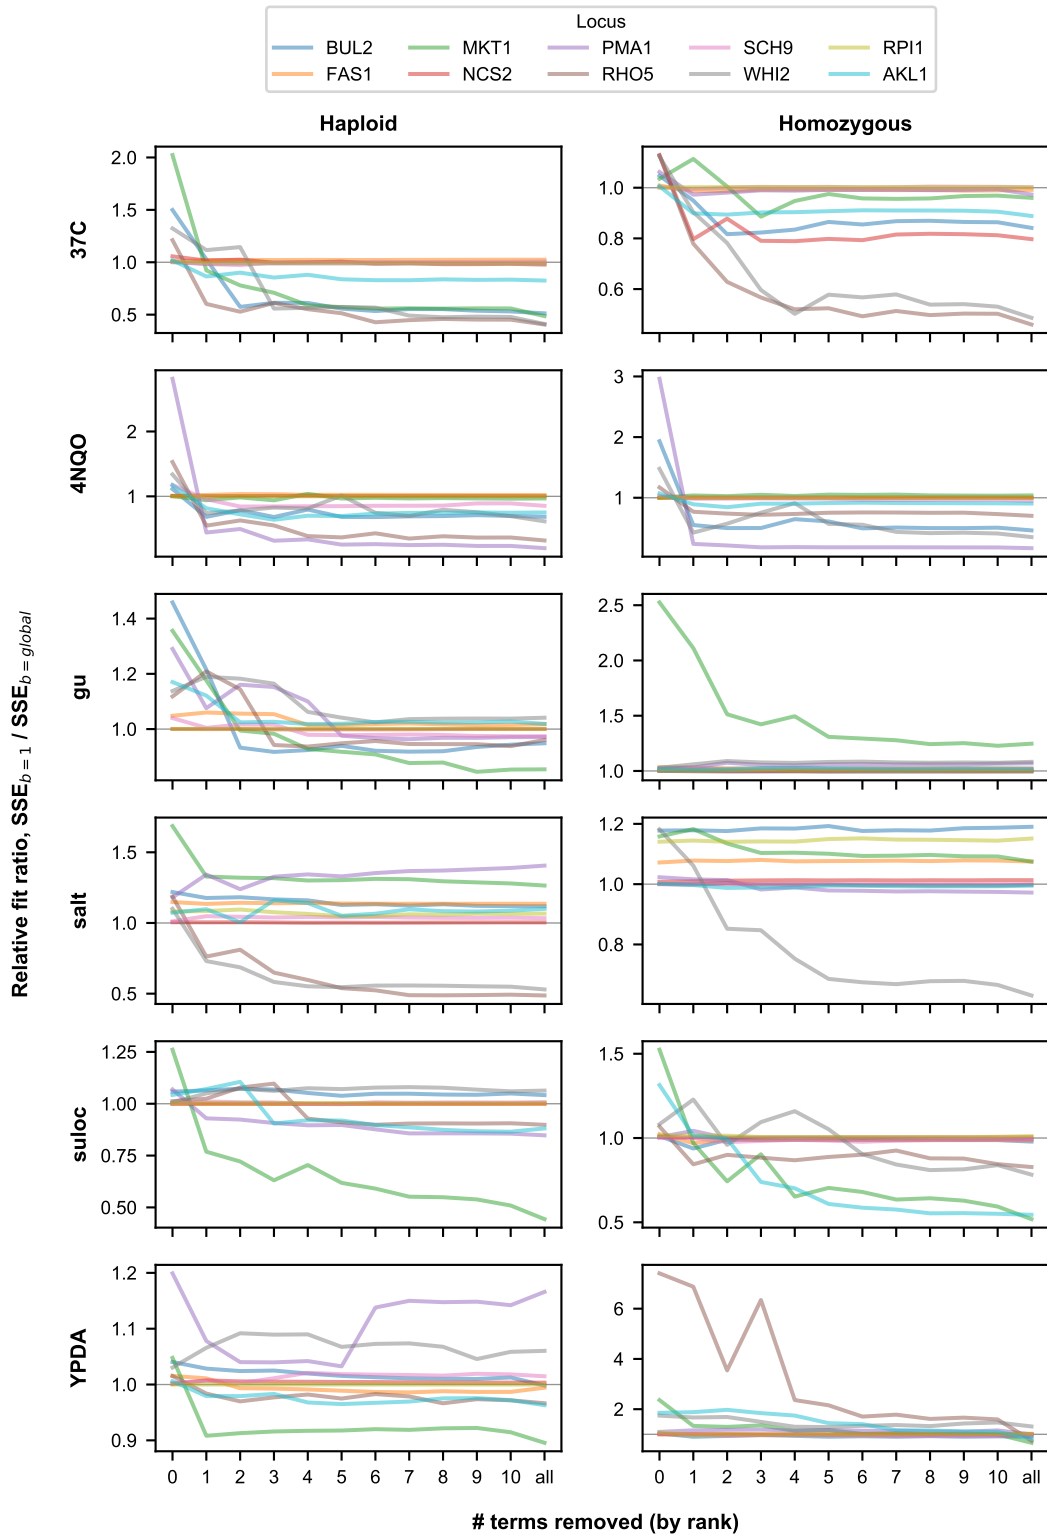


Figure S11: One panel for each ploidy and assay environment showing the change in the ratio of the SSEs for lines of slope $b=1$ and $b=global$ for each locus.

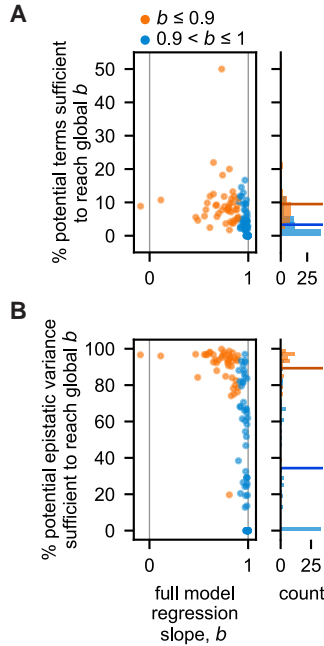


Figure S12: Scatterplot and histograms of regression slopes of FCTs for all data and the percentage of inferred epistatic (A) terms and (B) variance needed to recapitulate them. Horizontal colored lines in the histogram illustrate the mean.

694 and without a mutation at the focal locus and computed the regression slope.

695 We defined the number of terms sufficient to recapitulate the observed FCT as the number of
 696 added terms required to reach regression slope convergence within 0.01. More specifically, after
 697 adding each term, we asked whether the new regression slope differed from each of the previous
 698 three regression slopes by less than 0.01. If so, the number of terms required to reach that “plateau”
 699 was considered the number of terms sufficient to recapitulate the observed FCT. In a minority of
 700 cases, the final “plateau” slope differs from the full-model slope by greater than 0.01, but only in
 701 5 instances by greater than 0.02. Figure S12 presents the fraction of potential epistatic terms and
 702 potential epistatic variance sufficient to reach this plateau.

703 Note that, to permit more consistent comparisons, all loci were analyzed in the mutational
 704 direction that placed their regression slopes between -1 and 1. In other words, if plotting genotype
 705 fitness with A on the x axis and genotype fitness with a on the y axis gave a slope greater than 1,
 706 we would flip the axes such that the slopes would be equal to the reciprocal of the original slope
 707 (between 0 and 1).

708 Plots of ϕ_{wt}/ϕ_{mut} for all loci can be found in Figure S13 and Figure S14.

709 6.3 Quantifying the effect of landscape size in the analysis of fitness-correlated trends

710 The size of the fitness landscape we consider has two important effects on our ability to analyze the
 711 origins of fitness-correlated trends. First, as the number of mutations involved increases, the number
 712 of potential epistatic interactions between them increases exponentially. This creates more oppor-
 713 tunities for idiosyncratic interactions to exist and to produce apparent fitness-correlated trends.
 714 We note that this is an average effect: if we happened to choose precisely the set of mutations that

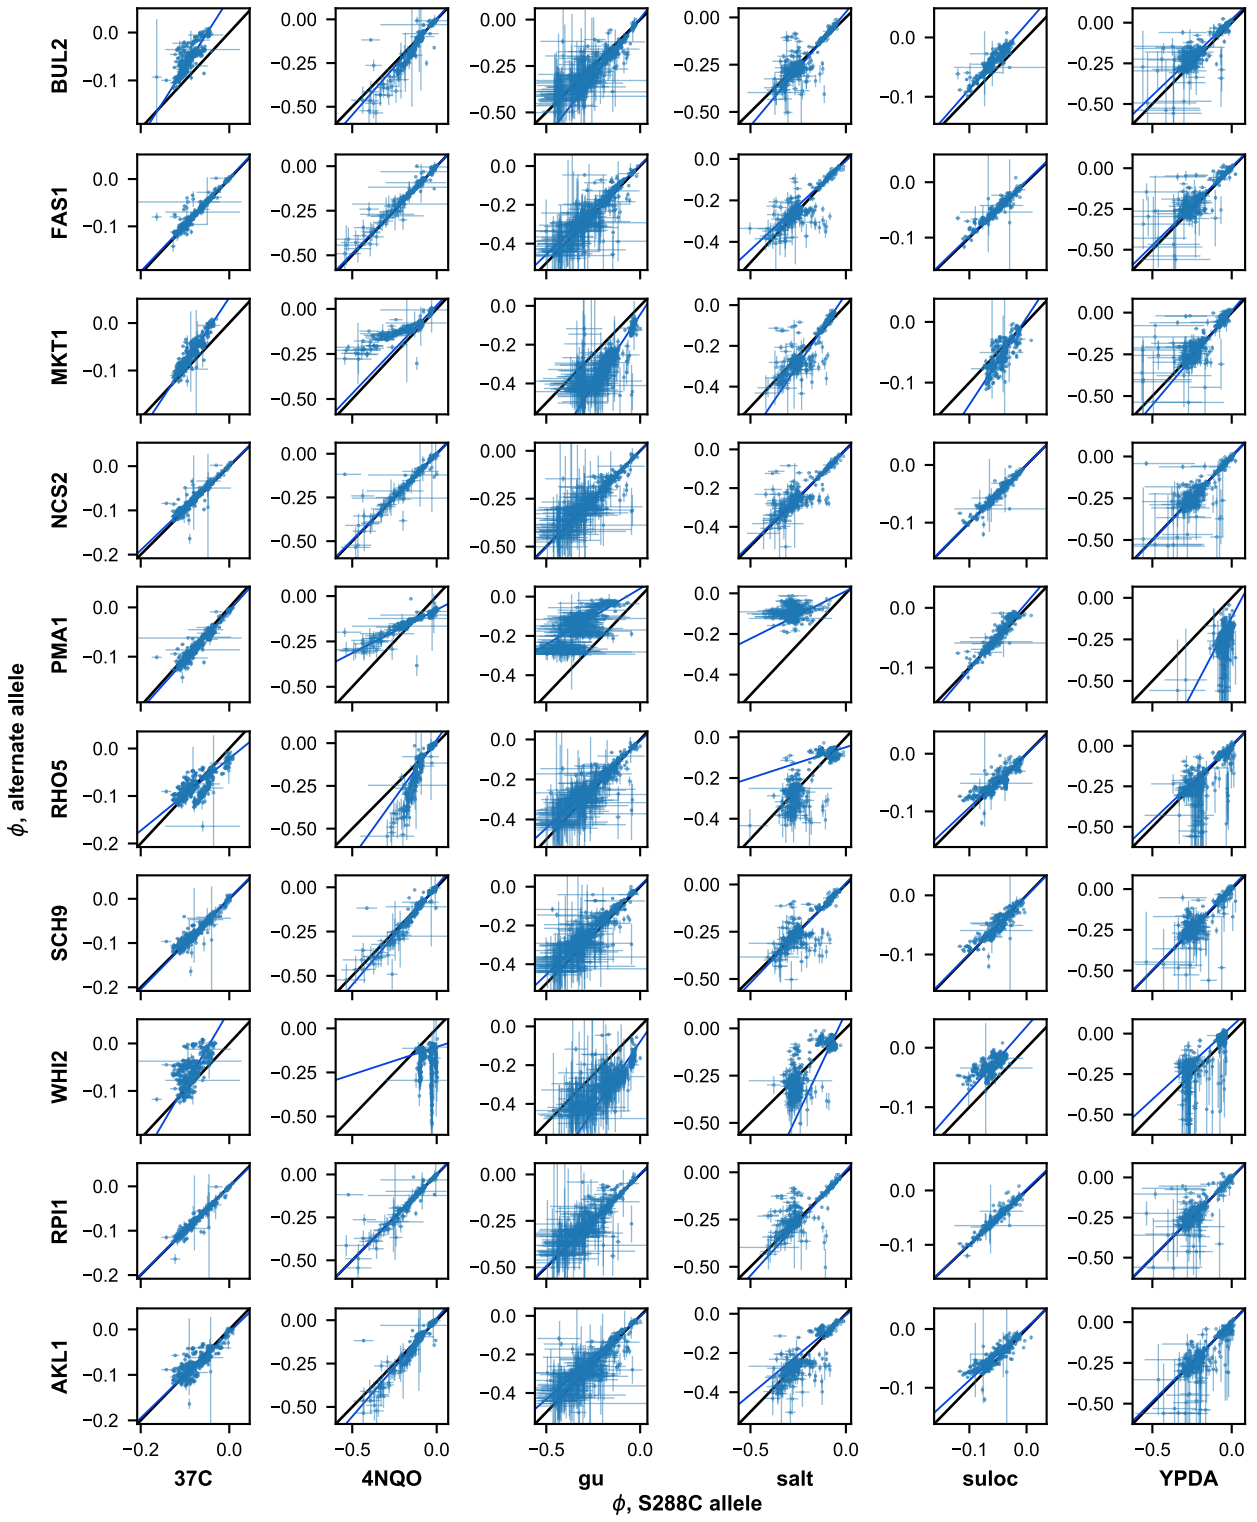


Figure S13: Scatterplots of ϕ_{wt}/ϕ_{mut} for all loci in haploid form.

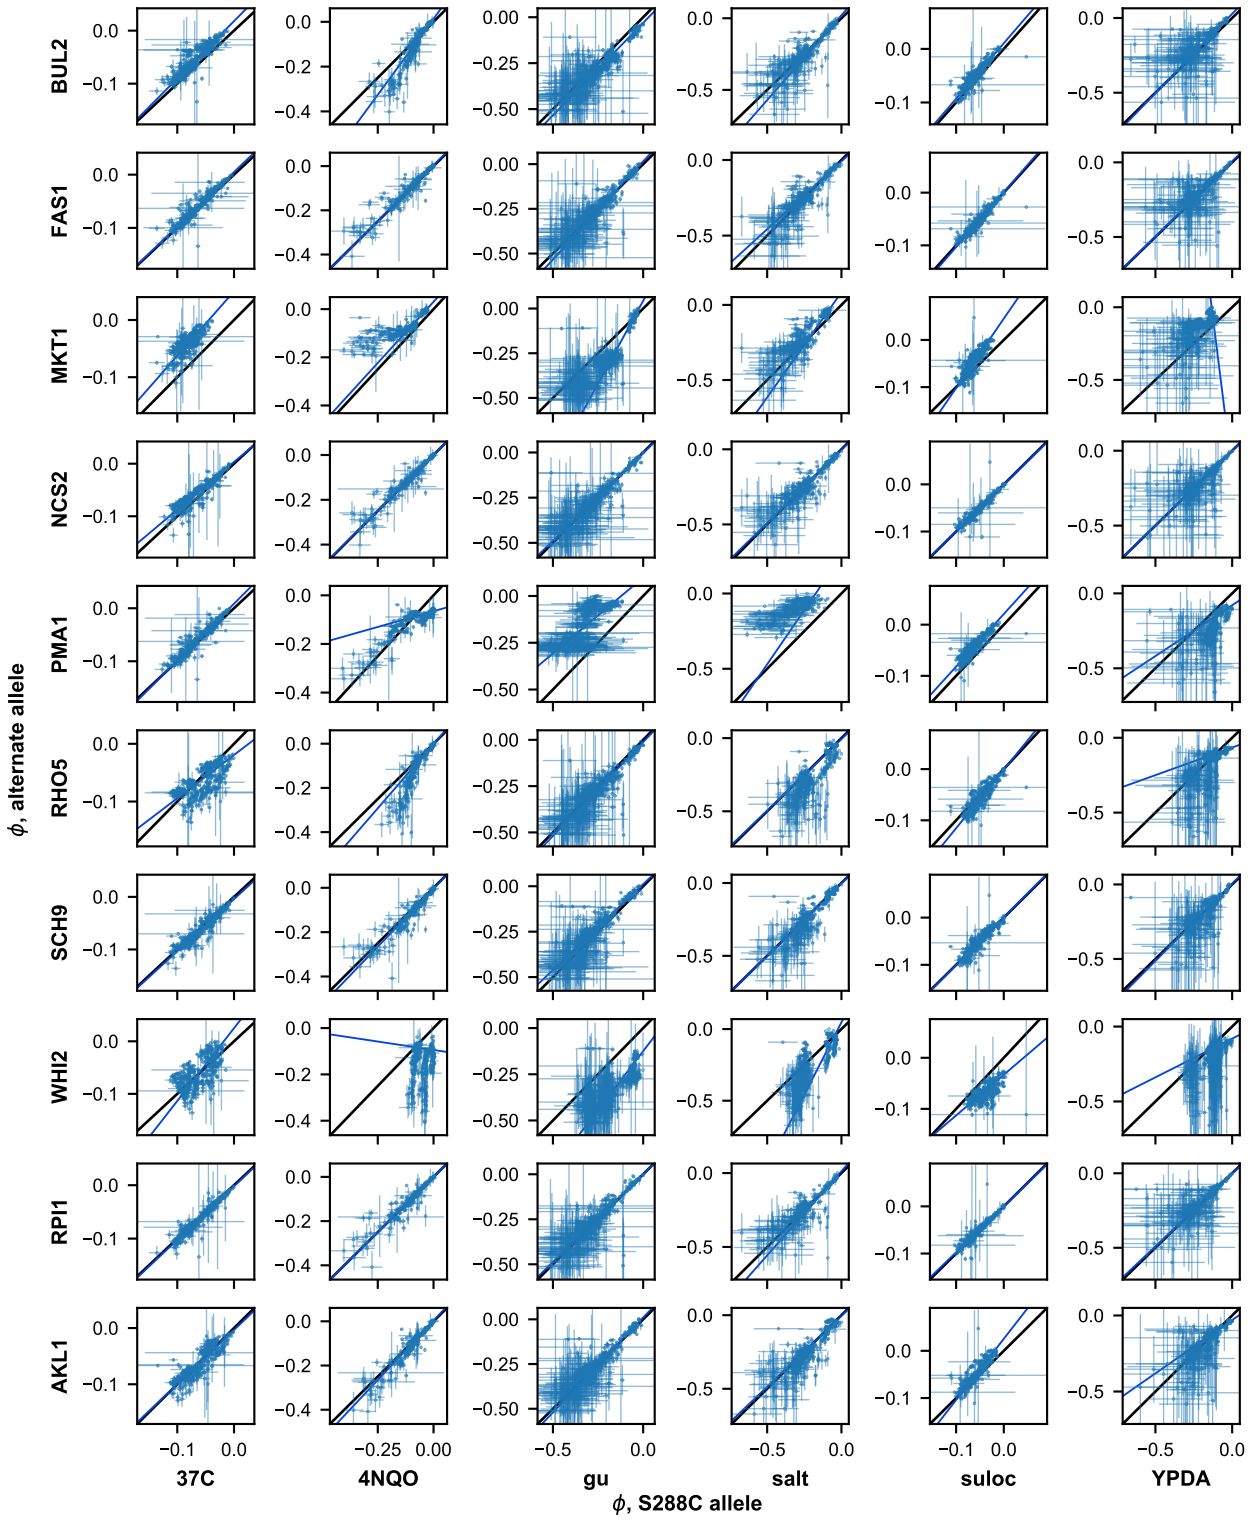


Figure S14: Scatterplots of ϕ_{wt}/ϕ_{mut} for all loci in homozygous form.

715 had the relevant idiosyncratic interactions, it may be possible to identify the relevant FCT in a
716 smaller landscape. In general, however, because theory argues that it is the accumulation of many
717 random idiosyncratic interactions that produces FCTs, we expect that larger landscapes become
718 more likely to reveal this effect. By random, we mean that idiosyncratic interactions do not obey
719 regular and predictable statistical patterns such as diminishing returns.

720 In addition to this, another key effect of landscape size is that the total number of genotypes,
721 and hence the total number of fitness measurements, also increases exponentially with the number
722 of mutations in the landscape. This reduces the influence of noise and improves our ability to
723 identify FCTs and the potential effects of idiosyncratic interactions in producing them. This is
724 critical, because linear regression analyses are known to be strongly affected by noise, which can
725 produce outliers: the variance on the slope estimate is (roughly) inversely proportional to the
726 number of data points used in the regression. Since increasing the number of loci considered in
727 fitness landscapes leads to an exponential increase in the total number of data points, we expect
728 that FCTs in significantly smaller landscapes (including landscapes like those examined in previous
729 studies) would therefore be more affected by noise.

730 To explore these effects of landscape size on the decomposition of fitness-correlated trends
731 (FCTs), we analyzed smaller sub-landscapes from the corresponding subsets of our data. By
732 definition, we cannot disentangle the potential role of idiosyncratic epistasis in creating an FCT
733 in a landscape consisting of only two loci. We therefore constructed landscapes with all possible
734 subsets of three or more of our mutations. For each subset, we analyzed the potential FCT using
735 our decomposition analysis (see Section 6.2). Specifically, for all subsets and all mutations that
736 had evidence of FCT in the full-dataset (i.e., $b \leq 0.9$), we computed the final ratio of sum-squared
737 errors (SSE) between a model with a slope of 1 (this is the idiosyncratic FCT model) and a model
738 with the global initial slope (the global FCT model), after removing all relevant epistatic terms.
739 The idiosyncratic model is supported when this final ratio is below 1. Note that we excluded from
740 this analysis subsets and mutations for which regressions were based on just 1 or 2 points.

741 We find that, at smaller subset sizes, there is a wide range of final relative fit ratios, indicating
742 that the same mutation can be found to display evidence for either the idiosyncratic model or
743 the global epistasis model driving FCTs. This spread of final SSE ratio can be explained by the
744 random effects of which mutations happen to be represented in each subset, as well as the increased
745 influence of noise on regression and on the inference of coefficients. However, we find that as the
746 subset size increases, the range narrows, with most relative fit ratios dropping below 1 (Figure S15
747 and Figure S16). This indicates that noise is particularly important in determining whether we can
748 distinguish between the idiosyncratic epistasis model and the global epistasis model, with smaller
749 subsets containing exponentially fewer points and hence far fewer measurements of the fitness effect
750 of mutations (or epistatic terms) with which to perform inference and regression. For our data,
751 with sparse interactions, a landscape of size greater than 8 appears sufficient to provide strong
752 support for the idiosyncratic model (Figure S15 and Figure S16).

753 To further confirm that noise is the primary driver of evidence towards the global epistasis
754 model (i.e., toward a relative fit ratio > 1), we investigated cases where the final relative fit ratio
755 remained above 1 even in our largest fitness landscapes. We found that these have a strong tendency
756 to be mutations in environmental/ploidy combinations with the greatest evidence for noise as
757 determined by the correlation between biological replicates (Figure S17). This suggests that these
758 outstanding cases pointing to global epistasis would be resolved toward the idiosyncratic epistasis
759 explanation with better measurements or with still larger landscapes. We also note that this finding
760 suggests that apparent differences between environments (e.g. with salt and YPDA environments
761 suggesting a larger role for global effects) may simply be an artifact of the inherently noisier
762 fitness measurements in these conditions. These lines of analysis also suggest that previous studies

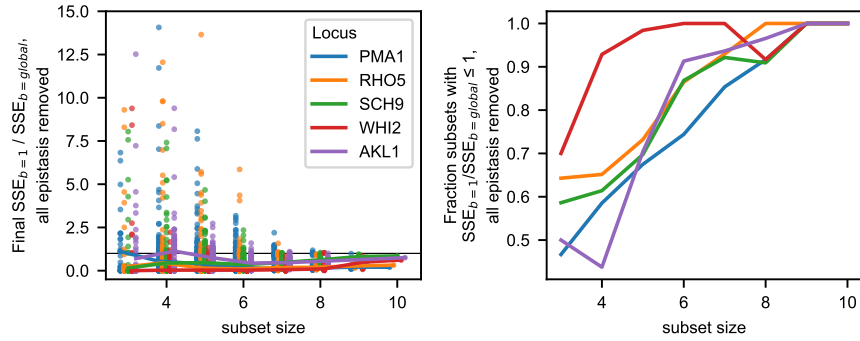


Figure S15: Effects of landscape size on the final SSE ratio (with values less than 1 indicating that FCTs are resolved in terms of idiosyncratic interactions) in 4NQO (haploid). In left panel, each point represents a subset of the full landscape of the corresponding size, with a particular focal mutation (indicated by the legend) having a fitness-correlated slope of $b \leq 0.9$ (polarity adopted such that b is ≤ 1). The relative fit (sum-squared error, SSE) ratio between regressions with fixed slope of $b=1$ and $b=\text{global}$ was computed after all epistatic terms were removed. At right, we show the fraction of subsets that have a final (all epistasis removed) relative fit ratio lower than 1 for each mutation, indicating support for the idiosyncratic model of fitness-correlated trends. Not shown are 16 points for which relative fit ratio is greater than 10. Lines show median ratios for each mutation.

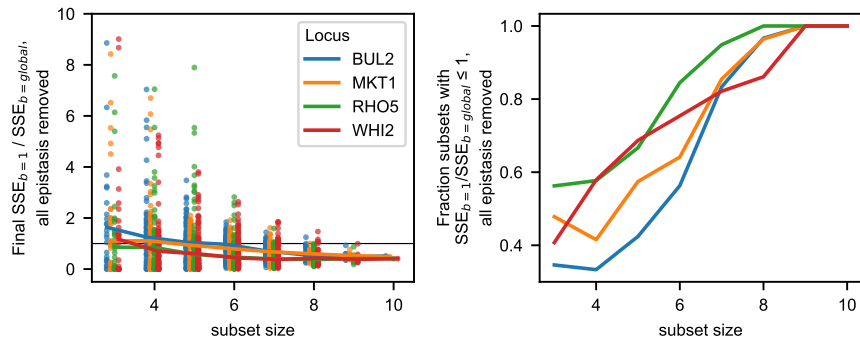


Figure S16: Effects of landscape size on the final SSE ratio (with values less than 1 indicating that FCTs are resolved in terms of idiosyncratic interactions) in 37C (haploid). In left panel, each point represents a subset of the full landscape of the corresponding size, with a particular focal mutation (indicated by the legend) having a fitness-correlated slope of $b \leq 0.9$ (polarity adopted such that b is ≤ 1). The relative fit (sum-squared error, SSE) ratio between regressions with fixed slope of $b=1$ and $b=\text{global}$ was computed after all epistatic terms were removed. At right, we show the fraction of subsets that have a final (all epistasis removed) relative fit ratio lower than 1 for each mutation, indicating support for the idiosyncratic model of fitness-correlated trends. Not shown are 15 points for which relative fit ratio is greater than 10. Lines show median ratios for each mutation.

763 with smaller landscape sizes might not have been able to decompose FCTs as being driven by
 764 idiosyncratic epistasis.




Figure S17: Final relative fit ratio as a function of reproducibility in biological replicates (i.e. the noise in individual fitness measurements). Each point represents the final sum-squared error ratio (i.e. the relative SSE ratio between regressions with a fixed slope of $b=1$ and $b=global$) for a given focal mutation (as indicated in legend) and environment (as indicated by arrows above, with haploids in black and homozygous diploids in blue). Note that SSE ratios greater than 1, which correspond to evidence for global epistasis, occur more frequently when the data is noisier. Only loci exhibiting an FCT in at least 3 of the 12 ploidy/environment combinations are presented.

7 Captions for Data Tables

7.1 Data Table S1

Primers used in genotyping, as well as search sequences used in parsing genotypes.

7.2 Data Table S2

Barcode to well to genotype map, and measured competitive fitness of each barcode in each ploidy and each environment.

The fitness values provided are joint inferred fitnesses from two technical replicates (two separate fitness assays were performed simultaneously), and the standard error of the estimate is obtained from the effect of an overdispersed binomial sampling error on this estimate (see Section 4.3 for more details). The estimated starting frequency of the barcode in the fitness assay in each technical replicate is also provided.

The HSL7-SPT7-FRS1 worksheet indicates whether each well was pure for one or the other allele, or considered impure at a stated threshold.

7.3 Data Table S3

Model parameters for each ploidy in each environment. We provide bootstrap 95% confidence intervals for the parameters as well.