

1 **Title:**

2 Extent and context dependence of pleiotropy revealed by high-throughput single-cell
3 phenotyping

4
5 **Author list:**

6 Kerry A Geiler-Samerotte^{a,b}

7 Shuang Li^a

8 Charalampos Lazaris^a

9 Austin Taylor^a

10 Naomi Ziv^a

11 Chelsea Ramjeawan^a

12 Annalise B Paaby^c

13 Mark L Siegal^a

14

15 **Author affiliations:**

16 ^a Center for Genomics and Systems Biology, Department of Biology, New York
17 University, New York, NY 10003

18 ^b Center for Mechanisms of Evolution, Biodesign Institutes, School of Life Sciences,
19 Arizona State University, Tempe, AZ 85287

20 ^c School of Biological Sciences, Georgia Institute of Technology, Atlanta, GA 30332

21

22

23 **Corresponding author:**

24 Kerry Geiler-Samerotte: kerry.samerotte@asu.edu

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43 **Abstract:**

44 Pleiotropy – when a single mutation affects multiple traits – is a controversial topic with
45 far-reaching implications. Pleiotropy plays a central role in debates about how complex
46 traits evolve and whether biological systems are modular or are organized such that every
47 gene has the potential to affect many traits. Pleiotropy is also critical to initiatives in
48 evolutionary medicine that seek to trap infectious microbes or tumors by selecting for
49 mutations that encourage growth in some conditions at the expense of others. Research in
50 these fields, and others, would benefit from understanding the extent to which pleiotropy
51 reflects inherent relationships among phenotypes that correlate no matter the perturbation
52 (vertical pleiotropy). Alternatively, pleiotropy may result from genetic changes that
53 impose correlations between otherwise independent traits (horizontal pleiotropy). We
54 distinguish these possibilities by using clonal populations of yeast cells to quantify the
55 inherent relationships between single-cell morphological features. Then, we demonstrate
56 how often these relationships underlie vertical pleiotropy and how often these
57 relationships are modified by genetic variants (QTL) acting via horizontal pleiotropy. Our
58 comprehensive screen measures thousands of pairwise trait correlations across hundreds
59 of thousands of yeast cells and reveals ample evidence of both vertical and horizontal
60 pleiotropy. Additionally, we observe that the correlations between traits can change with
61 the environment, genetic background and cell-cycle position. These changing
62 dependencies suggest a nuanced view of pleiotropy: biological systems demonstrate
63 limited pleiotropy in any given context, but across contexts (e.g., across diverse
64 environments and genetic backgrounds) each genetic change has the potential to
65 influence a larger number of traits. Our method suggests that exploiting pleiotropy for
66 applications in evolutionary medicine would benefit from focusing on traits with
67 correlations that are less dependent on context.

68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88

89 Introduction

90 Pleiotropy exists when a single mutation affects multiple traits (1,2). Often,
91 pleiotropy is defined instead as a single gene contributing to multiple traits, although
92 what is implied is the original definition — that a single change at the genetic level can
93 have multiple consequences at the phenotypic level (2). As our ability to survey the
94 influence of genotype on phenotype improves, examples of pleiotropy are growing (3-8).
95 For example, individual genetic variants have been associated with seemingly disparate
96 immune, neurological, and digestive symptoms in humans and mice (9,10). Genes
97 affecting rates of cell division across diverse environments and drug treatments have been
98 identified in microbes and cancers (11,12). A view emerging from genome-wide
99 association studies is that variation in complex traits is “omnigenic” in the sense that
100 many loci indirectly contribute to variation in many traits (13,14).

101
102 However, the extent of pleiotropy remains a major topic of debate because,
103 despite its apparent prevalence, pleiotropy is thought to be evolutionarily
104 disadvantageous. The more traits a mutation affects, the more likely it is that the mutation
105 will have a negative impact on at least one. Pervasive pleiotropy should therefore
106 constrain evolution (15), exacting what is known as a cost of complexity or cost of
107 pleiotropy (11,16-19). This cost may bias which mutations underlie adaptation, for
108 example, toward less-pleiotropic *cis*-regulatory changes over more-pleiotropic changes in
109 *trans*-acting factors (20,21), or toward changes to proteins that participate in relatively
110 few biological processes (22,23). Over long periods, the cost of pleiotropy may influence
111 the organization of biological systems, favoring a modular structure in which genetic
112 changes influencing one group of traits have minimal impact system-wide (24-29).

113 At stake in the ongoing debate about the extent of pleiotropy (30-33) are some of
114 modern biology’s prime objectives, including the prediction of complex phenotypes from
115 genotype data (18,34,35) and the prediction of how organisms will adapt to
116 environmental change (36,37). These predictions are more challenging if genetic changes
117 influence a large number of traits with complex interdependencies. Nonetheless,
118 understanding how a given mutation influences multiple traits could be powerful,
119 allowing prediction of some phenotypic responses given others (38,39). Indeed, recent
120 strategies in medicine called evolutionary traps aim to exploit pleiotropy, for example by
121 finding genetic changes that provide resistance to one treatment while promoting
122 susceptibility to another (40-42).

123 The lack of consensus about the extent of pleiotropy in natural systems is in part
124 due to poorly defined expectations for how to test for it experimentally. One key issue is
125 that defining a phenotype is not trivial (43,44). Consider a variant in the *apolipoprotein B*
126 gene that increases low-density lipoprotein (LDL) cholesterol levels as well as the risk of
127 heart disease. Elevated LDL promotes heart disease (45), so are these two phenotypes or
128 one? Alternatively, consider a mutation in the *phenylalanine hydroxylase* gene that
129 affects nervous system function and skin pigmentation. These dissimilar effects, both
130 symptoms of untreated phenylketonuria (PKU), originate from the same problem: a
131 deficiency in converting phenylalanine to tyrosine (46). Is it appropriate to call mutations
132 that have this single metabolic effect pleiotropic? Likewise, shall one call pleiotropic a
133 mutation that makes tomatoes both ripen uniformly and taste bad, when the effect of the

134 mutation is to reduce the function of a transcription factor that promotes chloroplast
135 development, which in turn necessarily affects both coloration and sugar accumulation
136 (47)?

137 The LDL, PKU and tomato cases are examples of vertical pleiotropy, *i.e.*
138 pleiotropy that results when one phenotype influences another or both are influenced by a
139 shared factor (5,43). The alternative to vertical pleiotropy is horizontal pleiotropy, in
140 which genetic differences induce correlations between otherwise independent
141 phenotypes. It might be tempting to discard vertical pleiotropy as less “genuine” (48) or
142 less important than horizontal pleiotropy, but that would be a mistake because vertical
143 pleiotropy reveals important information about the underlying biological systems that
144 produce the phenotypes in question. Consider the value in identifying yet-unknown
145 factors in heart disease by finding traits that correlate with it, or in understanding where
146 in a system an intervention is prone to produce undesirable side effects. Consider also
147 that the extent and nature of vertical pleiotropy speak directly to the question of
148 modularity: modularity is implied if vertical pleiotropy either is rare or manifests as small
149 groups of correlated traits that are isolated from other such groups. If there is modularity
150 then there can be horizontal pleiotropy, when particular genetic variants make links
151 between previously unconnected modules.

152 The above considerations suggest that a unified analysis that distinguishes and
153 compares horizontal and vertical pleiotropy is needed to make sense of the organization
154 and evolution of biological systems. However, existing methods of distinguishing
155 horizontal and vertical pleiotropy are problematic because judgments must be made about
156 which traits are independent from one another. Such judgments differ between
157 researchers and over time. Indeed, the tomato example can be viewed as a case of
158 horizontal pleiotropy transitioning recently to vertical pleiotropy as knowledge of the
159 underlying system advanced.

160 In this study, we propose and apply an empirical and analytical approach to
161 measuring pleiotropy that relies far less on subjective notions of what constitutes an
162 independent phenotype. The key principle is that the distinction between vertical and
163 horizontal pleiotropy lies in whether traits are correlated in the absence of genetic
164 variation (43). For vertical pleiotropy, the answer is yes: because one trait influences the
165 other or the two share an influence, non-genetic perturbations that alter one phenotype are
166 expected to alter the other. For horizontal pleiotropy, the answer is no: genetic variation
167 causes the trait correlation. In this study, we determined how traits correlate in the
168 absence of genetic variation by measuring single-cell traits in clonal populations of cells.

169 We used high-throughput morphometric analysis (49-53) of hundreds of
170 thousands of single cells of the budding yeast *Saccharomyces cerevisiae* to measure how
171 dozens of cell-morphology traits (thousands of pairs of traits) co-vary within clonal
172 populations and between such populations representing different genotypes. Within-
173 genotype correlations report on vertical pleiotropy, whereas between-genotype
174 correlations report on horizontal pleiotropy to the extent that they exceed the
175 corresponding within-genotype correlations. For one set of genotypes, we used 374
176 progeny of a cross of two natural isolates (54), which enabled not only the estimation of

177 vertical and horizontal pleiotropy but also the identification of quantitative trait loci
178 (QTL) with pleiotropic effects. For another set of genotypes, we used a collection of
179 mutation-accumulation lines, each of which contains a small number of unique
180 spontaneous mutations (55,56), which enabled a more direct test of the ability of
181 mutations to alter trait correlations.

182 The traits we study – morphological features of single cells – represent important
183 fitness-related traits (51,57,58) that contribute to processes such as cell division and
184 tissue invasion (*e.g.* cancer metastasis (59)). Cell-morphological features may correlate
185 across cells for a variety of vertical or horizontal reasons. Vertical reasons include: (1)
186 inherent geometric constraints (*e.g.* on cell circumference and area); (2) constraints
187 imposed by gene-regulatory networks (*e.g.* if the genes influencing a group of traits are
188 all under control of the same transcription factor); and (3) constraints induced by
189 developmental processes (*e.g.* as a yeast cell divides or “buds”, many morphological
190 features are affected). Horizontal pleiotropy might be evident because genetic variants
191 each affecting two or more traits (that are otherwise weakly correlated) are segregating in
192 the progeny of the cross between two natural isolates. Alternatively, horizontal pleiotropy
193 might be evident because a particular allele strengthens the trait correlation so that
194 genetic variation affecting one trait is more likely to affect another when that allele is
195 present. These alternatives can be distinguished by examining trait correlations in two
196 subsets of progeny strains defined by which natural isolate’s allele they possess at a QTL
197 of interest.

198 In addition to genetic variation, non-genetic variation may also alter the
199 correlations between traits. We rely on non-genetic heterogeneity within clonal
200 populations to serve as perturbations that reveal inherent trait correlations. However, the
201 correlations themselves might be heterogeneous within these populations. For example,
202 the dependencies between morphological features may change as cells divide. To control
203 for this possibility, we performed our trait mapping and subsequent analysis after binning
204 cells into three stages (unbudded, small-budded and large-budded cells). We further
205 examined whether trait correlations change across the cell cycle by using a machine-
206 learning approach to more finely bin the imaged cells into 48 stages of division.

207
208 Collectively, the results we present here demonstrate that both types of pleiotropy,
209 vertical and horizontal, are prevalent for single-cell morphological traits, suggesting that
210 biological systems occupy a middle ground between extreme modularity and extreme
211 interconnectedness. Perhaps more surprisingly, we find that trait correlations are often
212 context dependent, and can be altered by mutations as well as cell-cycle state and drug
213 treatments. The dynamic nature of trait correlations encourages caution when attempting
214 to quantify and interpret the extent of pleiotropy in nature or when making predictions
215 about correlated phenotypic responses to the same selection pressure, as is done when
216 crafting evolutionary traps. However, applying our approach may suggest which trait
217 correlations are less context dependent and therefore more useful in setting such traps.

218
219
220
221

222 **Results:**

223

224 ***QTLs with pleiotropic effects influence yeast single-cell morphology***

225 To detect genes with pleiotropic effects on cell morphology, we measured 167
226 single-cell morphological features (*e.g.* cell size, bud size, bud angle, distance from
227 nucleus to bud neck; **Table S1**) in 374 yeast strains that were generated in a previous
228 study from a mating between two wild yeast isolates (54,60). These wild isolates, one
229 obtained from soil near an oak tree, the other from a wine barrel, differ by 0.006 SNPs
230 per site (61) and have many heritable differences in single cell morphology (62). For
231 example, we find that yeast cells from the wine strain, on average, are smaller, are
232 rounder, and have larger nuclei during budding than yeast cells from the oak strain (**Fig**
233 **S1**).

234 To measure their morphologies, we harvested exponentially growing cells from
235 three replicate cultures of each of these 374 recombinant strains, and imaged on average
236 800 fixed, stained cells per strain using high-throughput microscopy in a 96-well plate
237 format (**Fig S2**). We used control strains present on each plate to correct for plate-to-plate
238 variation (see *Methods*), and quantified morphological features using CalMorph software
239 (53), which divides cells into three categories based on their progression through the cell
240 cycle (*i.e.* unbudded, small-budded, and large-budded cells) and measures phenotypes
241 specific to each category.

242 A simple way to measure pleiotropy would be to identify QTL that contribute to
243 variation in these phenotypes, and then to count the number of phenotypes to which each
244 QTL contributes. However, such a measure is sensitive to the statistical thresholds that
245 are used, and therefore risks yielding false inferences about trait modularity. Using a
246 liberal threshold would cause false-positive cases of pleiotropy (less apparent
247 modularity), whereas a conservative threshold would cause failures to detect pleiotropic
248 QTL when they exist (more apparent modularity). Such a measure also assumes the
249 counted traits are somehow independent except for correlations induced by genetic
250 variants. Statistically independent traits could be constructed (and then counted) as
251 principal components of the original traits, but the concern about too-liberal or too-
252 conservative QTL-detection thresholds would remain. Moreover, as we explore
253 extensively below, trait correlations are hierarchical (differing within and between
254 genotypes and conditions), making application of principal components analysis
255 problematic. For these reasons, we do not focus on counting the number of phenotypes
256 influenced by a given locus. Still, to begin to dissect vertical and horizontal pleiotropy we
257 must start with candidate examples of pleiotropic loci.

258 To detect QTL, we used 225 markers spread throughout the genome (54) and
259 Haley-Knott regression implemented in the R package R/qtl (63,64). We used a standard
260 permutation-based method to estimate statistical significance (63-65), with permutations
261 performed separately for each trait such that the per-trait probability of detecting a false
262 positive is 0.05. With this cutoff, we identified 41 QTL that contribute to variation in 155
263 of the surveyed morphological features (**Fig 1A**). This approach does not correct for the
264 testing of multiple traits. When we do so using a false-discovery rate set to 5%, results do
265 not change qualitatively. Indeed, the majority of QTL-trait associations that are
266 eliminated using this more-stringent cutoff involve QTL that are detected regardless of
267 this correction (**Fig 1A**; 80% of red points are present at QTL that also possess black

268 points). This observation suggests many of the associations detected with the less-
269 stringent cutoff are not spurious. We therefore present subsequent analyses using the
270 QTL-trait associations based on the less-stringent cutoff, but we also report analyses
271 using the reduced set of associations based on the more-stringent cutoff to establish that
272 qualitative conclusions did not change.

273 Most of the QTL we detect are pleiotropic, meaning each contributes to variation
274 in more than one morphological feature (**Fig 1A**; 36/41 QTL influence multiple traits,
275 20/26 after correcting for testing multiple traits). The median number of traits to which
276 each QTL contributes is 5 (5.5 after correction). This finding provides some support for
277 the idea that biological systems demonstrate limited pleiotropy, in that the median
278 number of traits affected per QTL is low (5/167). This median number of traits is similar
279 to that found in previous high-dimensional QTL screens (17,31), and in analyses of gene
280 knockouts in yeast and mouse (11). Further evidence that biological systems demonstrate
281 limited pleiotropy, and thus a modular organization, comes from previous studies of this
282 mapping family that show these same QTL do not contribute to variation in sporulation
283 efficiency (60).

284 However, as noted above, conclusions drawn about modularity from studies that
285 count traits are subject to criticism: some QTL influences may be too small to detect,
286 even with less-stringent significance thresholds, creating the appearance of modularity
287 even if every QTL influences every trait to at least some small degree. Further, we detect
288 some QTL that influence large numbers of traits, up to 73 (68 after correction). Although
289 such QTL mitigate to some extent concerns about detection power, they highlight other
290 potential problems: some QTL might contain multiple genetic differences that impact
291 different traits, and some morphological features might be inherently correlated and
292 therefore should not be counted as independent traits. Next, we focus on the first of these
293 problems by asking whether these QTL represent single genes that contribute to
294 phenotypic variation in many morphological features. We return to the issue of trait
295 independence after that.

296

297 ***Single genes with pleiotropic effects influence yeast single-cell morphology***

298 When a QTL affects multiple traits, it might not mean that variation in a single
299 gene is contributing to variation in these traits but instead that linked genes are
300 contributing to variation in distinct, individual traits. We partitioned genotype-phenotype
301 associations on the same chromosome into separate QTL when they were greater than 5
302 cM apart, except in genome regions where many genotype-phenotype associations are
303 present and there is no clear break point. A distance of 5 cM corresponds on average to 8
304 protein-coding genes in the yeast genome. The largest QTL we detected spans 17 cM,
305 which is roughly half the window size utilized in a previous study of this same QTL
306 mapping family (63). This approach reduces but does not eliminate the possibility that
307 QTL represent the action of linked loci.

308 For a small number of QTL with high pleiotropy (highlighted in **Fig 1A**), we
309 sought to test whether the effects on different morphological features were due to the
310 action of a single gene. We performed these tests by swapping the parental versions of
311 candidate genes (*i.e.* we genetically modified the wine strain to carry the oak version of a
312 given gene, and vice versa). We used the *delitto perfetto* technique to perform these
313 swaps (66), such that the only difference between a parental genome and the swapped

314 genome is the coding sequence of the single candidate gene plus up to 1 kb of flanking
315 sequence (see *Methods*). Candidate genes were selected based on descriptions of the
316 single-cell morphologies of their knockout mutants (67) and the presence of at least one
317 non-synonymous amino acid difference between the wine and oak alleles (62).

318 When a candidate gene contributes to the morphological differences between the
319 wine and oak parents, we expect yeast strains that differ at only that locus to recapitulate
320 some of the morphological differences between the wine and oak parents. Indeed, this is
321 what we observe for *PXLI*, a candidate for the QTL on chromosome 11, and *HOF1*, a
322 candidate for the QTL on chromosome 13 (**Fig 1B**; compare each plot on the right to the
323 leftmost plot; see also **Table S2**). This influence is most pervasive for *HOF1*; both the
324 oak and the wine alleles have a strong effect on the morphology of the opposite parent,
325 and their effects recapitulate the parental difference to a large extent. The pervasive
326 influence of *HOF1* on various morphological features is consistent with the fact that this
327 gene's product affects actin-cable organization and is involved in both polar cell growth
328 and cytokinesis (68). The effect of *PXLI* on cell morphology is also apparent across
329 many single-cell features, although only the oak allele has a strong effect that
330 recapitulates the parental difference. We evaluated *RASI*, a candidate for the QTL on
331 chromosome 15, but initial tests indicated that it did not have a significant impact on
332 most morphological features (**Table S2**). We also attempted to swap alleles for a
333 candidate gene corresponding to the QTL on chromosome 8, but were unsuccessful (see
334 *Methods*).

335 A previous screen for QTL influencing single-cell morphology in the progeny of a
336 genetically distinct pair of yeast strains (a different vineyard strain and a laboratory
337 strain) found some of the same pleiotropic QTL that we detect in the wine and oak cross
338 (69) (compare their Table 2 to our **Table S1**). In particular, we both find a QTL in the
339 same position on chromosome 15 that influences many morphological features related to
340 nucleus size, shape, and position in the cell (**Fig 1A**; orange). We also both detect a QTL
341 near base pair 100,000 on chromosome 8 that influences cell size and shape (**Fig 1A**;
342 pink). In the previous screen, the genetic basis of this QTL was shown to be a single
343 nucleotide change within the *GPA1* gene (69).

344 The main conclusion from our gene-swapping experiments, which is consistent
345 with the previous cell-morphology QTL study (69) as well as with comprehensive
346 surveys of how gene deletions affect the morphology of a laboratory yeast strain (11,49),
347 is that single genes with pleiotropic effects on cell morphology are readily detected in
348 budding yeast. Moreover, the morphological traits involved were previously shown to
349 influence fitness (51,57,58), which raises the question: why do so many genetic analyses
350 (including ours) detect pleiotropy (5,9-12,14) when other work suggests that pleiotropy
351 exacts a cost (17,18,20,21)?

352 353 ***Dissecting pleiotropy using clonal populations of cells***

354 One hypothesis to explain pervasive pleiotropy may be that the phenotypes we
355 chose to measure are not independent. Instead, many of these single-cell morphological
356 features may be inherently related such that perturbing one will have unavoidable
357 consequences on another and thus any associated limitation of adaptation will be
358 unavoidable as well. In other words, the hypothesis is that much of the pleiotropy we
359 observe is vertical pleiotropy. A test of this hypothesis is to ask whether traits that are

360 jointly affected by the same QTL are correlated in the absence of genetic differences. Our
361 dataset provides a unique opportunity to perform such a test because we quantified
362 single-cell traits for, on average, 800 clonal cells per yeast strain (**Fig S2**).

363 We leverage the hierarchical structure and large sample size of our dataset to
364 obtain precise estimates of the correlations that exist within and between strains. Thereby
365 we learn about the underlying relationships between morphological traits, which we use
366 to distinguish vertical from horizontal pleiotropy. Because we are studying clonal
367 families without a complicated pedigree structure, these within- and between-strain
368 correlations are equivalent to the so-called environmental and genetic correlations of
369 quantitative genetics (70). Here, we use a simple (and fast) method that is appropriate for
370 two-level hierarchical data to partition the total correlation into a pooled within-strain
371 component (r_w) and a between-strain component (r_B) (71). One caveat of this correlation-
372 partitioning approach is that r_B is effectively the correlation between strain means, which
373 can bias estimates of genetic covariance (70). This bias is most pronounced at small
374 sample sizes (70), so our large sample sizes allay concern. Nonetheless, for a subset of
375 traits, we tested whether estimates obtained from correlation partitioning are similar to
376 those obtained from mixed-effect linear models that specify the variance-covariance
377 structure of the experimental design. Environmental correlations estimated using both
378 methods are nearly identical (**Fig S3**). Genetic correlations estimated by correlation
379 partitioning are sometimes slightly smaller in magnitude than those obtained by linear
380 modeling (**Fig S3**). This bias is conservative; it may prevent us from identifying cases
381 where the environmental and genetic correlations significantly differ but will not tend to
382 create such cases. Despite this reduced power, we rely on the correlation-partitioning
383 approach, which is substantially faster, because our goal is to estimate environmental and
384 genetic correlations for thousands of trait pairs.

385 Unlike the mapping analysis, which considered phenotypes across all three
386 classes of cell type (unbudded, small-budded and large-budded), this correlation-
387 partitioning analysis can only be applied to pairs of phenotypes that can be measured in
388 the same cell. Three of the 36 pleiotropic QTL exclusively affect traits from different cell
389 types. For example, a QTL on chromosome 4 affects the shape of the nucleus in
390 unbudded cells as well as in large-budded cells. The correlation between these traits
391 cannot be partitioned into a within-strain component because these traits are never
392 measured in the same single cell. Excluding these three QTL leaves 33 pleiotropic QTL.

393 The 167 single-cell morphological features we measured represent 5645 pairs of
394 traits (378, 1081, and 4186 pairs of morphological features pertaining to unbudded,
395 small-budded, and large-budded cells respectively). Because some of these traits are
396 related, these thousands of trait pairs are not independent. This dependence prevents us
397 from reliably counting the absolute number of traits that are influenced by vertical vs.
398 horizontal pleiotropy. Still, partitioning correlations into a non-genetic (r_w) and a genetic
399 (r_B) component for thousands of trait pairs enables us to: 1) analyze a network describing
400 the degree to which morphological traits are interconnected or modular, and 2) detect
401 examples of horizontal pleiotropy if they exist. This approach differs from that of
402 previous studies of pleiotropy that used principal component analysis (PCA) to
403 understand which traits are correlated (17). Performing PCA on the individual-cell data is
404 not the same as controlling for r_w because PCA would ignore the strain groupings, which
405 can then dominate the analysis (the classic “heterogeneous subgroup” problem in

406 correlation analysis (72)). As a consequence, PCA can miss cases of horizontal pleiotropy
407 and obscure, rather than reveal, inherent trait relationships.

408

409 *Inherent relationships between traits contribute to pleiotropy*

410 We focus first on vertical pleiotropy by analyzing correlations that exist in the
411 absence of any genetic differences (r_W). In the analyses that follow, when we refer to r_W
412 (or r_B), we mean the magnitude of the correlation, as the sign has no relevance for
413 arbitrary pairs of traits. The distribution of r_W across traits that are influenced by the same
414 QTL reflects the degree to which that QTL acts via vertical pleiotropy. The overall
415 pattern of r_W values (*i.e.*, whether there are isolated clusters of highly correlated traits
416 versus a densely interconnected network of traits) reflects the modularity of the
417 underlying biological system. These within-strain correlations are estimated with
418 extremely high precision because of our large sample size of hundreds of thousands of
419 clonal cells (800 per each of 374 strains).

420 Most pairs of single-cell morphological traits are not strongly correlated across
421 clonal cells (**Fig 2A**). Median r_W is < 0.1 , and 74% of pairs have $r_W < 0.2$. Even if we
422 allow for nonlinear correlations by transforming data using a nonparametric model that
423 finds the fixed point of maximal correlation (73), r_W is less than 0.2 for roughly 65% of
424 pairs. These observations suggest that most of the morphological traits we surveyed are
425 not inherently related; *i.e.* for any individual cell, the value of one trait does not predict
426 well the values of most other traits.

427 Nonetheless, the distribution of r_W has a prominent right tail (**Fig 2A**) indicating
428 that some morphological features are strongly correlated across clonal cells. These
429 correlated features are more likely to be influenced by pleiotropic QTL. Among pairs
430 represented by this right tail (specifically, those with $r_W > 0.2$), 75% consist of traits that
431 share at least one QTL influence; the same is true for only 36% of pairs with $r_W < 0.2$.
432 These percentages are similar after changing our QTL detection threshold to correct for
433 having tested multiple phenotypes (66% and 21%, respectively). Further, the number of
434 pleiotropic QTL influencing both traits in a pair correlates with that pair's r_W (Pearson's r
435 is 0.52 before correction and 0.54 after). These results suggest that inherent correlations
436 among morphological features often cause genetic perturbations to one feature to have
437 consequences on another. In other words, we observe evidence of vertical pleiotropy.

438 Next, we studied each of the 33 pleiotropic QTL one at a time, asking whether
439 they influence pairs of traits with higher r_W than expected by chance. Most QTL have a
440 higher median r_W for the pairs of traits they influence than the median r_W given by all
441 possible pairs of traits (**Fig 2B**). This difference suggests that vertical pleiotropy drives a
442 large portion of the pleiotropy we detect.

443 We also used network analysis to move beyond the pairwise comparisons in **Fig**
444 **2A** and ask if morphological traits tend to be clustered into modules. Traits with higher
445 r_W do indeed tend to group into clusters in networks in which the single-cell
446 morphological traits are nodes and the r_W magnitudes are edge weights (**Fig 2C**). This
447 need not have been the case; pairs of traits with high r_W could have been distributed
448 throughout the network without necessarily being clustered near other high r_W pairs.
449 Instead, networks representing single-cell morphological features demonstrate more
450 clustering than do random networks drawn from the same values of r_W (**Fig 2D**; for
451 corresponding figures from unbudded and small-budded trait networks, see **Fig S4**). This

452 observation might indicate that morphological phenotypes have a modular organization,
453 whereby phenotypes within a module exert influence on one another, but exert less
454 influence on phenotypes from other modules. However, this observation could also result
455 from human bias when enumerating phenotypes that can be measured, in the sense that
456 phenotypes that bridge modules might somehow be absent from the data set. The
457 comprehensive nature of CalMorph diminishes this concern. A related concern is that
458 apparent modules are formed by trivially related phenotypes, such as the radius and
459 diameter of a circular object, but we do not find such trivial relationships among the
460 CalMorph phenotypes. Even a high correlation between the length and area of the
461 nucleus implies a constraint on nuclear aspect ratio.

462 Our analysis of within-strain trait correlations so far suggests that natural
463 variation contributing to variation in multiple single-cell morphological features often
464 acts via vertical pleiotropy. Still, there are hints of another mechanism at play. Some
465 QTL tend to influence traits that are not clustered in the correlation network (*e.g.* **Fig**
466 **2E**). And many pleiotropic QTL influence some pairs of traits with negligible r_W (**Fig**
467 **2B**). To investigate how often pleiotropy is not predicted by the degree to which
468 morphological features correlate in the absence of genetic variation, in the next section
469 we compare trait correlations present across clones (r_W) to those present between
470 genetically diverse strains (r_B).

471

472 ***Many traits are more strongly correlated across strains than they are across clones***

473 When genetic changes that perturb one trait have collateral effects on another, we
474 expect the way traits correlate across genetically diverse strains to reflect trait
475 correlations across clones (*i.e.* $r_B = r_W$). When this condition is met, pleiotropy can be
476 viewed as an expected consequence of inherent relationships between traits, *i.e.* vertical
477 pleiotropy. On the other hand, if a QTL influences two traits that do not correlate across
478 clones, it may cause these traits to correlate across strains in which this QTL is
479 segregating. In this case, we expect r_B will be greater than r_W , suggesting horizontal
480 pleiotropy.

481 After correcting for testing thousands of trait pairs, r_B significantly exceeds r_W in
482 25% of all trait pairs, and 43% of pairs in which at least one pleiotropic QTL influences
483 both traits (**Fig 3**; left panel; 43% of points are above the envelope, which represents a
484 Bonferroni-corrected significance threshold of $p < 0.01$). This percentage grows to 53%
485 in the smaller set of QTL that are detected after correcting for testing many traits. In the
486 majority of cases in which r_B significantly differs from r_W , r_B is greater than r_W (**Fig 3**;
487 left panel; 78% of points outside the envelope are above it). The magnitude of the
488 increase in r_B vs. r_W tends to scale with the number of pleiotropic QTL that jointly
489 influence both traits in a pair (**Fig 3**; left panel; colors get warmer farther above the
490 envelope). These observations are consistent with the hypothesis that QTL acting via
491 horizontal pleiotropy increase r_B relative to r_W .

492 However, horizontal pleiotropy is not the only reason traits may correlate
493 differently across strains versus across clones. We find significant deviations in r_B
494 relative to r_W in 14% of pairs for which no pleiotropic QTL influence both traits (**Fig 3**;
495 right panel), or 16% such pairs when using the smaller set of QTL that are detected after
496 correcting for testing many traits. This observation may suggest the presence of
497 pleiotropic genetic variants that we did not have statistical power to detect in our QTL

498 screen. But an alternate explanation for the observed increases in r_B over r_W is that
499 perhaps we sometimes underestimate r_W .

500 One reason r_W could be underestimated is that single-cell measurements are
501 noisier than group-level averages. To test this possibility, we randomly assigned
502 individual cells to groups (pseudo-strains) having the same number of cells as the actual
503 strains, and found that in these permuted data, r_B and r_W never significantly differ (**Fig 3**;
504 insets). Because detection of r_W was not underpowered relative to r_B , we conclude that
505 measurement noise does not meaningfully obscure r_W . Another reason r_W could be
506 underestimated is if trait correlations across strains are more linear than those across
507 clones. To test this possibility, for every pair of traits we transformed the single-cell trait
508 measurements using a nonparametric model that finds their maximal correlation (73).
509 This transformation abrogated significant differences in r_B relative to r_W for fewer than
510 5% of affected trait pairs. Another reason r_W might be less than r_B is if there tends to be
511 less phenotypic variation within strains than between strains. Contrary to this prediction,
512 every morphological trait we surveyed varies more within strains than between strains.
513 Indeed, broad-sense heritability of the traits did not exceed 15% (Fig S1B), reflecting that
514 within-strain phenotypic variation (*e.g.* variation in a cell's progress through the division
515 cycle) accounted for at least 85% of the total variation. A final reason r_W could be poorly
516 estimated is if non-genetic heterogeneity across different subpopulations within clonal
517 populations causes variation in r_W . Therefore, next we investigated whether the
518 relationship between single-cell features varies for clonal cells in different stages of the
519 cell-division cycle.

520

521 *Inferring a cell's progress through division from fixed cell images*

522 Pairs of traits for which r_B is strong whereas r_W is not should reflect horizontal
523 pleiotropy, but a closer examination of some of these pairs revealed traits that should
524 correlate due to simple geometric constraints. For example, cell size and the width of the
525 bud neck should correlate due to the constraint that, even at its maximum, bud neck width
526 cannot be larger than the diameter of the cell. When measured in small-budded cells,
527 these two traits are correlated across yeast strains ($r_B = 0.40$) but are significantly less
528 correlated across clones ($r_W = 0.15$). Given the simple geometric constraint coupling the
529 width of the bud neck to the cell's size, why is there a discrepancy between r_B and r_W ?
530 We reasoned that this discrepancy exists because the correlation between cell size and
531 neck width is disrupted during particular moments of cell division; *e.g.* the width of the
532 bud neck starts small even for large cells (**Fig 4A**; cell micrographs outlined in blue show
533 two cells in the progress of budding). If the relationship between morphological features
534 varies during cell division, r_W may represent a poor summary statistic.

535 How often does the relationship between morphological traits change during cell
536 division? Our single cell measurements are primed to address this question: we fixed
537 cells during exponential growth and imaged hundreds of thousands of single cells,
538 thereby capturing the full spectrum of morphologies as cells divide. A remaining
539 challenge is sorting these images according to progress through cell division, and then re-
540 measuring the correlation between morphological features within narrow windows along
541 that progression.

542 We performed this sorting using the Wishbone algorithm (74). This algorithm
543 extracts developmental trajectories from high-dimensional phenotype data (typically

544 single-cell transcriptome data). We applied Wishbone separately to cells belonging to
545 each of the three cell types defined by morphometric analysis (unbudded, small-budded,
546 and large-budded cells). The trends describing how morphological features vary across
547 Wishbone-defined cell-division trajectories are consistent with previous observations of
548 how morphology changes as yeast cells divide (75,76) (**Fig 4A**; line plots). For example,
549 Wishbone sorts fixed-cell images in such a way that cell area increases throughout the
550 course of cell division (**Fig4A**; upper left panel), and nuclear elongation occurs just
551 before nuclear division (**Fig4A**; lower left panel). These trajectories also match our own
552 observations of how morphological features change as live cells divide, which we tracked
553 by imaging at 1-minute intervals one of the 374 progeny strains that we had engineered to
554 express a fluorescently tagged nuclear protein (HTB2-GFP) (**Fig 4A**; micrographs). We
555 chose this particular strain because it does not deviate from the average morphology of all
556 374 recombinants by more than one standard deviation for any of the phenotypes we
557 measure.

558 To further validate Wishbone's performance, we asked whether it could
559 reconstruct the time series of live-cell images from the HTB2-GFP strain. We obtained
560 time series for 78 single dividing cells, each imaged over at least 20 timepoints.
561 Quantifying morphological phenotypes from live-cell images in a high-throughput
562 fashion proved difficult because the morphometric software was optimized for fixed-cell
563 images and as cells grow and bud, the cells and their nuclei can move out of the focal
564 plane. Also, although we used short exposure times when imaging GFP fluorescence,
565 there are concerns about photo-toxicity and associated growth and morphology defects
566 (77). For these reasons, we expect Wishbone to perform better on fixed-cell images than
567 on time series of live cell images. Still, Wishbone's cell-division trajectories recapitulate
568 the time course. When we align time series data across live cells by centering on each
569 cell's average predicted progress through division, Spearman's r is 0.65, 0.91, and 0.77
570 for time series corresponding to each of the three cell types (**Fig 4B**; see **Fig S5** for
571 recapitulation of 78 individual time series). These correlations are substantially higher
572 than those obtained by repeating the merging procedure after randomly permuting each
573 time series (corresponding Spearman's r of 0.42, 0.43. and 0.56). These observations
574 suggest that Wishbone is effective at properly assigning single-cell images to their
575 position in the cell cycle.

576 577 ***Cell cycle state can influence the relationship between morphological features***

578 To identify cases where significant differences in r_B vs. r_W might result because
579 r_W is sensitive to cell-cycle state, we first assigned each imaged yeast cell from the QTL-
580 mapping population to one of 16 equal-sized bins based on Wishbone's estimation of
581 how far that cell had progressed through division. Because we did this separately for each
582 of the previously defined cell stages (unbudded, small-budded, and large-budded), this
583 additional binning finely partitions cell division into 48 (16 x 3) stages. To hold genotype
584 representation constant across each of the 48 bins, we performed binning separately for
585 each of the 374 mapping-family strains, then merged like bins across strains. We then
586 performed correlation partitioning on each bin separately.

587 Binning cells by cell-cycle state typically decreased the amount of phenotypic
588 variation per bin, which we expect in turn to obscure the correlation between traits.
589 Consider an extreme example: if there is no phenotypic variation remaining for a given
590 trait, it cannot covary with any other traits. Indeed, for most pairs of traits, the binning

591 procedure either decreases r_W or does not have a dramatic effect on it; decreases in r_W are
592 especially evident for trait pairs where variation of at least one of the traits shows a
593 relatively large decrease upon binning (**Fig 4C**). However, for some pairs of traits,
594 despite the decrease in phenotypic variation for at least one trait, the correlation between
595 traits improves upon binning. For example, binning by cell division increases the
596 correlation between cell size and the width of the bud neck (**Fig 4D**; leftmost plot) such
597 that it approaches r_B . This increased correlation is consistent with our hypothesis that the
598 process of cell division was obscuring the dependency of bud neck width on cell size.
599 Examining more pairs of traits for which binning tends to increase r_W (**Fig 4C**; red,
600 orange, and yellow points) reveals additional cases where the process of cell division
601 decouples traits that are otherwise correlated, and where binning reveals the underlying
602 correlation (**Fig 4D**; leftmost three plots).

603 Despite the evidence that cell asynchrony alters some trait correlations, many
604 cases remain where heterogeneity in cell-cycle state does not explain the observed
605 discrepancy between r_W and r_B (**Fig 4D**; rightmost three plots). We previously
606 demonstrated that r_B significantly exceeds r_W in 24% of all trait pairs (1389/5645) (**Fig**
607 **3**). For almost half of these pairs (689 pairs), binning by cell division does not resolve the
608 discrepancy between r_B and r_W to any extent; in other words, r_W does not increase in any
609 of the 16 bins. For an additional 193 pairs, binning by cell division resolves the
610 discrepancy by at most 5% in any bin. These results imply that cell-cycle heterogeneity
611 does not cause the discrepancy between r_W and r_B in the majority of cases, suggesting
612 that the elevation of r_B over r_W could be explained by QTL demonstrating horizontal
613 pleiotropy.

614 ***Many QTL demonstrate horizontal pleiotropy***

616 To test horizontal pleiotropy further, we asked whether pleiotropic QTL cause
617 increases in r_B relative to r_W . Not all pleiotropic QTL affect pairs of traits for which r_B is
618 significantly greater than r_W , so we focus this analysis on the 27 out of 33 QTL that do.
619 We divided our yeast strains into sets in which a given QTL is not segregating, then re-
620 measured the difference between r_B and r_W . More specifically, for each QTL, we split the
621 374 phenotyped yeast strains into two groups based on whether they inherited the wine or
622 the oak parent's allele at the genotyped marker closest to the estimated QTL location.
623 Then we repeated correlation partitioning on each subset of strains and compared the
624 results to those obtained from the complete set. For each QTL, we focused on trait pairs
625 in which: (1) both traits are affected by this QTL, and (2) r_B is significantly greater than
626 r_W . Across all such pairs, median r_B tends to decrease upon eliminating allelic variation at
627 the marker nearest the QTL (**Fig 5A**). No similar reduction in r_B is observed when we
628 focus on pairs of traits that are not affected by each QTL (**Fig 5A**) and no similar
629 reduction is observed in r_W (median reduction in r_W is 0.0001).

630 There appear to be two ways in which a QTL may affect r_B . In some cases,
631 eliminating genetic variation at the marker nearest a QTL decreases r_B in both resulting
632 subpopulations. Such cases are consistent with a straightforward scenario in which
633 horizontal pleiotropy results when a QTL that influences two or more traits (that are
634 otherwise weakly correlated) is segregating in a population (**Fig 5B**; top row shows that
635 the correlation is strongest in the mixed population where both oak and wine alleles are
636 segregating). In other cases, eliminating allelic variation at a QTL site decreases r_B in

637 only one of the two resulting subpopulations (*i.e.* the subpopulation possessing either the
638 oak or the wine allele). This observation demonstrates that horizontal pleiotropy can
639 emerge by virtue of a QTL allele strengthening a correlation between two traits so that
640 genetic variation affecting one trait is more likely to affect the other when that allele is
641 present (78,79) (**Fig 5B**; bottom row).

642 How many cases where r_B significantly exceeds r_W can be explained, to some
643 extent, by horizontal pleiotropy (*i.e.* a QTL increasing the between-genotype
644 correlation)? For every trait pair where r_B significantly exceeds r_W and at least one QTL
645 influences both traits in the pair (1108 pairs total), eliminating allelic variation at the
646 marker nearest at least one of the shared QTL causes r_B to decrease in one or both of the
647 resulting subpopulations (**Fig 5C**: solid black line in rightmost plot). About 60% of these
648 decreases affect both subpopulations (*e.g.* **Fig 5B**; top row) and 40% affect only one
649 subpopulation (*e.g.* **Fig 5B**; bottom row). These decreases in r_B appear to resolve the
650 discrepancies in r_B vs. r_W more often and to a greater extent than does accounting for
651 cell-cycle heterogeneity (**Fig 5C**; leftmost plot). Some QTL have larger impacts on r_B
652 than do others (**Fig 5C**). Eliminating allelic variation near a QTL on chromosome 13
653 decreases r_B in the largest number of traits pairs (658). Subtracting the influence of a
654 QTL on chromosome 15 decreases r_B to the greatest extent; the average decrease across
655 342 affected trait pairs is 0.07.

656 One caveat that remains is whether these examples of pleiotropy represent single
657 genetic changes that influence multiple traits, or the presence of multiple nearby genetic
658 variants within each QTL. Our finding that there are two types of horizontal pleiotropy
659 (**Fig 5B**), provides some insight. The first type of horizontal pleiotropy (**Fig 5B**; upper
660 row) may result from the presence of multiple genetic variants segregating together
661 because recombination has not broken them apart. However, the second type of
662 horizontal pleiotropy (**Fig 5B**; lower row) cannot be explained in the same way because a
663 strong trait correlation exists in the absence of allelic variation at that QTL (blue points).
664 Therefore, although multiple closely linked variants might underlie the difference in the
665 trait correlation between the oak and wine alleles of the QTL, they would act at the level
666 of the correlation rather than of individual traits.

667

668 ***Spontaneous mutations alter the relationships between morphological features***

669 Our finding that some QTL alleles appear to strengthen correlations between
670 otherwise weakly correlated traits (**Fig 5B**; lower panel) lends credence to the idea that
671 the relationships between phenotypes, and thus the extent of phenotypic modularity (or
672 integration), are mutable traits (80). This finding has implications for evolutionary
673 medicine, in particular evolutionary traps, *e.g.* strategies to contain microbial populations
674 by encouraging them to evolve resistance to one treatment so that they become
675 susceptible to another (40-42). These traps will fail if targeted correlations can be broken
676 by mutations. To test whether spontaneous mutations can alter trait correlations, we
677 analyzed the cell-morphology phenotypes of a collection of yeast mutation-accumulation
678 (MA) lines (55). These MA lines were derived from repeated passaging through
679 bottlenecks, which dramatically reduced the efficiency of selection and thereby allowed
680 retention of the natural spectrum of mutations irrespective of effect on fitness (56). We
681 previously imaged these lines in high throughput (>1000 clonal cells imaged per each of
682 94 lines) (51).

683 Because MA lines contain private mutations unique to each strain, they are not
684 amenable to QTL mapping and between-strain trait correlations have less meaning.
685 Instead, we focused on within-strain correlations, which we expected to be consistent
686 across strains because of the limited number of mutations distinguishing the strains (an
687 average of 4 single-nucleotide mutations per line (56)), except if a rare mutation does
688 indeed alter the correlation. To determine if such correlation-altering mutations exist, we
689 calculated within-strain correlations for each strain separately and asked, for each trait
690 pair, whether any strains had extreme correlations relative to the other strains. For most
691 trait pairs, the MA lines trait correlations did not vary much from each other or from that
692 of the ancestor strain (**Fig 6**). However, in several instances, we observed a trait-pair
693 correlation dramatically outside the range of the other trait pairs and more than four
694 standard deviations from the mean (**Fig 6A**). Some mutations appear to influence many
695 trait-trait relationships (mutations found in blue- and purple-colored strains in **Fig 6B &**
696 **C**), whereas others influence fewer (mutations found in magenta-colored strain in **Fig**
697 **6C**). Mutations that alter trait correlations are not necessarily the result of rare events
698 such as aneuploidies or copy number variants; all five strains highlighted in **Fig 6** do not
699 possess these types of mutations and instead possess at least one single nucleotide
700 mutation (56).

701 Given that in the small sampling of spontaneous mutations captured by the MA
702 strain collection, we found several that appear to alter the relationship between
703 morphological features, we think such mutations are common enough to merit further
704 consideration in evolutionary models. The mutations in the outlier lines provide candidate
705 correlation-altering mutations for future mechanistic studies as well.

706 *Different environments alter the relationships between morphological features*

707 We have used non-genetic heterogeneity within clonal populations to uncover
708 inherent trait correlations. One might consider achieving the same aim by using instead
709 the non-genetic perturbations represented by different environmental treatments.
710 However, our results suggest that trait correlations can be highly context dependent,
711 changing across cell cycle state (**Fig 4**) and genetic background (**Fig 5B lower panel &**
712 **Fig 6**). If trait correlations change across environments, then the intricacies of the
713 environment-specific effects would need to be incorporated into any inferences about
714 vertical and horizontal pleiotropy, adding a complicating dimension to the analysis.

715 To investigate the potential utility of across-environment trait correlations for
716 distinguishing horizontal from vertical pleiotropy, and to further explore the context
717 dependence of trait correlations, we analyzed trait correlations across a range of
718 concentrations of the Hsp90-inhibiting drug geldanamycin (GdA). We showed previously
719 that GdA affects cell morphology (51), so it presents an opportunity to analyze how
720 correlations among these traits vary across environments. We performed this analysis
721 using a subset of the yeast strains from our QTL mapping family (**Fig S6**), partitioning
722 trait correlations into a pooled within-strain component (r_w) and a between-strain
723 component (r_B).

724 GdA alters the correlations between morphological traits. The impact of GdA on
725 r_w increases with the concentration of GdA (**Fig 7**), suggesting that more extreme
726 environmental differences are more likely to result in changes in r_w . We conclude that
727 looking across diverse environments is not a good way to understand the inherent
728

729 relationships between traits that exist in a single environment. Indeed, previous studies of
730 pleiotropy have treated growth parameters in different environments as different traits
731 (81) rather than as a way to estimate inherent trait correlations.
732

733 **Discussion:**

734 Although evolutionary biologists and medical geneticists alike appreciate that
735 organismal traits can rarely be understood in isolation, the extent and implications of
736 pleiotropy have remained difficult to assess. A common approach to measuring
737 pleiotropy has been to count phenotypes influenced by individual genetic loci (18,34,35).
738 For example, the median number of skeletal traits affected per QTL in a mouse cross was
739 six (out of 70 traits measured); this small median fraction of traits suggests that variation
740 in skeletal morphology is modular (17,31). Of course, for a count of traits to be
741 meaningful the full trait list must be comprehensive, and correlations between traits must
742 be properly accounted for (18,34,35). We aimed for comprehensiveness in a very similar
743 way to the studies of mouse skeletal traits, by systematic phenotyping of a large number
744 of morphological traits. However, we addressed the need for a principled approach to
745 separating inherent trait correlations from those induced by genetic differences in a new
746 way: by extending the analysis to include within-genotype correlations and thereby
747 enabling an operational definition of the distinction between vertical and horizontal
748 pleiotropy.

749 Our comprehensive analysis of how thousands of trait pairs co-vary within and
750 between mapping strains yields an unprecedentedly quantitative and nuanced view of
751 pleiotropy. We found support for modularity, not only in the low median number of traits
752 affected per QTL (five out of 167), but also in the way that within-genotype correlations
753 grouped traits into relatively isolated clusters (**Fig 2**). We also found ample evidence of
754 horizontal pleiotropy layered on top of that modularity, with many cases of between-
755 genotype trait correlations that exceeded within-genotype correlations (**Fig 3**).

756 Our results do not speak directly to whether modularity results from selection
757 against pleiotropy in nature because we survey only two natural genetic backgrounds
758 (wine and oak). In other words, the presence of modularity is not necessarily evidence
759 that it is adaptive or that it is maintained by natural selection. Future work comparing MA
760 lines to a larger collection of natural isolates might help answer questions about the
761 extent to which selection purges pleiotropic mutations.

762 Our partitioning of between-strain (genetic) and within-strain (environmental)
763 correlations relates to another approach to understanding trait interdependencies, the
764 estimation of the so-called **G** matrix. This genetic variance-covariance matrix
765 summarizes the joint pattern of heritable variation in a population of the traits that
766 compose its rows and columns, and is central to understanding how trait correlations
767 constrain evolution. The **G** matrix arises in the multivariate breeder's equation, which
768 describes the responses to selection of correlated traits (82). If breeding is the goal, the
769 distinction between vertical and horizontal pleiotropy is not so important, because both
770 can impede selection. Indeed, any philosophical concern about what constitutes a
771 biologically meaningful trait is irrelevant to the breeder, who actually cares about
772 particular traits (*e.g.*, milk yield and fat content).

773 **G** matrices are not only relevant to breeders, but to evolutionary biologists as
774 well, and it is worthwhile to place our results into this context. A major evolutionary

775 question in the **G**-matrix literature is whether the **G** matrix itself can evolve. In other
776 words, do short-term responses to selection (as captured in the breeder's equation) predict
777 long-term responses or do constraints shift through time, perhaps in a way that facilitates
778 (or is part of) adaptation (83)? Our results with MA lines add to evidence that the **G**
779 matrix readily changes (84), in that individual mutations have major effects on particular
780 trait correlations (*e.g.* **Fig 6A**). Our QTL-mapping results also support this view, in that
781 some cases of horizontal pleiotropy appear to be caused by alleles that alter trait
782 correlations (*e.g.* **Fig 5B**; bottom panel).

783 Another prominent question in the **G**-matrix literature is the extent to which the **P**
784 matrix, which includes all sources of phenotypic variation and covariation, predicts the **G**
785 matrix, which only includes additive genetic effects (*i.e.*, those that respond to selection).
786 If **P** predicts **G** well, as proposed by Cheverud (85), then inference of selection responses
787 from patterns of trait covariation in a population would suffice when genetic analysis
788 would be difficult or costly. Our results do not speak directly to this question, because we
789 did not estimate **G** itself and instead estimated genetic correlations that include non-
790 additive effects. However, our results are informative from another angle, which is the
791 comparison of genetic and environmental correlations. As we showed (**Fig 3**), although
792 there are cases in which the environmental and genetic correlations have different signs,
793 the environmental correlations do tend to match the signs of the genetic correlations and
794 predict their magnitudes to some extent as well, consistent with similarity between **P** and
795 **G**. Future experiments using clones embedded in a more complicated crossing scheme
796 could properly partition **P** into **G**, **E**, and the non-additive genetic components, to address
797 Chevrud's conjecture (85) more directly. There are only a few reports of comparisons of
798 **E** matrices (86), but we encourage increased attention to the **E** matrix to understand
799 inherent trait correlations and to contextualize **G** in a way that diminishes concerns about
800 which traits are biologically meaningful and therefore merit status as the matrix's rows
801 and columns.

802 A major and unforeseen conclusion of our work is the extent to which context is
803 crucial. We have shown that trait correlations change through the cell-division cycle, in
804 different genetic backgrounds, and across a drug gradient. It is likely that
805 macroenvironmental differences alter trait correlations as well (87). These findings
806 provide insight as to how biological systems appear to be modular, as evolutionary theory
807 predicts (24-29), yet in other studies appear to be highly interconnected (13,14). Our
808 results suggest that biological systems are modular, but that these modules change across
809 contexts such that the potential phenotypic impacts of a genetic change can be extensive.

810 These results support the idea that predicting the phenotypic impact of a genetic
811 change requires a paradigm shift (88,89), away from merely mapping the relationships
812 between traits and toward unfurling the range of contexts across which those
813 relationships persist. Future work in this direction will not only advance understanding of
814 the evolution of complex traits but will have practical benefits. For example, our
815 approach demonstrates a potentially fruitful way to design evolutionary traps: studying
816 within-genotype correlations across contexts to identify particularly immutable
817 correlations between traits.

818

819 **Acknowledgments:**

820 We are grateful to Dmitri Petrov, Grant Kinsler and Michael Lynch for helpful
821 discussions. We also thank Barak Cohen and David Hall for providing strains used in this
822 study. This work was supported by National Institutes of Health grant R35GM118170 (to
823 MLS), National Institutes of Health fellowship F32GM103166 (to KGS), National
824 Institutes of Health grant R35GM133674 (to KGS), a New York University Graduate
825 School of Arts and Science Dean's Dissertation Fellowship (to SL), and National
826 Institutes of Health grant R35GM119744 (to ABP). The funders had no role in study
827 design, data collection and analysis, decision to publish, or preparation of the manuscript.
828

829 **Data Availability:**

830 All data presented in this study are available at Open Science Framework
831 (DOI 10.17605/OSF.IO/B7NY5).

832

833 **Materials and Methods:**

834 *Measuring the morphology of single yeast cells*

835 Recombinant yeast strains were generated from a cross between the oak parent (BC233:
836 *SPS2:EGFP:kanMX4/SPS2:EGFP:kanMX4*) and the wine parent (BC240:
837 *SPS2:EGFP:natMX4/SPS2:EGFP:natMX4*) then genotyped at 225 markers in a previous
838 study (54,60); each resulting recombinant strain is a homozygous diploid. We prepared
839 yeast cells from these strains for microscopy using published methods (50-52,90).
840 Briefly, yeast strains were grown in minimal media with 0.08% glucose in 96-well plates
841 (91), harvested during exponential phase, fixed in 4% paraformaldehyde, stained for cell-
842 surface manno-protein (with FITC-concanavalin A) and nuclear DNA (with DAPI),
843 sonicated, mounted on 96-well glass-bottom microscopy plates, and imaged with a Nikon
844 Eclipse TE-2000E epifluorescence automated microscope using a 40× objective and
845 appropriate fluorescence filters. Three biological replicate experiments were performed,
846 typically yielding a total of between 500 to 1,000 imaged cells per strain (**Fig S2**).

847

848 *Statistical analysis and processing of cell image data*

849 Cell image processing was performed similarly to previous studies (50-52,90). Imaged
850 cells were analyzed for quantitative morphological traits using the CalMorph software
851 package (53), which reports on hundreds of morphological features that are each specific
852 to one of three cell types: unbudded, small-budded, and large-budded cells. We excluded
853 phenotypes for which >10% of cells had missing values, leaving 167 morphological
854 features. Any cell that was not scored for all features pertaining to its type was
855 eliminated. Each morphological trait was transformed via a Box-Cox transformation of
856 the raw data with the value of lambda that makes the residuals of a linear regression of
857 phenotype on strain most normal using the EnvStats package in R (92). Internal controls
858 (several wells representing the wine and oak parents) were present on every 96-well plate
859 and were used to correct for effects on phenotypic variation that resulted from differences
860 among replicate experiments, such as differences in the brightness of the cell stain. We
861 calculated the mid-parent value for each phenotype on every plate, then calculated the
862 average mid-parent value across all plates. For each phenotype, we found the difference
863 between the plate-specific mid-parent value and the average mid-parent value across all
864 plates. Then we subtracted this difference from each plate for the corresponding
865 phenotype. After correction, any cell with a morphological feature that deviated from the

866 average by more than 5 standard deviations was then eliminated, as investigation of such
867 cells typically revealed these were CalMorph miscalls or cellular debris.

868

869 *QTL mapping*

870 QTL interval mapping was performed similarly to previous studies (63) using the R/qtl
871 package (64). We performed a QTL scan using the function “scanone”, which finds at
872 most one QTL per chromosome, followed by the “scantwo” function which allowed us to
873 identify potential second additive QTL per chromosome. The yeast strains, which are
874 homozygous diploids, were modeled as doubled haploids and QTL models were fit using
875 Haley-Knott regression. When comparing QTL across traits, QTL greater than 5 cM apart
876 on the same chromosome were counted as separate QTL. We estimate that a region of 5
877 cM contains on average 8 genes, since there are 6,746 genes in the yeast genome (67),
878 and the map length we calculated using R/qtl is 4076 cM. In some cases, we detected a
879 QTL in between two others on the same chromosome and within 5 cM of both. In these
880 cases, there were typically many QTL found within a narrow region without any gaps of
881 greater than 2 cM. We counted these as single QTL that affect many traits. Using this
882 method, the largest QTL we detect spans 17 cM. A summary of all significant QTL
883 effects, including their chromosomal locations in cM and which QTL on the same
884 chromosome we considered unique, is provided in **Table S1** (also see **Fig 1A**).

885

886 To determine significance thresholds, we employed the method of (65) as implemented in
887 R/qtl (64). QTL were assigned a p-value based on a trait-specific empirical distribution of
888 genome-wide LOD score maximums from 10,000 (1,000 for the two dimensional scan)
889 randomly permuted datasets. We used a p-value cutoff of 0.05 to determine significant
890 QTL. We also obtained a more stringent set of significant QTLs by correcting for the
891 testing of multiple traits by controlling the false discovery rate across phenotypes. For
892 each trait we took the position and p-value of the maximum LOD score on every
893 chromosome. We calculated qvalues for this set of loci using the R ‘qvalue’ package (93)
894 and used a 0.05 q-value threshold to call significant QTL.

895

896 *Candidate gene swaps*

897 All yeast transformations were performed using the lithium acetate (94) and *delitto*
898 *perfetto* (66) methods. For each candidate gene, the gene was first deleted from haploid
899 variants of both the wine and oak parental strains and replaced with a selectable marker,
900 the yeast gene encoding orotidine-5'-phosphate decarboxylase (*URA3*). Gene knockouts
901 were confirmed by growth on plates lacking uracil and DNA sequencing of the affected
902 region. Next, the *URA3* selectable marker was replaced with the other parent's version of
903 the candidate gene. These candidate gene ‘swaps’ were selected by growth on 5-
904 Fluoroorotic acid and confirmed by sequencing of the affected region. For each candidate
905 gene, we swapped a region containing the coding sequence plus 5 – 750 bp up and
906 downstream. We used the following regions of homology to define the boundaries of
907 each swapped segment:

908

909 ~300bp upstream of *PXLI*: TTATAATTGTGGTTTAGCGTTTCATAGTCGC

910 ~300bp downstream of *PXLI*: CCTTATTCTCTATTCTTAGGCTCCTGTTCC

911 ~5bp upstream of *HOF1*: GAAAGAATGAGCTACAGTTATGAAGCTTG

912 ~ 300bp downstream of *HOF1*: GTATTCGTAACAAGTGAAGTCTAATGATAT
913 ~ 750bp upstream of *RASI*: CGACTAAAGGAATTATACCATCATGCATC
914 ~ 300bp downstream of *RASI*: GCATTTCTAAAAACAGAGCTTTTGCCG

915

916 These regions of homology were chosen by searching for regions of higher GC content
917 nearby the start and end of each gene's coding sequence. In addition, we attempted to
918 swap the wine and oak parents' versions of the *GPAL* gene on chromosome 8. Despite
919 trying various regions of homology, we could not successfully replace *GPAL* with the
920 *URA3* selectable marker in the oak parent. *GPAL* is known to be essential in some genetic
921 backgrounds (95).

922

923 Though the recombinant strains we studied are homothallic diploids, the strains in **Fig 1B**
924 (both the parental strains and the strains possessing the gene swaps) are haploid. Because
925 the analyses in **Fig 1B** compare pairs of strains (*e.g.* the oak haploid parent to the wine
926 haploid parent, or the wine haploid parent to the wine haploid parent possessing the oak
927 allele of *PXLI*), we only considered experiments where both strains in the pair were
928 imaged in the same replicate experiment. To account for differences among replicate
929 experiments, for each phenotype, we subtracted the value in one strain from the value in
930 the other to calculate the phenotypic difference between strains in that replicate
931 experiment; the reported value is the average of these differences across replicate
932 experiments (**Table S2, Fig 1B**).

933

934 *Calculation of correlation coefficients*

935 We used WABA II as implemented in the multilevel package in R (71) to
936 calculate cell-level (r_W) and strain-level (r_B) Pearson correlation coefficients for each pair
937 of traits. We used an r-to-z transformation to determine whether differences in r_B vs. r_W
938 are significant at a Bonferroni corrected p-value of 0.01 (this is a z-score cutoff of 4.63,
939 given 5645 pairs of traits were tested). To assess whether correlations across single cells
940 generally result in different values than correlations across group-level averages, we
941 assigned yeast cells to groups (pseudo-strains) randomly, maintaining the same number
942 of cells per strain as in the actual data. To assess whether results would differ if we
943 allowed for non-linear correlations, we transformed the single-cell data using a
944 nonparametric model that finds the fixed point of maximal correlation, implemented in
945 the R package *acepack* (73). To assess whether results from WABA differed from those
946 obtained using a standard quantitative genetics model (**Fig S3**), we implemented the latter
947 using the *nlme* package in R (96) to specify a mixed-effects model with cells nested
948 within strains. We specified a covariance structure that allows covariance between two
949 traits but no covariance between cells or between strains. We used this model to calculate
950 the environmental and genetic correlations for 350 pairs of randomly chosen traits.

951

952 *Live imaging single cells as they divide*

953 For live imaging the morphology of dividing yeast cells, we chose one of the
954 recombinant yeast strains, F2_292. This strain was chosen because it does not deviate
955 from the average morphology of all 374 recombinants by more than one standard
956 deviation for any of the phenotypes we measured. F2_292 was transformed to express a
957 fusion protein of GFP and a nuclear protein (histone H2B encoded by *HTB2*). Two

958 independent transformants were imaged in the GFP channel (for nuclei) and in brightfield
959 (for cell outlines). We prepared live cells for imaging following published methods
960 (91,97,98), in a similar way to that described above, except cells were neither fixed nor
961 stained. Cells were taken during mid-log phase growth, seeded in 96-well glass bottom
962 microscopy plates containing minimal media with 0.08% glucose, and imaged over a
963 period of 3 hours. In each of four replicate experiments, cells were imaged either every
964 minute, every 90 seconds, or every 2 minutes. We used short exposure times (afforded by
965 the highly abundant HTB2-GFP) and took only a single image per well per timepoint to
966 reduce photo-toxicity. We processed images with CalMorph then matched cells across
967 timepoints by their centroid locations in the imaging fields. Overall we obtained time
968 series for 78 cells that each: (1) were longer than 20 timepoints, (2) contained no gaps
969 where the cell was not phenotyped for many consecutive timepoints, and (3) contained no
970 images that appeared to be very out of focus potentially resulting in misestimation of
971 phenotype values. Because CalMorph divides cells into unbudded, small-budded and
972 large-budded stages, these 78 time series are also divided this way (11, 23, and 44 cells,
973 respectively).

974
975 We used the Wishbone algorithm implemented in python (74) to estimate progression
976 through the cell-division cycle. Wishbone recapitulates each of these 78 time series (**Fig**
977 **S5**) with Spearman correlations between the actual and inferred image orders that average
978 0.42, 0.85, 0.40 across all unbudded, small-budded or large-budded series, respectively.
979 The lower correlations between Wishbone's predicted progress through division and time
980 for the unbudded and large-budded cells may result because each time series captured
981 only a part of the cell-division cycle and, during some stretches in the cycle, there are
982 fewer morphological changes taking place. To estimate Wishbone's accuracy across a
983 longer stretch of time, we merged the Wishbone predictions within the classes of
984 unbudded, small-budded or large-budded cell time series. To do so, we had to contend
985 with the fact that the first timepoint for each imaged cell often represents a different
986 moment in division. For example, some time series for unbudded cells start from an
987 image that is already far along the division process (**Fig S5**; values close to 1 on the
988 vertical axis) while others start from a cell image that has just begun its division cycle
989 (**Fig S5**; values close to zero on the vertical axis). Therefore, we aligned the time series
990 by subtracting from each the difference between Wishbone's estimate of the average
991 percent progress through division and the average time elapsed.

992
993 Note that, because this merging procedure utilized information from Wishbone, it
994 imposes a correlation between time and Wishbone's estimated progress through division.
995 To reduce the impact of this induced correlation, we eliminated the cell images in the
996 middle of each time series, which represent the images that are most affected by this
997 induced correlation. Eliminating 25% or 50% of cell images in this way reduced the
998 correlations by at most 0.05, suggesting these correlations are not driven by our merging
999 procedure.

1000

1001 *Assigning cells to a bin based on progression through cell division*

1002 We used Wishbone to estimate how far each fixed-cell image had progressed through cell
1003 division. Wishbone software requires input about which "start" cell has features

1004 resembling those present at the start of the cell cycle. To identify such features, we used
1005 the data from the live-imaged cell time series. We plotted how single-cell features change
1006 over the course of live imaging, and chose several features that correlate best with
1007 progress through cell division (e.g. cell size, bud size, location of the nucleus). Complete
1008 datasets provided at Open Science Framework (DOI 10.17605/OSF.IO/B7NY5) include
1009 information on which fixed-cell image was chosen as the start cell.

1010

1011 Using Wishbone's estimation of how far each fixed cell had progressed through division,
1012 we assigned each cell to one of 16 equal-sized bins. We did this separately for each of the
1013 374 yeast strains, then merged like bins across strains, such that genetic diversity was
1014 constant across each of the final 16 bins. We obtained very similar results to those
1015 reported in **Figs 4C, 4D**, and **5C** when we used 8 instead of 16 bins. The names of the
1016 traits plotted in **Fig 4** represent succinct summaries of single-cell morphologies
1017 quantified using CalMorph (53). For fuller descriptions of these traits, see the following
1018 trait designations in the CalMorph software manual: **Fig 4A** upper left: C11.1 in
1019 unbudded cells, C101 in budded cells; **Fig 4A** lower left: D184 in small-budded cells,
1020 D182 in unbudded and large-budded cells; **Fig 4A** upper right: C12.2; **Fig 4A** lower
1021 right: D116; **Fig 4C** upper left: C101 and C109 in small-budded cells; **Fig 4C** upper
1022 middle: C11.2 and D132 in small-budded cells; **Fig 4C** upper right: C105 and C113 in
1023 small-budded cells; **Fig 4C** lower left: C114 and D145 in large-budded cells; **Fig 4C**
1024 lower middle: C109 and C126 in large-budded cells; **Fig 4C** lower right: D14.2 and D169
1025 in large-budded cells.

1026

1027 *Eliminating genetic variation at the marker nearest a QTL*

1028 For each of the 27 QTL suspected of horizontal pleiotropy (i.e. pleiotropic QTL that
1029 influence at least one pair of traits for which r_B significantly exceeds r_W), we divided the
1030 374 phenotyped yeast strains into two groups based on whether they inherited the wine or
1031 the oak parent's allele at the genotyped marker closest to the QTL. In some cases, a QTL
1032 spans multiple markers; for example, a QTL on chromosome 15 that influences 64
1033 morphological features spans 14 cM and 4 markers (**Table S1**). These 64 genotype-
1034 phenotype associations are mainly clustered around the ninth marker on chromosome 15,
1035 though a few are closer to the eighth, tenth, or eleventh. To avoid redundancy, for QTL
1036 spanning multiple markers we study the one that is most represented. After dividing
1037 strains into two groups based on which allele they inherited at that marker, we performed
1038 correlation partitioning separately for each group of strains.

1039

1040 The names of the traits plotted in **Fig 5B** represent succinct summaries of single-cell
1041 morphologies quantified using CalMorph. For fuller descriptions of these traits, see the
1042 following trait designations in the CalMorph software manual: upper: D128 and C114 in
1043 large-budded cells; lower: D197 and D17.1 in large-budded cells.

1044

1045 *Quantifying trait correlations within each MA line*

1046 Mutation accumulation occurred in a diploid laboratory yeast strain with genotype *ade2*,
1047 *lys2-801*, *his3-ΔD200*, *leu2-3.112*, *ura3-52* (55). Resulting diploid MA lines were
1048 sporulated to create haploids, which were sequenced in a previous study (56). We
1049 previously imaged these haploid lines in high throughput (>1000 clonal cells imaged per

1050 each of 94 lines) (51). Fewer morphological traits were analyzed in that study than in the
1051 current study, such that there were only 3731 pairs of traits to survey, as opposed to 5645
1052 in the QTL-mapping family. We calculated Pearson correlations between every pair of
1053 traits, separately within each MA line. The names of the traits plotted in **Fig 6A** represent
1054 succinct summaries of single-cell morphologies quantified using CalMorph. For fuller
1055 descriptions of these traits, see the following trait designations in the CalMorph software
1056 manual: upper left: D185 and D186 in large-budded cells; upper right: C102 and D132 in
1057 small-budded cells; lower left: C108 and D167 in large-budded cells; lower right: D135
1058 and D169 in large-budded cells.

1059

1060 *Quantifying trait correlations across drug concentrations*

1061 For this analysis, we imaged the single-cell morphologies of 78 of the 374 strains that
1062 comprised our QTL mapping family. We chose these strains because they were stored
1063 together on a single 96-well plate (the rest of the 96 wells represent blanks or internal
1064 controls), removing concerns about batch effects. We imaged these strains after
1065 exponential growth in three concentrations of geldanamycin (GdA) (8.5 μ M, 25 μ M and
1066 100 μ M). We chose these concentrations because of their wide-ranging impacts on cell
1067 growth rate (51). We obtained single-cell morphology measurements for cells grown in
1068 the lowest concentrations of GdA from our previous study (51) and collected data for
1069 cells grown in higher concentrations following the procedures outlined in that study,
1070 which was very similar to those outlined above, but with a control for the solvent in
1071 which GdA is dissolved. Specifically, cells exposed to GdA were compared to cells
1072 imaged in identical conditions (containing the same concentration of the solvent DMSO)
1073 but lacking GdA. GdA \pm paired experiments are performed side-by-side, with cells
1074 grown in each condition being imaged in adjacent wells on a 384-well microscopy plate.

1075

1076 Resulting morphological data were analyzed following similar procedures as described
1077 above. Briefly, each trait was transformed via a Box-Cox transformation of the raw data
1078 based on the residuals of a linear model with strain, environment, and replicate as effects.
1079 Two replicates were performed for both the 8.5 and 100 μ M environments, and a single
1080 replicate for the 25 μ M environment. Internal controls (wells representing the wine and
1081 oak parents) were used to correct for effects on phenotypic variation that resulted from
1082 differences among replicate experiments. We used WABA II to calculate cell-level (r_w)
1083 correlation coefficients in each of the three GdA concentrations, as well as the
1084 corresponding three control conditions. To calculate the impact of GdA on r_w , we
1085 compared r_w in each drug vs. control condition.

1086

1087

1088 **Author Contributions:**

1089 KGS: Conceptualization, Funding acquisition, Formal analysis, Investigation, Writing – original
1090 draft, Writing – review & editing

1091 SL: Investigation, Writing – review & editing

1092 CL: Conceptualization, Formal analysis, Writing – review & editing

1093 AT: Investigation, Writing – review & editing

1094 NZ: Formal analysis, Writing – review & editing

1095 CR: Investigation, Writing – review & editing

1096 AP: Conceptualization, Writing – review & editing

1097 MLS: Conceptualization, Funding acquisition, Formal analysis, Writing – original draft, Writing –
1098 review & editing

1099

1100

1101

1102 **References:**

1103

1104 1. Plate L. Genetics and evolution. In: Festschrift zum sechzigsten Geburtstag
1105 Richard Hertwigs. Jena: Verlag von Gustav Fischer; 1910. pp. 536–610.

1106 2. Stearns FW. One Hundred Years of Pleiotropy: A Retrospective. *Genetics*;
1107 2010 Nov 1;186(3):767–73.

1108 3. Richardson TG, Harrison S, Hemani G, Smith GD. An atlas of polygenic risk
1109 score associations to highlight putative causal relationships across the human
1110 phenome. *Elife*. 2019;8:e43657.

1111 4. Simons YB, Bullaughey K, Hudson RR, Sella G. A population genetic
1112 interpretation of GWAS findings for human quantitative traits. *PLoS Biol*.
1113 2018 Mar;16(3):e2002985.

1114 5. Tyler AL, Asselbergs FW, Williams SM, Moore JH. Shadows of complexity:
1115 what biological networks reveal about epistasis and pleiotropy. *Bioessays*.
1116 2009 Feb 1;31(2):220–7.

1117 6. Visscher PM, Wray NR, Zhang Q, Sklar P, McCarthy MI, Brown MA, et al.
1118 10 Years of GWAS Discovery: Biology, Function, and Translation. *Am J*
1119 *Hum Genet*. 2017 Jul 6;101(1):5–22.

1120 7. Visscher PM, Yang J. A plethora of pleiotropy across complex traits. *Nat*
1121 *Genet*. 2016 Jun 28;48(7):707–8.

1122 8. Chesmore K, Bartlett J, Williams SM. The ubiquity of pleiotropy in human
1123 disease. *Hum Genet*. 137(1):39–44.

1124 9. Sivakumaran S, Agakov F, Theodoratou E, Prendergast JG, Zgaga L,
1125 Manolio T, et al. Abundant pleiotropy in human complex diseases and traits.
1126 *Am J Hum Genet*. 2011 Nov 11;89(5):607–18.

1127 10. White JK, Gerdin A-K, Karp NA, Ryder E, Buljan M, Bussell JN, et al.
1128 Genome-wide generation and systematic phenotyping of knockout mice
1129 reveals new roles for many genes. *Cell*. 2013 Jul 18;154(2):452–64.

1130 11. Wang Z, Liao B-Y, Zhang J. Genomic patterns of pleiotropy and the
1131 evolution of complexity. *Proc Natl Acad Sci USA*. 2010 Oct
1132 19;107(42):18034–9.

- 1133 12. Housman G, Byler S, Heerboth S, Lapinska K, Longacre M, Snyder N, et al.
1134 Drug resistance in cancer: an overview. *Cancers (Basel)*. 2014 Sep
1135 5;6(3):1769–92.
- 1136 13. Rockman MV. The QTN program and the alleles that matter for evolution: all
1137 that's gold does not glitter. *Evolution*. 2012 Jan;66(1):1–17.
- 1138 14. Boyle EA, Li YI, Pritchard JK. An Expanded View of Complex Traits: From
1139 Polygenic to Omnigenic. *Cell*. 2017 Jun;169(7):1177–86.
- 1140 15. Fisher RA. *The Genetical Theory of Selection*. Oxford: Clarendon; 1930.
- 1141 16. Orr HA. Adaptation and the cost of complexity. *Evolution*. 2000;54(1):13.
- 1142 17. Wagner GP, Kenney-Hunt JP, Pavlicev M, Peck JR, Waxman D, Cheverud
1143 JM. Pleiotropic scaling of gene effects and the 'cost of complexity'. *Nature*.
1144 2008 Mar 27;452(7186):470–2.
- 1145 18. Wagner GP, Zhang J. The pleiotropic structure of the genotype–phenotype
1146 map: the evolvability of complex organisms. *Nat Rev Genet*. 2011 Mar
1147 1;12(3):204–13.
- 1148 19. McGee LW, Sackman AM, Morrison AJ, Pierce J, Anisman J, Rokyta DR.
1149 Synergistic Pleiotropy Overrides the Costs of Complexity in Viral
1150 Adaptation. *Genetics*. 2016 Jan;202(1):285–95.
- 1151 20. Wittkopp PJ, Haerum BK, Clark AG. Regulatory changes underlying
1152 expression differences within and between *Drosophila* species. *Nat Genet*.
1153 2008 Mar 1;40(3):346–50.
- 1154 21. Stern DL. Perspective: evolutionary developmental biology and the problem
1155 of variation. *Evolution*. 2009 Jan 20;54(4):1079.
- 1156 22. He X, Zhang J. Toward a Molecular Understanding of Pleiotropy. *Genetics*.
1157 2006 Aug 1;173(4):1885–91.
- 1158 23. Papakostas S, Vøllestad LA, Bruneaux M, Aykanat T, Vanoverbeke J, Ning
1159 M, et al. Gene pleiotropy constrains gene expression changes in fish adapted
1160 to different thermal conditions. *Nat Commun*. 2014 Jun 3;5:4071.
- 1161 24. Altenberg L. Modularity in Evolution: Some Low-Level Questions. In:
1162 Callebaut W, Rasskin-Gutman D, editors. *Modularity: Understanding the*
1163 *Development and Evolution of Natural Complex Systems*. Cambridge: MIT
1164 Press; 2005. pp. 99–128.
- 1165 25. Wagner GP, Altenberg L. Perspective: Complex Adaptations and the
1166 Evolution of Evolvability. *Evolution*. 1996 Jun;50(3):967.

- 1167 26. Welch JJ, Waxman D. Modularity and the cost of complexity. *Evolution*.
1168 2003 Aug;57(8):1723–34.
- 1169 27. Wagner GP, Pavlicev M, Cheverud JM. The road to modularity. *Nat Rev*
1170 *Genet*. 2007 Dec;8(12):921–31.
- 1171 28. Melo D, Porto A, Cheverud JM, Marroig G. Modularity: genes, development
1172 and evolution. *Annu Rev Ecol Evol Syst*. 2016;47(1):463–86.
- 1173 29. Collet JM, McGuigan K, Allen SL, Chenoweth SF, Blows MW. Mutational
1174 Pleiotropy and the Strength of Stabilizing Selection Within and Between
1175 Functional Modules of Gene Expression. *Genetics*. 2018 Apr;208(4):1601–
1176 16.
- 1177 30. Paaby AB, Rockman MV. Pleiotropy: what do you mean? Reply to Zhang
1178 and Wagner. *Trends Genet*. 2013 Jul;29(7):384.
- 1179 31. Wagner GP, Zhang J. Universal pleiotropy is not a valid null hypothesis:
1180 reply to Hill and Zhang. *Nat Rev Genet*. 2012 Apr 1;13(4):296–6.
- 1181 32. Hill WG, Zhang X-S. Assessing pleiotropy and its evolutionary
1182 consequences: pleiotropy is not necessarily limited, nor need it hinder the
1183 evolution of complexity. *Nat Rev Genet*. Nature Publishing Group; 2012 Apr
1184 1;13(4):296.
- 1185 33. Zhang J, Wagner GP. On the definition and measurement of pleiotropy.
1186 *Trends Genet*. 2013 Jul;29(7):383–4.
- 1187 34. Gibson G. Decanalization and the origin of complex disease. *Nat Rev Genet*.
1188 2009 Feb;10(2):134–40.
- 1189 35. Nuzhdin SV, Friesen ML, McIntyre LM. Genotype-phenotype mapping in a
1190 post-GWAS world. *Trends Genet*. 2012 Sep;28(9):421–6.
- 1191 36. Li Y, Venkataram S, Agarwala A, Dunn B, Petrov DA, Sherlock G, et al.
1192 Hidden Complexity of Yeast Adaptation under Simple Evolutionary
1193 Conditions. *Curr Biol*. 2018 Feb 19;28(4):515–6.
- 1194 37. Gorter FA, Aarts MGM, Zwaan BJ, de Visser JAGM. Local Fitness
1195 Landscapes Predict Yeast Evolutionary Dynamics in Directionally Changing
1196 Environments. *Genetics*. 2018 Jan;208(1):307–22.
- 1197 38. Baym M, Stone LK, Kishony R. Multidrug evolutionary strategies to reverse
1198 antibiotic resistance. *Science*. 2016 Jan 1;351(6268):aad3292–2.
- 1199 39. Hartley SW, Monti S, Liu C-T, Steinberg MH, Sebastiani P. Bayesian
1200 methods for multivariate modeling of pleiotropic SNP associations and
1201 genetic risk prediction. *Front Genet*. 2012;3:176.

- 1202 40. Chen G, Mulla WA, Kucharavy A, Tsai H-J, Rubinstein B, Conkright J, et al.
1203 Targeting the adaptability of heterogeneous aneuploids. *Cell*. 2015 Feb
1204 12;160(4):771–84.
- 1205 41. Nichol D, Jeavons P, Fletcher AG, Bonomo RA, Maini PK, Paul JL, et al.
1206 Steering Evolution with Sequential Therapy to Prevent the Emergence of
1207 Bacterial Antibiotic Resistance. *PLoS Comput Biol*. 2015
1208 Sep;11(9):e1004493.
- 1209 42. Kaznatcheev A, Peacock J, Basanta D, Marusyk A, Scott JG. Fibroblasts and
1210 alectinib switch the evolutionary games played by non-small cell lung cancer.
1211 *Nat Ecol Evol*. 2019 Mar;3(3):450–6.
- 1212 43. Paaby AB, Rockman MV. The many faces of pleiotropy. *Trends Genet*. 2013
1213 Feb;29(2):66–73.
- 1214 44. Walsh B, Blows MW. Abundant Genetic Variation + Strong Selection =
1215 Multivariate Genetic Constraints: A Geometric View of Adaptation. *Annu*
1216 *Rev Ecol Evol Syst*. 2009 Nov 9;40(1):41–59.
- 1217 45. Ference BA, Ginsberg HN, Graham I, Ray KK, Packard CJ, Bruckert E, et al.
1218 Low-density lipoproteins cause atherosclerotic cardiovascular disease. 1.
1219 Evidence from genetic, epidemiologic, and clinical studies. A consensus
1220 statement from the European Atherosclerosis Society Consensus Panel. *Eur*
1221 *Heart J*. 2017 Aug 21;38(32):2459–72.
- 1222 46. Al Hafid N, Christodoulou J. Phenylketonuria: a review of current and future
1223 treatments. *Transl Pediatr*. 2015 Oct;4(4):304–17.
- 1224 47. Powell ALT, Nguyen CV, Hill T, Cheng KL, Figueroa-Balderas R, Aktas H,
1225 et al. Uniform ripening encodes a Golden 2-like transcription factor
1226 regulating tomato fruit chloroplast development. *Science*. 2012 Jun
1227 29;336(6089):1711–5.
- 1228 48. Grüneberg H. An analysis of the “pleiotropic” effects of a new lethal
1229 mutation in the rat (*Mus norvegicus*). *Proc R Soc Lond B Biol Sci*. 1938 Mar
1230 3;125(838):123–144.
- 1231 49. Ohya Y, Sese J, Yukawa M, Sano F, Nakatani Y, Saito TL, et al. High-
1232 dimensional and large-scale phenotyping of yeast mutants. *Proceedings of the*
1233 *National Academy of Sciences*. 2005 Dec 27;102(52):19015–20.
- 1234 50. Bauer CR, Li S, Siegal ML. Essential gene disruptions reveal complex
1235 relationships between phenotypic robustness, pleiotropy, and fitness. *Mol*
1236 *Syst Biol*. 2015 Jan;11(1):773–3.

- 1237 51. Geiler-Samerotte KA, Zhu YO, Goulet BE, Hall DW, Siegal ML. Selection
1238 Transforms the Landscape of Genetic Variation Interacting with Hsp90. *PLoS*
1239 *Biol.* 2016 Oct;14(10):e2000465.
- 1240 52. Richardson JB, Uppendahl LD, Traficante MK, Levy SF, Siegal ML. Histone
1241 variant HTZ1 shows extensive epistasis with, but does not increase robustness
1242 to, new mutations. *PLoS Genet.* 2013;9(8):e1003733.
- 1243 53. Negishi T, Nogami S, Ohya Y. Multidimensional quantification of subcellular
1244 morphology of *Saccharomyces cerevisiae* using CalMorph, the high-
1245 throughput image-processing program. *J Biotechnol.* 2009 May 20;141(3-
1246 4):109–17.
- 1247 54. Gerke J, Lorenz K, Cohen B. Genetic interactions between transcription
1248 factors cause natural variation in yeast. *Science.* 2009 Jan 23;323(5913):498–
1249 501.
- 1250 55. Hall DW, Mahmoudizad R, Hurd AW, Joseph SB. Spontaneous mutations in
1251 diploid *Saccharomyces cerevisiae*: another thousand cell generations. *Genet*
1252 *Res (Camb).* 2008 Jun;90(3):229–41.
- 1253 56. Zhu YO, Siegal ML, Hall DW, Petrov DA. Precise estimates of mutation rate
1254 and spectrum in yeast. *Proc Natl Acad Sci USA.* 2014 Jun 3;111(22):E2310–
1255 8.
- 1256 57. Cooper TF, Ostrowski EA, Travisano M. A negative relationship between
1257 mutation pleiotropy and fitness effect in yeast. *Evolution.* 2007
1258 Jun;61(6):1495–9.
- 1259 58. Yang M, Ohnuki S, Ohya Y. Unveiling nonessential gene deletions that
1260 confer significant morphological phenotypes beyond natural yeast strains.
1261 *BMC Genomics.* 2014;15(1):932.
- 1262 59. Sahai E. Mechanisms of cancer cell invasion. *Curr Opin Genet Dev.* 2005
1263 Feb;15(1):87–96.
- 1264 60. Gerke JP, Chen CTL, Cohen BA. Natural isolates of *Saccharomyces*
1265 *cerevisiae* display complex genetic variation in sporulation efficiency.
1266 *Genetics.* 2006 Oct;174(2):985–97.
- 1267 61. Liti G, Carter DM, Moses AM, Warringer J, Parts L, James SA, et al.
1268 Population genomics of domestic and wild yeasts. *Nature.* 2009 Mar
1269 19;458(7236):337–41.
- 1270 62. Skelly DA, Merrihew GE, Riffle M, Connelly CF, Kerr EO, Johansson M, et
1271 al. Integrative phenomics reveals insight into the structure of phenotypic
1272 diversity in budding yeast. *Genome Res.* 2013 Sep;23(9):1496–504.

- 1273 63. Ziv N, Shuster BM, Siegal ML, Gresham D. Resolving the Complex Genetic
1274 Basis of Phenotypic Variation and Variability of Cellular Growth. *Genetics*.
1275 2017 May 11;:genetics.116.195180.
- 1276 64. Broman KW, Wu H, Sen S, Churchill GA. R/qtl: QTL mapping in
1277 experimental crosses. *Bioinformatics*. 2003 May 1;19(7):889–90.
- 1278 65. Churchill GA, Doerge RW. Empirical threshold values for quantitative trait
1279 mapping. *Genetics*. 1994 Nov;138(3):963–71.
- 1280 66. Stuckey S, Storici F. Gene knockouts, in vivo site-directed mutagenesis and
1281 other modifications using the delitto perfetto system in *Saccharomyces*
1282 *cerevisiae*. *Meth Enzymol*. 2013;533:103–31.
- 1283 67. Cherry JM, Hong EL, Amundsen C, Balakrishnan R, Binkley G, Chan ET, et
1284 al. *Saccharomyces* Genome Database: the genomics resource of budding
1285 yeast. *Nucleic Acids Res*. 2012 Jan;40:D700–5.
- 1286 68. Graziano BR, Yu H-YE, Alioto SL, Eskin JA, Ydenberg CA, Waterman DP,
1287 et al. The F-BAR protein Hof1 tunes formin activity to sculpt actin cables
1288 during polarized growth. *Mol Biol Cell*. 2014 Jun;25(11):1730–43.
- 1289 69. Nogami S, Ohya Y, Yvert G. Genetic Complexity and Quantitative Trait Loci
1290 Mapping of Yeast Morphological Traits. *PLoS Genet*. 2007 Feb 23;3(2):e31.
- 1291 70. Lynch M, Walsh B. Correlations between characters. In: *Genetics and*
1292 *Analysis of Quantitative Traits*. Sinauer Associates Incorporated; 1998. pp.
1293 629–56.
- 1294 71. Bliese P. Multilevel modeling in R (2.6). 2013. Available from: [https://cran.r-](https://cran.r-project.org/doc/contrib/Bliese_Multilevel.pdf)
1295 [project.org/doc/contrib/Bliese_Multilevel.pdf](https://cran.r-project.org/doc/contrib/Bliese_Multilevel.pdf)
- 1296 72. Sockloff AL. Behavior of the Product-Moment Correlation Coefficient When
1297 Two Heterogeneous Subgroups Are Pooled. *Educational and Psychological*
1298 *Measurement*. 1975;35(2):267–276.
- 1299 73. Breiman L, Friedman JH. Estimating Optimal Transformations for Multiple
1300 Regression and Correlation. *Journal of the American Statistical Association*.
1301 2012 Mar 12;80(391):580–98.
- 1302 74. Setty M, Tadmor MD, Reich-Zeliger S, Angel O, Salame TM, Kathail P, et
1303 al. Wishbone identifies bifurcating developmental trajectories from single-cell
1304 data. *Nat Biotechnol*. 2016 Jun;34(6):637–45.
- 1305 75. Turner JJ, Ewald JC, Skotheim JM. Cell size control in yeast. *Curr Biol*. 2012
1306 May 8;22(9):R350–9.

- 1307 76. Wang R, Kamgoue A, Normand C, Léger-Silvestre I, Mangeat T, Gadad O.
1308 High resolution microscopy reveals the nuclear shape of budding yeast during
1309 cell cycle and in various biological states. *J Cell Sci.* 2016 Dec
1310 15;129(24):4480–95.
- 1311 77. Logg K, Bodvard K, Blomberg A, Käll M. Investigations on light-induced
1312 stress in fluorescence microscopy using nuclear localization of the
1313 transcription factor Msn2p as a reporter. *FEMS Yeast Res.* 2009
1314 Sep;9(6):875–84.
- 1315 78. Pavlicev M, Cheverud JM. Constraints Evolve: Context Dependency of Gene
1316 Effects Allows Evolution of Pleiotropy. *Annu Rev Ecol Evol Syst.*
1317 2015;46(1):413:434.
- 1318 79. Pavlicev M, Wagner GP, Noonan JP, Hallgrímsson B, Cheverud JM.
1319 Genomic correlates of relationship QTL involved in fore- versus hind limb
1320 divergence in mice. *Genome Biol Evol.* 2013;5(10):1926–36.
- 1321 80. Porto A, Schmelter R, VandeBerg JL, Marroig G, Cheverud JM. Evolution of
1322 the genotype-to-phenotype map and the cost of pleiotropy in mammals.
1323 *Genetics.* 2016;204(4):1601–12.
- 1324 81. Dudley AM, Janse DM, Tanay A, Shamir R, Church GM. A global view of
1325 pleiotropy and phenotypically derived gene function in yeast. *Mol Syst Biol.*
1326 2005 Jan 1;1(1):E1–E11.
- 1327 82. Lande R. Quantitative genetic analysis of multivariate evolution, applied to
1328 brain:body size allometry. *Evolution.* 2017 May 31;33(1 Part 2):402–16.
- 1329 83. Arnold SJ, Bürger R, Hohenlohe PA, Ajie BC, Jones AG. Understanding the
1330 evolution and stability of the g-matrix. *Evolution.* 2008 Oct;62(10):2451–61.
- 1331 84. Björklund M, Husby A, Gustafsson L. Rapid and unpredictable changes of
1332 the G-matrix in a natural bird population over 25 years. *J Evol Biol.* 2012 Dec
1333 13;26(1):1–13.
- 1334 85. Cheverud JM. A comparison of genetic and phenotypic correlations.
1335 *Evolution.* 2017 Jun 5;42(5):958–68.
- 1336 86. Arnold SJ, Phillips PC. Hierarchical comparison of genetic variance-
1337 covariance matrices. II coastal-inland divergence in the garter snake,
1338 *thamnophis elegans*. *Evolution.* 2017 May 31;53(5):1516–27.
- 1339 87. Sgrò CM, Hoffmann AA. Genetic correlations, tradeoffs and environmental
1340 variation. *Heredity.* 2004;93:241–8.

- 1341 88. Eguchi Y, Bilollikar G, Geiler-Samerotte K. Why and how to study genetic
1342 changes with context-dependent effects. *Curr Opin Genet Dev.* 2019 Oct;58–
1343 59:95–102.
- 1344 89. Pavlicev M, Wagner GP. Evolutionary Systems Biology: Shifting Focus to
1345 the Context-Dependency of Genetic Effects. In: Martin LB, Ghalambor CK,
1346 Woods HA, editors. *Integrative Organismal Biology.* 2014. pp. 91–108.
- 1347 90. Levy SF, Siegal ML. Network hubs buffer environmental variation in
1348 *Saccharomyces cerevisiae*. *PLoS Biol.* 2008 Nov 4;6(11):e264.
- 1349 91. Ziv N, Siegal ML, Gresham D. Genetic and nongenetic determinants of cell
1350 growth variation assessed by high-throughput microscopy. *Mol Biol Evol.*
1351 2013 Dec;30(12):2568–78.
- 1352 92. Millard SP. *EnvStats.* New York, NY: Springer Science & Business Media;
1353 2013.
- 1354 93. Storey JD, Bass A, Dabney A, Robinson D. *qvalue: Q-value estimation for*
1355 *false discovery rate control. R package version 2.16.0.* 2019.
- 1356 94. Gietz RD, Woods RA. Transformation of yeast by lithium acetate/single-
1357 stranded carrier DNA/polyethylene glycol method. *Meth Enzymol.*
1358 2002;350:87–96.
- 1359 95. Miyajima I, Nakafuku M, Nakayama N, Brenner C, Miyajima A, Kaibuchi K,
1360 et al. GPA1, a haploid-specific essential gene, encodes a yeast homolog of
1361 mammalian G protein which may be involved in mating factor signal
1362 transduction. *Cell.* 1987 Sep 25;50(7):1011–9.
- 1363 96. Pinheiro J, Bates D, DebRoy S, Sarkar D, R Core Team. *nlme: linear and*
1364 *nonlinear mixed effects models. R package version 3.1-137.* 2018.
- 1365 97. Levy SF, Ziv N, Siegal ML. Bet hedging in yeast by heterogeneous, age-
1366 correlated expression of a stress protectant. *PLoS Biol.* 2012;10(5):e1001325.
- 1367 98. Li S, Giardina DM, Siegal ML. Control of nongenetic heterogeneity in
1368 growth rate and stress tolerance of *Saccharomyces cerevisiae* by cyclic AMP-
1369 regulated transcription factors. *PLoS Genet.* 2018 Nov;14(11):e1007744.
- 1370 99. Wickham H. *ggplot2.* New York, NY: Springer New York; 2009.

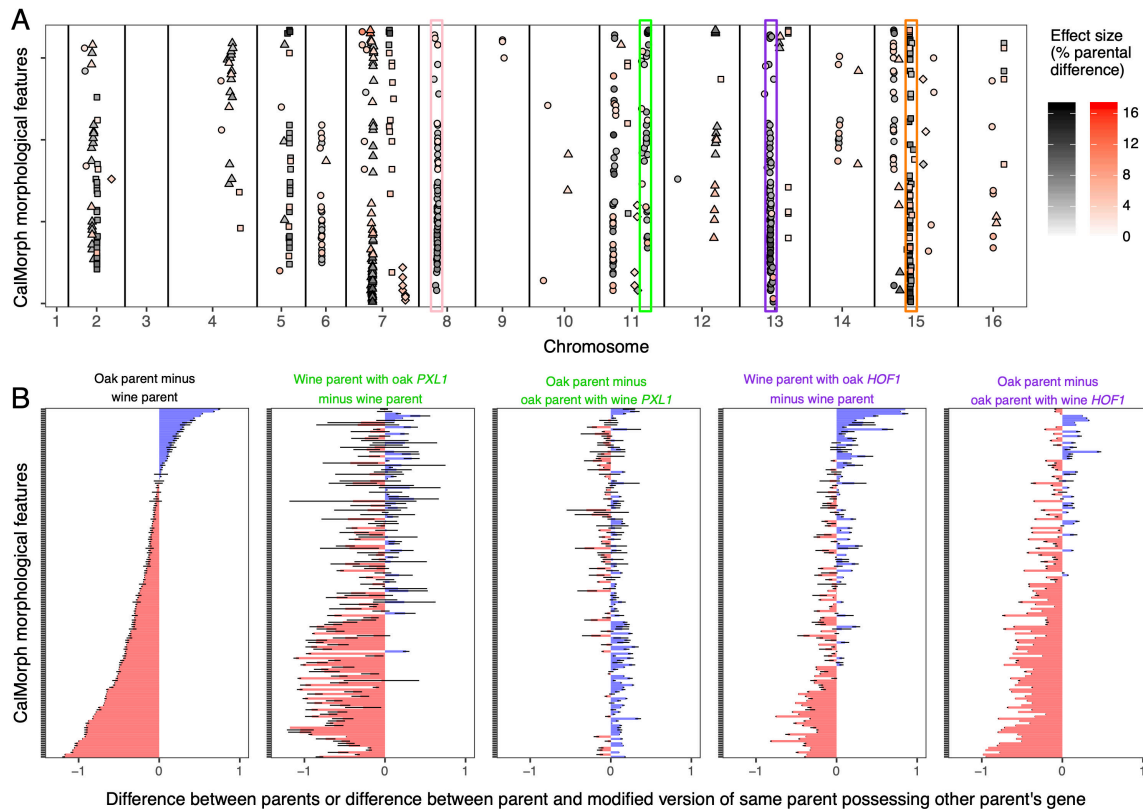


Figure 1: Pleiotropic QTL influence yeast single-cell morphology. The vertical axes in all plots represent the 155 CalMorph morphological traits for which we detect QTL. These traits are sorted, from top to bottom, based on the difference between the oak and wine parental strains. **(A)** Of 41 QTL that contribute to variation in single-cell morphology, 36 contribute to variation in multiple features. The horizontal axis indicates the chromosomal location of each QTL (in cM). Differently shaped points indicate separate QTL that are more than 5 cM apart on the same chromosome. The darkness of a point represents the effect size of a QTL; effect sizes range from 1.3% (lightest points) to 17.5% (darkest points) of the difference between parents. All points represent genotype-phenotype associations detected using a per-trait genome-wide type one error rate of 5%. The points highlighted in black are significant after correcting for testing multiple phenotypes using a false discovery rate of 5%. The QTL highlighted in pink, green, purple, and orange are very pleiotropic, contributing to 57, 30, 73, or 64 morphological features, respectively. **(B)** Gene-swapping experiments demonstrate that single genes contribute to multiple morphological features. The horizontal axis represents the relative phenotypic differences between the wine and oak parents (leftmost column) or one of these strains versus a derivative strain that differs in a single gene. The relative phenotypic differences between a pair of strains are calculated by scaling each trait to have a mean of 0 and standard deviation of 1 across all cells in both strains, and then subtracting the average value in one strain from that in the other. To control for variation among replicate experiments, this scaling was done independently for each replicate experiment in which both strains were imaged. Error bars represent 95% confidence intervals inferred from the replicate experiments. The two gene replacements shown, *PXL1* and *HOF1*, are respectively located within the QTL highlighted in green and purple in panel A. When calculating the difference between strains, we always subtracted the trait values of the strain possessing more wine genes from those of the strain possessing more oak genes, such that the effects of the wine or oak gene replacements appear in the same direction on all plots.

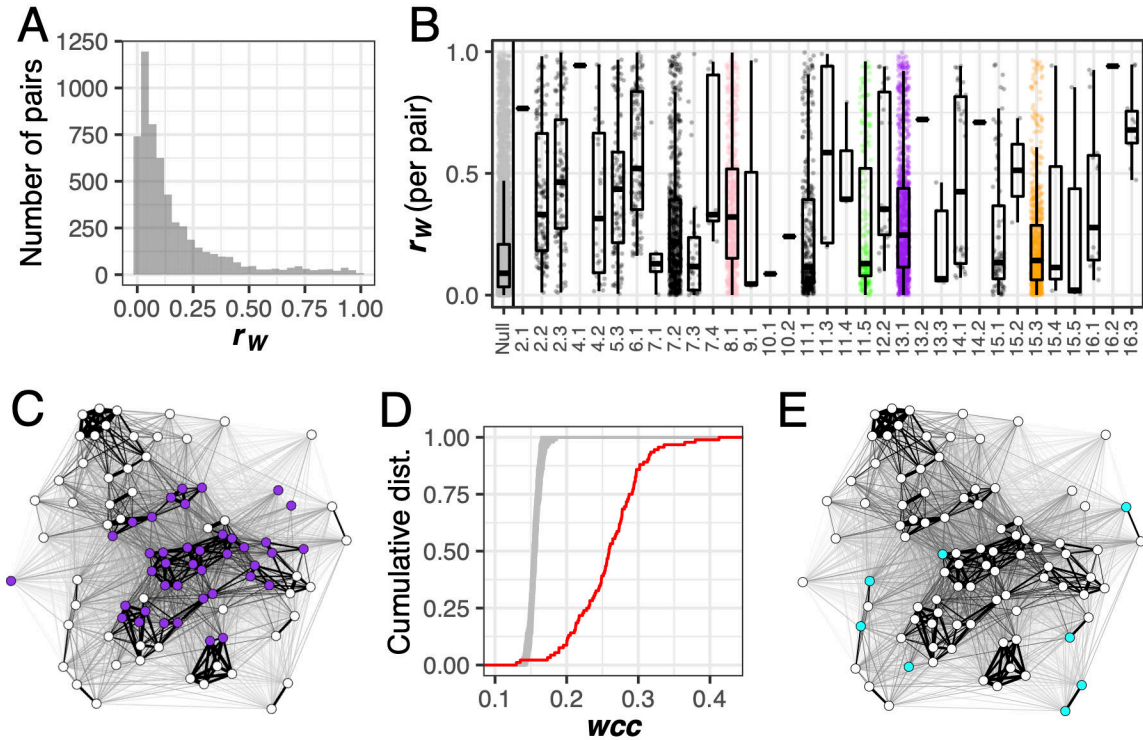


Figure 2: Pairs of traits with high correlation across clones are overrepresented among those influenced by pleiotropic QTL. Within-genotype correlations (r_w) are calculated for 5645 pairs of morphological traits. **(A)** Histogram showing distribution of r_w . **(B)** Each point represents r_w for a pair of traits. The null distribution displays r_w for all 5645 pairs of traits. Every other distribution displays r_w for pairs of traits influenced by a single QTL. Horizontal axis labels represent the chromosome on which a QTL resides, followed by the order in which that QTL appears on the chromosome. Colored points correspond to those QTL highlighted in the same color in **Fig 1A**. Each boxplot shows the median (center line), interquartile range (IQR) (upper and lower hinges), and highest value within $1.5 \times \text{IQR}$ (whiskers). **(C)** Force-directed network visualizing how pairs of morphological features correlate across clones. Each node represents a single-cell morphological trait measured in large-budded cells. For networks representing traits from unbudded and small budded cells, see **Fig S4**. The thickness of the line connecting each pair of nodes is proportional to r_w . Node position in the network is determined using the Fruchterman-Reingold algorithm. Purple nodes correspond to traits influenced by a QTL on chromosome 13 containing the *HOF1* gene. **(D)** Cumulative distributions of weighted clustering coefficients (wcc) in a network created using measured values of r_w (red line) or in 100 permuted networks (grey lines) for traits corresponding to large-budded cells. Permutations were performed by sampling r_w , without replacement, and reassigning each value to a random pair of traits. For distributions summarizing wcc in networks representing traits from unbudded and small budded cells, see **Fig S4**. **(E)** The same network as in panel C with colored nodes corresponding to traits influenced by the leftmost QTL on chromosome 15.




Figure 3: Natural genetic variation affects the correlation between morphological features. The absolute value of the between-strain correlation (r_B), made negative when r_B and r_w have opposite signs, is plotted against the absolute value of the within-strain correlation (r_w), for each pair of traits. The plot at left shows pairs of traits that share at least one QTL influence. The color of each point represents the number of pleiotropic QTL that influence both traits in that pair. The plot at right shows pairs of traits that share no QTL influence. The dashed line represents a Bonferroni-corrected significance threshold of $p < 0.01$. Insets represent the results of correlation partitioning performed after randomly assigning individual cells to groups (pseudo-strains) having the same numbers of cells as the actual strains.

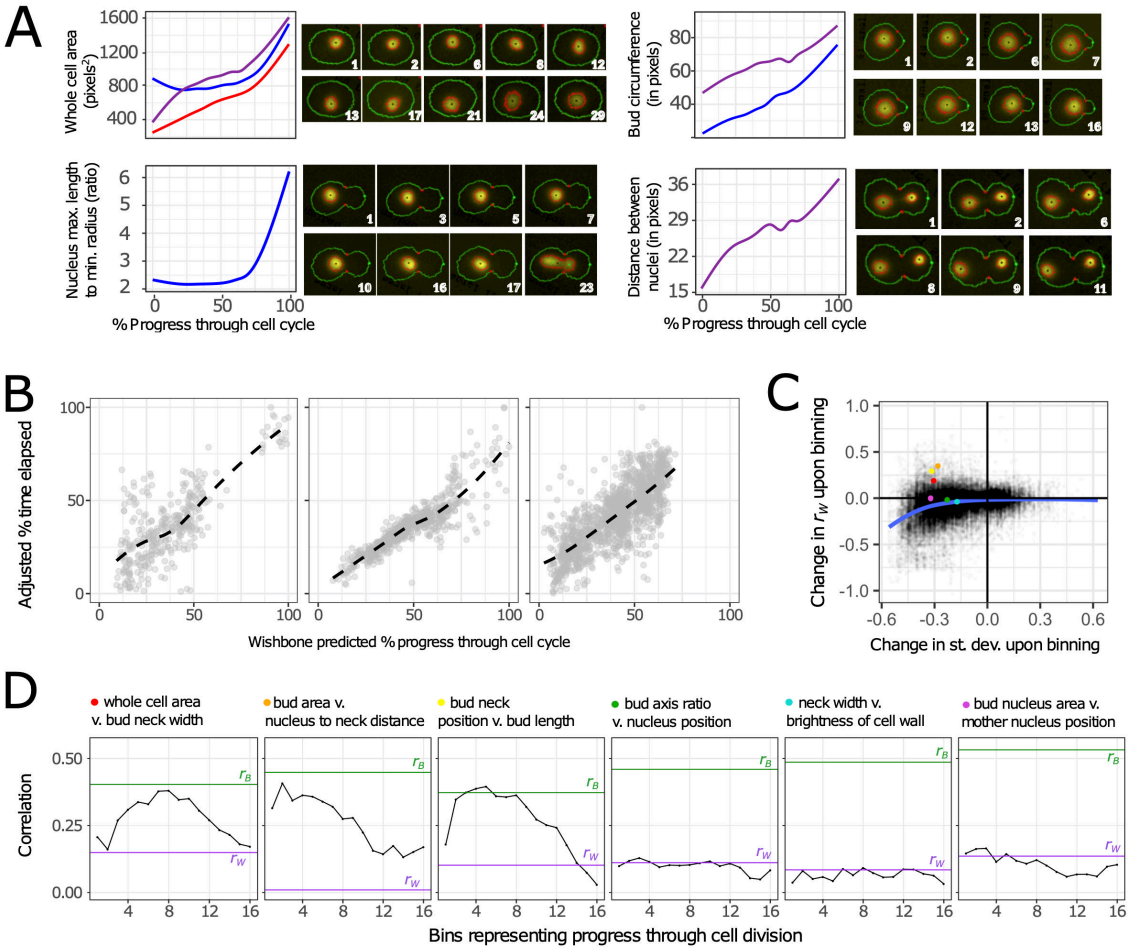


Figure 4: Morphological features vary as cells divide. The morphological features of unbudded (red), small-budded (blue), and large-budded (purple) cells change as these cells progress through the cell cycle. **(A)** Variation of four traits through the cell cycle. Line plots represent fixed-cell images from all 374 mapping family strains, positioned on the horizontal axis based on progression through the cell cycle as calculated by Wishbone (74). Regression lines are smoothed with cubic splines, calculated with the “gam” method in the R package ggplot2 (99), to depict trends describing how each displayed trait varies across the estimated growth trajectory. The displayed trends match those observed in micrographs of live cells progressing through division. Each series of micrographs displays a different live cell imaged over several minutes, which are displayed in the lower right corner of each micrograph. **(B)** Centered data for 11, 23, and 44 unbudded, small-budded and large-budded cells, respectively, show how Wishbone sorts live cells in a way that recapitulates the actual time series. Each point in these plots represents a cell image from a single timepoint. The horizontal axis represents Wishbone’s estimation of how far that cell has progressed through division. The vertical axis displays time, as a percentage of the total time elapsed and adjusted in a way that controls for every cell having started at a different place in the cell division cycle at time zero (see *Methods*). Trend lines are smooth fits using the “loess” method in the R package ggplot (99). **(C)** The correlation between some morphological features changes throughout the course of cell division. The scatterplot shows how binning influences both the phenotypic correlation (vertical axis) and phenotypic variation (horizontal axis) across clones. Each point represents these values for a pair of traits as measured in 1 of 16 bins. The value on the horizontal axis represents whichever trait in each pair had the larger decrease in standard deviation, as such decreases are likely to reduce the correlation on the vertical axis. The blue line shows a smooth fit by loess regression. Colored points on the scatterplot correspond to bin 5 for each pair of traits represented by the line plots in panel **D**. **(D)** These line plots show three pairs of traits for which binning increases r_W such that it approaches r_B (leftmost three plots), and three pairs of traits for which r_W does not approach r_B even after binning (rightmost three plots). In each plot, r_B is shown as the horizontal green line, r_W (without binning) is shown as the horizontal purple line, and r_W for each bin is shown in black.

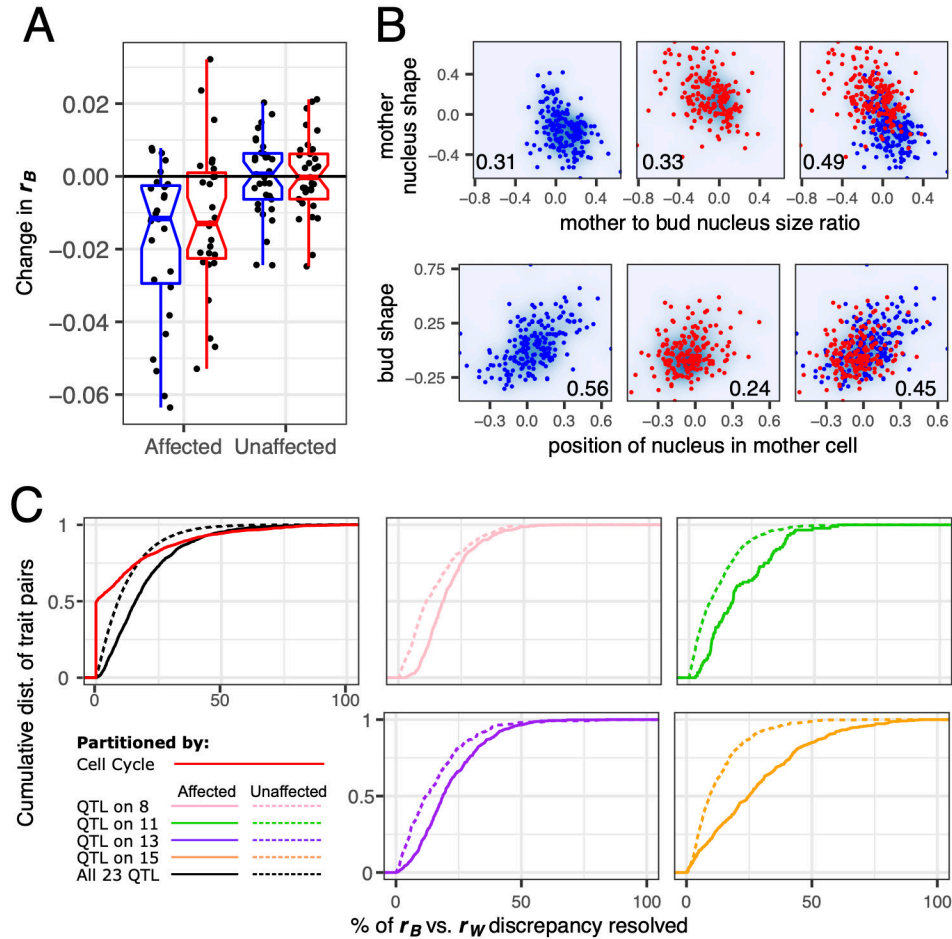


Figure 5: Many QTL demonstrate horizontal pleiotropy. (A) Eliminating allelic variation at the site of each QTL tends to reduce r_B . The vertical axis represents how r_B changes upon eliminating allelic variation at each QTL site. Each point represents the median change in r_B for all pairs of traits that are affected or unaffected by one of the 27 QTL suspected of horizontal pleiotropy. Boxplots summarize these changes in r_B when re-measured across strains possessing the wine (red) or the oak (blue) allele at the marker closest to the QTL. (B) The upper and lower series of three plots demonstrate two different ways that a QTL can increase the correlation between traits. Each point represents a yeast strain possessing either the wine (red) or the oak (blue) allele at a marker closest to a QTL on chromosome 15 (upper) or 8 (lower). In the upper plots, the QTL increases the correlation between nucleus shape and size ratio when it is segregating across strains. In the lower plots, the oak allele strengthens a correlation between bud shape and the position of the nucleus in the mother cell that is weak in the wine subpopulation. Numbers in the lower corner of each plot represent r_B for the strains displayed. (C) Cumulative distributions display the extent to which binning cells or splitting strains resolves the difference between r_B and r_W . When calculating percent resolved (horizontal axes) we always plot the value in whichever subset (e.g. wine or oak) this percent is greatest. If subsetting always worsens the discrepancy between r_B vs. r_W , we score this as 0% resolution. Only pairs of traits for which r_B is significantly greater than r_W are considered. The pink, green, purple and orange lines show the effect of splitting strains by whether they inherited the wine or oak allele at the marker closest to each of four QTL (colors correspond to QTL in Fig 1A). In these plots, comparing the solid vs dotted lines shows that splitting strains resolves the discrepancy between r_B and r_W more often for pairs in which both traits are affected by the QTL than pairs in which both traits are unaffected. The black lines in the leftmost plot summarize these effects across 27 QTL, displaying for each trait pair, the largest resolution in the r_B vs. r_W discrepancy observed across all QTL that affect the pair of traits (solid line) or all QTL that do not (dotted line). The red line shows the effect of binning cells by their progress through division, displaying the largest resolution in the r_B vs. r_W difference across all 16 bins.

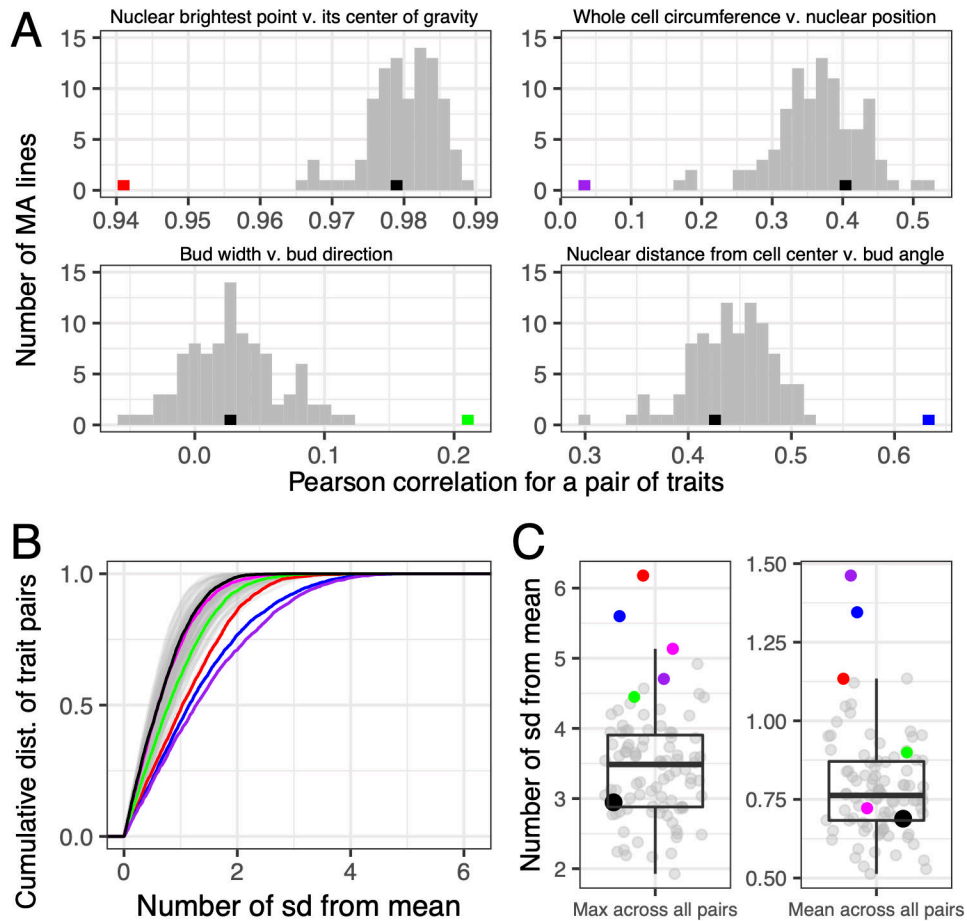


Figure 6: Some MA lines display unique relationships between certain pairs of traits. In all plots, black represents the ancestor of the MA lines and colors represent MA lines with trait correlations that differ from other lines (strains: black = HAncestor, green = DHC81H1, red = DHC41H1, magenta = DHC40H1, blue = DHC66H1, purple = DHC84H1; see Table S2 in Geiler-Samerotte et al 2016 (51)). **(A)** Histograms display the number of MA lines with Pearson correlations corresponding to the values on the horizontal axis for four example pairs of traits; the number of bins is set to 30. **(B)** This plot displays, for each of the 94 MA lines, the cumulative distribution of the number of standard deviations away from the mean correlation across all trait pairs. **(C)** Plots display, for each MA line, the maximum deviation from the mean observed for any pair of traits (left) and the average standard deviation observed across all pairs of traits (right).

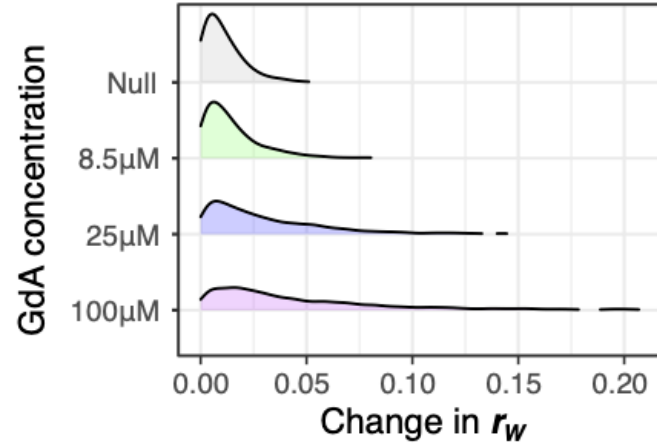


Figure 7: The correlations between traits changes depending on drug concentration. Plots display the density of trait pairs for which the within-strain correlation (r_w) changes by the amount shown on the horizontal axis. To calculate how each drug treatment changes r_w , we subtracted r_w observed for a pair of traits in the drug condition from that observed in a paired experiment that lacked the drug. The absolute value of this change is displayed. These changes are smallest in the null condition, which represents the change in r_w observed across replicate experiments lacking the drug. For clarity, we exclude regions of the plot for which density is less than 0.01.

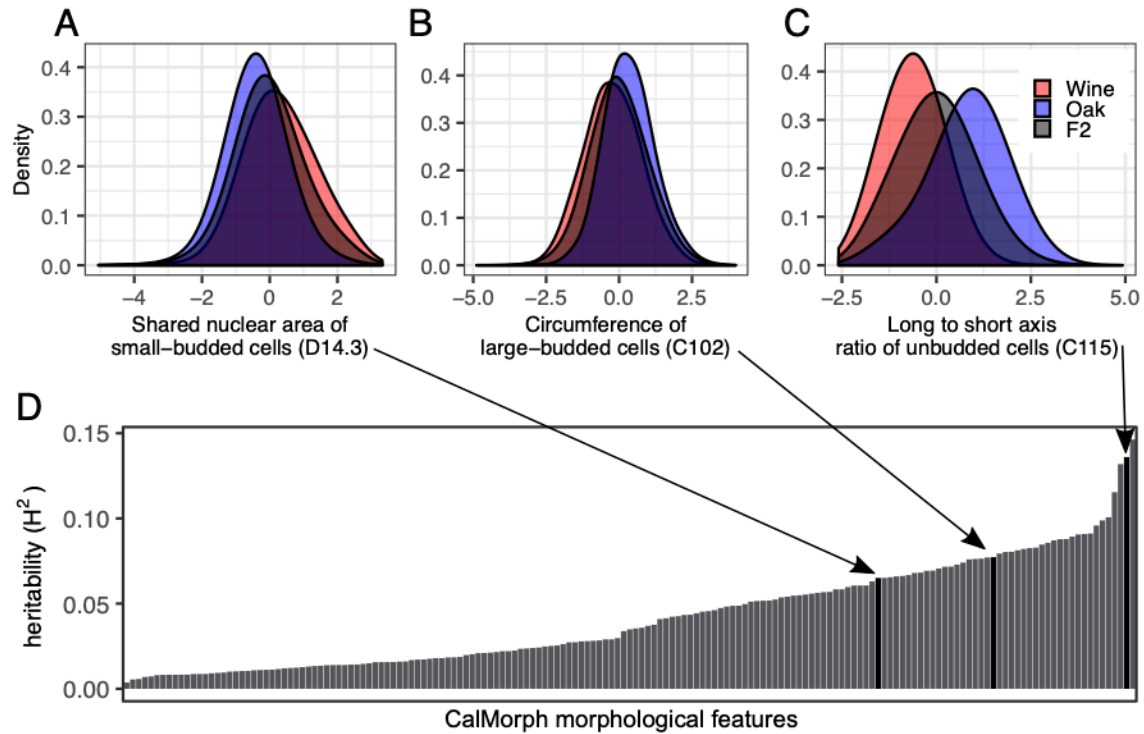


Figure S1: Morphological differences exist between the parents of the QTL mapping family. (A–C) Each density plot displays the distribution of phenotype values from yeast cells corresponding to the wine parent (red), the oak parent (blue), or all of the 374 progeny (grey) for the trait listed on the horizontal axis. Trait names in parentheses correspond to those listed in the CalMorph manual (53). Each distribution represents at minimum 5,000 cells from three replicate experiments; distributions corresponding to progeny strains represent many more cells (70,000 – 200,000 depending on whether the trait was measured in unbudded, small-budded, or large-budded cells). (D) The broad sense heritability for each of the 155 morphological features for which QTL were detected. Heritability is low because cell morphology varies across the cell cycle, and so the amount of non-genetic phenotypic variation is high.




Figure S2: Total numbers of cells imaged per each of 374 progeny strains. Each point represents, for one of the 374 progeny strains, the number of unbudded, small-budded, or large-budded cells for which images passed filtering (see *Methods*). Each boxplot shows the median (center line), interquartile range (IQR) (upper and lower hinges), and highest value within $1.5 \times$ IQR (whiskers).




Figure S3: Comparison of correlation estimates obtained from correlation partitioning with those obtained from a mixed-effect linear model. Each point represents one of 350 randomly sampled trait pairs of the 5645 total. Vertical axes display trait correlations estimated using the correlation-partitioning approach; horizontal axes display trait correlations estimated using a mixed-effect linear model that specifies the variance-covariance structure of the experimental design.

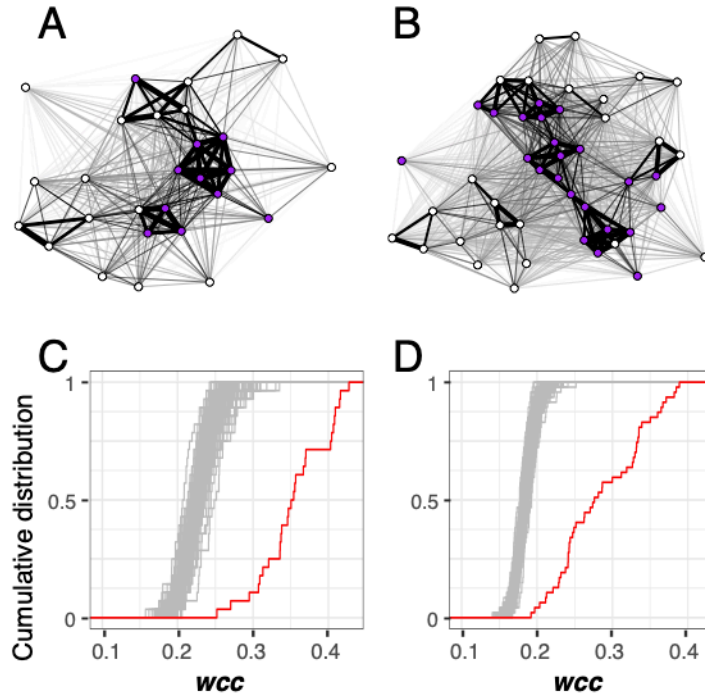


Figure S4: Single-cell morphological traits have higher weighted clustering coefficients (wcc) than expected given the distribution of r_w . (A – B) Force-directed networks visualizing how pairs of morphological features correlate across clones in unbudded (panel A) and small-budded (panel B) cells. Each node represents a single-cell morphological trait. The thickness of the line connecting each pair of nodes is proportional to r_w . Node position in the network is determined using the Fruchterman-Reingold algorithm. Purple nodes correspond to traits influenced by a QTL on chromosome 13 containing the *HOF1* gene. (C – D) Cumulative distributions of weighted clustering coefficients (wcc) in a network created using measured values of r_w (red line) or in 100 permuted networks (grey lines) for traits corresponding to unbudded (panel C) or small-budded (panel D) cells. Permutations were performed by sampling r_w , without replacement, and reassigning each value to a random pair of traits.

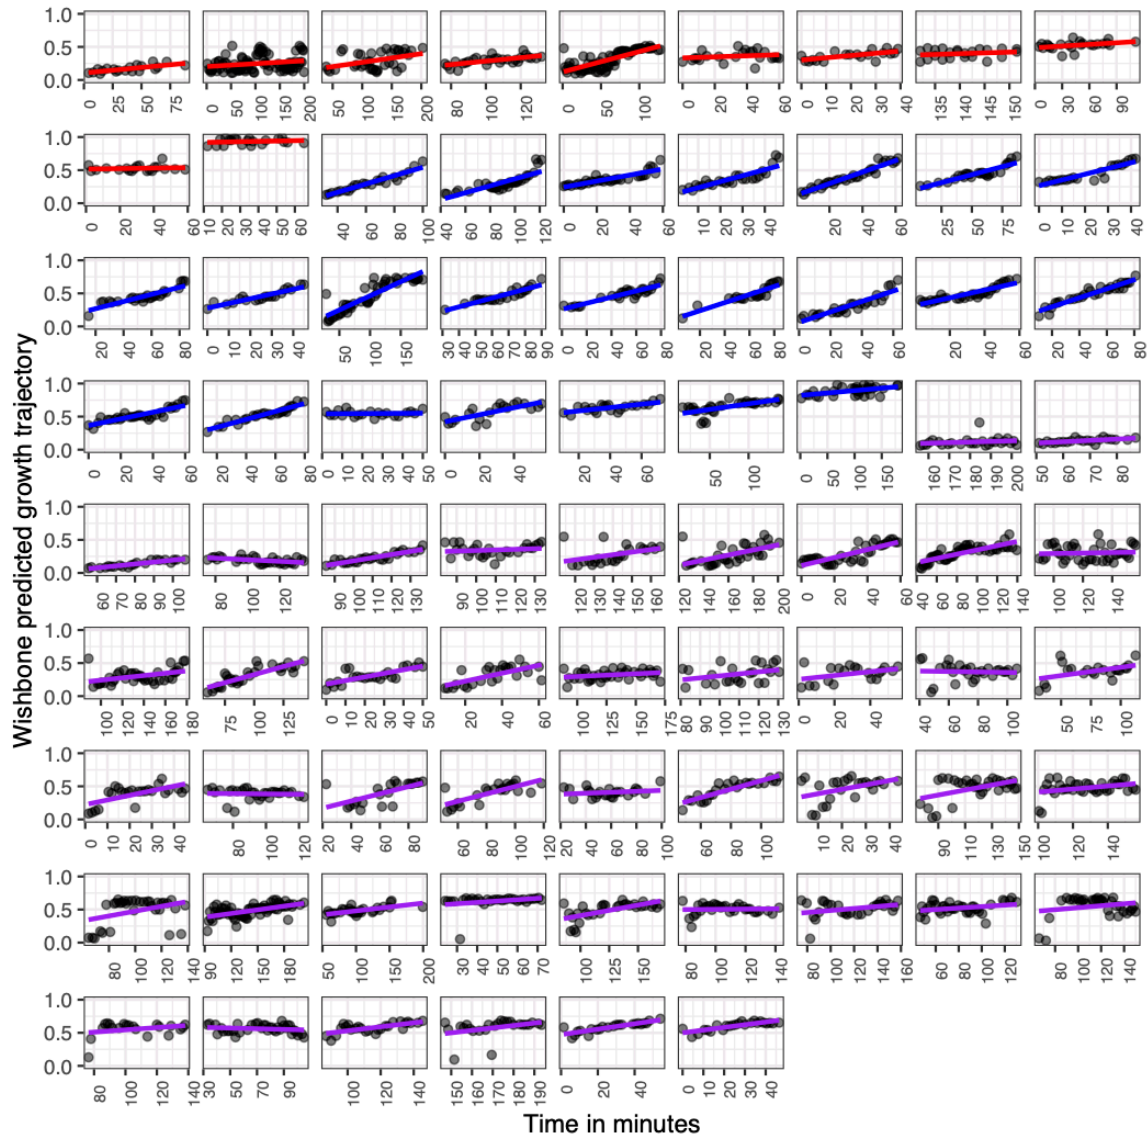


Figure S5: Wishbone recapitulates time series data obtained in live images of 78 cells undergoing exponential growth. Each point represents a cell image. Horizontal axes display the minute that image was captured during a three-hour window of exponential growth. Vertical axes display Wishbone's prediction of how far that cell image has passed through the cell cycle. Linear regression lines are calculated with the "lm" method in the R package ggplot2 (99), and are colored red for images corresponding to unbudded cells, blue for small-budded cells and purple for large-budded cells. Plots are organized by cell type and then from earliest to latest average predicted progress through cell division.

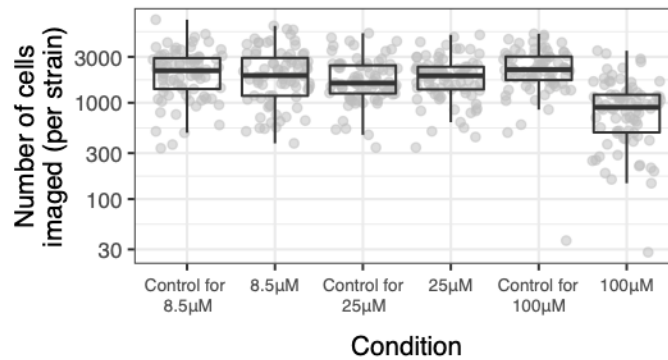


Figure S6: Total numbers of cells imaged per strain in varying concentrations of GdA. Each point represents, for one of the strains, the number of unbudded, small-budded, or large-budded cells for which images passed filtering (see *Methods*). Each boxplot shows the median (center line), interquartile range (IQR) (upper and lower hinges), and highest value within $1.5 \times$ IQR (whiskers).

S1 Table. Chromosomal locations, effects sizes and phenotypes affected by quantitative trait loci described in this study.

S2 Table. Impact of gene swaps on single-cell morphological traits including the corrected phenotypic difference between strains for each phenotype, and its standard deviation and standard error across replicate experiments.